

Hyperparameter Tuning Cookbook

A guide for scikit-learn, PyTorch, river, and spotPython

Thomas Bartz-Beielstein

Jun 19, 2023

Table of contents

Preface	3
Citation	3
1 Introduction: Hyperparameter Tuning	5
1.1 The Hyperparameter Tuning Software SPOT	6
1.2 Spot as an Optimizer	7
1.3 Example: <code>Spot</code> and the Sphere Function	8
1.3.1 The Objective Function: Sphere	8
1.4 Spot Parameters: <code>fun_evals</code> , <code>init_size</code> and <code>show_models</code>	10
1.5 Print the Results	12
1.6 Show the Progress	12
2 Multi-dimensional Functions	14
2.1 Example: <code>Spot</code> and the 3-dim Sphere Function	14
2.1.1 The Objective Function: 3-dim Sphere	14
2.1.2 Results	15
2.1.3 A Contour Plot	16
2.2 Conclusion	18
2.3 Exercises	18
2.3.1 The Three Dimensional <code>fun_cubed</code>	18
2.3.2 The Ten Dimensional <code>fun_wing_wt</code>	19
2.3.3 The Three Dimensional <code>fun_runge</code>	19
2.3.4 The Three Dimensional <code>fun_linear</code>	19
3 Isotropic and Anisotropic Kriging	20
3.1 Example: Isotropic <code>Spot</code> Surrogate and the 2-dim Sphere Function	20
3.1.1 The Objective Function: 2-dim Sphere	20
3.1.2 Results	21
3.2 Example With Anisotropic Kriging	21
3.2.1 Taking a Look at the <code>theta</code> Values	22
3.3 Exercises	23
3.3.1 <code>fun_branin</code>	23
3.3.2 <code>fun_sin_cos</code>	24
3.3.3 <code>fun_runge</code>	24
3.3.4 <code>fun_wingwt</code>	24

4	Using sklearn Surrogates in spotPython	25
4.1	Example: Branin Function with spotPython's Internal Kriging Surrogate . . .	25
4.1.1	The Objective Function Branin	25
4.1.2	Running the surrogate model based optimizer Spot:	26
4.1.3	Print the Results	26
4.1.4	Show the Progress and the Surrogate	26
4.2	Example: Using Surrogates From scikit-learn	27
4.2.1	GaussianProcessRegressor as a Surrogate	28
4.3	Example: One-dimensional Sphere Function With spotPython's Kriging	30
4.3.1	Results	35
4.4	Example: Sklearn Model GaussianProcess	36
4.5	Exercises	42
4.5.1	DecisionTreeRegressor	42
4.5.2	RandomForestRegressor	42
4.5.3	linear_model.LinearRegression	42
4.5.4	linear_model.Ridge	43
4.6	Exercise 2	43
5	Sequential Parameter Optimization: Using scipy Optimizers	44
5.1	The Objective Function Branin	44
5.2	The Optimizer	45
5.3	Print the Results	46
5.4	Show the Progress	46
5.5	Exercises	47
5.5.1	dual_annealing	47
5.5.2	direct	47
5.5.3	shgo	48
5.5.4	basinhopping	48
5.5.5	Performance Comparison	48
6	Sequential Parameter Optimization: Gaussian Process Models	49
6.1	Gaussian Processes Regression: Basic Introductory scikit-learn Example . .	49
6.1.1	Train and Test Data	50
6.1.2	Building the Surrogate With Sklearn	50
6.1.3	Plotting the SklearnModel	50
6.1.4	The spotPython Version	51
6.1.5	Visualizing the Differences Between the spotPython and the sklearn Model Fits	52
6.2	Exercises	53
6.2.1	Schonlau Example Function	53
6.2.2	Forrester Example Function	53
6.2.3	fun_runge Function (1-dim)	54
6.2.4	fun_cubed (1-dim)	55

6.2.5	The Effect of Noise	55
7	Expected Improvement	57
7.1	Example: <code>Spot</code> and the 1-dim Sphere Function	57
7.1.1	The Objective Function: 1-dim Sphere	57
7.1.2	Results	58
7.2	Same, but with EI as <code>infill_criterion</code>	58
7.3	Non-isotropic Kriging	59
7.4	Using <code>sklearn</code> Surrogates	61
7.4.1	The <code>spot</code> Loop	61
7.4.2	<code>spot</code> : The Initial Model	63
7.4.3	Init: Build Initial Design	63
7.4.4	Evaluate	66
7.4.5	Build Surrogate	66
7.4.6	A Simple Predictor	66
7.5	Gaussian Processes regression: basic introductory example	66
7.6	The Surrogate: Using scikit-learn models	69
7.7	Additional Examples	71
7.7.1	Optimize on Surrogate	75
7.7.2	Evaluate on Real Objective	75
7.7.3	Impute / Infill new Points	75
7.8	Tests	75
7.9	EI: The Famous Schonlau Example	76
7.10	EI: The Forrester Example	78
7.11	Noise	81
7.12	Cubic Function	84
7.13	Factors	90
8	Hyperparameter Tuning and Noise	92
8.1	Example: <code>Spot</code> and the Noisy Sphere Function	92
8.1.1	The Objective Function: Noisy Sphere	92
8.2	Print the Results	96
8.3	Noise and Surrogates: The Nugget Effect	96
8.3.1	The Noisy Sphere	96
8.4	Exercises	99
8.4.1	Noisy <code>fun_cubed</code>	99
8.4.2	<code>fun_runge</code>	100
8.4.3	<code>fun_forrester</code>	100
8.4.4	<code>fun_xsin</code>	100
9	Handling Noise: Optimal Computational Budget Allocation in <code>Spot</code>	101
9.1	Example: <code>Spot</code> , OCBA, and the Noisy Sphere Function	101
9.1.1	The Objective Function: Noisy Sphere	101

9.2	Print the Results	111
9.3	Noise and Surrogates: The Nugget Effect	112
9.3.1	The Noisy Sphere	112
9.4	Exercises	115
9.4.1	Noisy <code>fun_cubed</code>	115
9.4.2	<code>fun_runge</code>	115
9.4.3	<code>fun_forrester</code>	115
9.4.4	<code>fun_xsin</code>	116
10	HPT: sklearn SVC on Moons Data	117
10.1	Step 1: Setup	117
10.2	Step 2: Initialization of the Empty <code>fun_control</code> Dictionary	118
10.3	Step 3: SKlearn Load Data (Classification)	118
10.4	Step 4: Specification of the Preprocessing Model	120
10.5	Step 5: Select Model (<code>algorithm</code>) and <code>core_model_hyper_dict</code>	121
10.6	Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	123
10.6.1	Modify hyperparameter of type numeric and integer (boolean)	124
10.6.2	Modify hyperparameter of type factor	124
10.6.3	Optimizers	124
10.7	Step 7: Selection of the Objective (Loss) Function	125
10.7.1	Predict Classes or Class Probabilities	125
10.8	Step 8: Calling the SPOT Function	125
10.8.1	Preparing the SPOT Call	125
10.8.2	The Objective Function	126
10.8.3	Run the <code>Spot</code> Optimizer	126
10.8.4	Starting the Hyperparameter Tuning	127
10.9	Step 9: Results	128
10.9.1	Show variable importance	130
10.9.2	Get Default Hyperparameters	131
10.9.3	Get SPOT Results	131
10.9.4	Plot: Compare Predictions	132
10.9.5	Detailed Hyperparameter Plots	135
10.9.6	Parallel Coordinates Plot	136
10.9.7	Plot all Combinations of Hyperparameters	136
11	HPT: PyTorch With fashionMNIST	137
11.1	Step 1: Setup	137
11.2	Step 2: Initialization of the Empty <code>fun_control</code> Dictionary	139
11.3	Step 3: PyTorch Data Loading	139
11.3.1	Load fashionMNIST Data	139
11.4	Step 4: Specification of the Preprocessing Model	140

11.5	Step 5: Select Model (<code>algorithm</code>) and <code>core_model_hyper_dict</code>	140
11.5.1	The Search Space	141
11.5.2	Configuring the Search Space With <code>spotPython</code>	141
11.6	Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	143
11.6.1	Modify hyperparameter of type numeric and integer (boolean)	143
11.6.2	Modify hyperparameter of type factor	144
11.6.3	Optimizers	144
11.7	Step 7: Selection of the Objective (Loss) Function	144
11.7.1	Evaluation	144
11.7.2	Metric	145
11.8	Step 8: Calling the SPOT Function	145
11.8.1	Preparing the SPOT Call	145
11.8.2	The Objective Function <code>fun_torch</code>	146
11.8.3	Starting the Hyperparameter Tuning	146
11.9	Step 9: Tensorboard	151
11.10	Step 10: Results	151
11.10.1	Show variable importance	153
11.10.2	Get the Tuned Architecture (SPOT Results)	153
11.10.3	Get Default Hyperparameters	154
11.10.4	Evaluation of the Default and the Tuned Architectures	154
11.10.5	Detailed Hyperparameter Plots	157
11.10.6	Parallel Coordinates Plot	159
11.10.7	Plot all Combinations of Hyperparameters	159

12 HPT: PyTorch With cifar10 Data 160

12.1	Step 1: Setup	160
12.2	Step 2: Initialization of the Empty <code>fun_control</code> Dictionary	162
12.3	Step 3: PyTorch Data Loading	162
12.3.1	Load Data Cifar10 Data	162
12.4	Step 4: Specification of the Preprocessing Model	163
12.5	Step 5: Select Model (<code>algorithm</code>) and <code>core_model_hyper_dict</code>	163
12.5.1	Implementing a Configurable Neural Network With <code>spotPython</code>	163
12.5.2	The Search Space	164
12.5.3	Configuring the Search Space With <code>spotPython</code>	164
12.6	Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	165
12.6.1	Step 5: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	166
12.6.2	Modify hyperparameter of type factor	166
12.6.3	Optimizers	166
12.7	Step 7: Selection of the Objective (Loss) Function	167
12.7.1	Evaluation	167

12.7.2	Metric	167
12.8	Step 8: Calling the SPOT Function	168
12.8.1	Preparing the SPOT Call	168
12.8.2	The Objective Function <code>fun_torch</code>	168
12.8.3	Starting the Hyperparameter Tuning	169
12.9	Step 9: Tensorboard	173
12.10	Step 10: Results	173
12.10.1	Show variable importance	175
12.10.2	Get the Tuned Architecture (SPOT Results)	175
12.10.3	Evaluation of the Tuned Architecture	176
12.10.4	Cross-validated Evaluations	177
12.10.5	Detailed Hyperparameter Plots	178
12.10.6	Parallel Coordinates Plot	180
12.10.7	Plot all Combinations of Hyperparameters	180
13	HPT: River	181
13.1	Step 1: Setup	181
13.1.1	<code>river</code> Hyperparameter Tuning: HATR with Friedman Drift Data	181
13.2	Step 2: Initialization of the <code>fun_control</code> Dictionary	182
13.3	Step 3: Load the Friedman Drift Data	182
13.4	Step 4: Specification of the Preprocessing Model	183
13.5	Step 5: Select <code>algorithm</code> and <code>core_model_hyper_dict</code>	184
13.6	Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	187
13.6.1	Modify hyperparameter of type factor	187
13.6.2	Modify hyperparameter of type numeric and integer (boolean)	187
13.7	Step 7: Selection of the Objective (Loss) Function	187
13.8	Step 8: Calling the SPOT Function	188
13.8.1	Prepare the SPOT Parameters	188
13.8.2	Run the <code>Spot</code> Optimizer	189
13.9	Step 9: Results	190
13.9.1	Show variable importance	192
13.9.2	Build and Evaluate HTR Model with Tuned Hyperparameters	192
13.9.3	The Large Data Set (<code>k=0.2</code>)	193
13.9.4	Get Default Hyperparameters	194
13.9.5	Get SPOT Results	197
13.9.6	Visualize Regression Trees	201
13.9.7	Spot Model	201
13.9.8	Detailed Hyperparameter Plots	203
13.9.9	Parallel Coordinates Plots	203
13.9.10	Plot all Combinations of Hyperparameters	203

14 HPT: PyTorch With spotPython and Ray Tune on CIFAR10	205
14.1 Step 1: Setup	206
14.2 Step 2: Initialization of the <code>fun_control</code> Dictionary	207
14.3 Step 3: PyTorch Data Loading	207
14.4 Step 4: Specification of the Preprocessing Model	208
14.5 Step 5: Select Model (<code>algorithm</code>) and <code>core_model_hyper_dict</code>	209
14.5.1 The <code>Net_Core</code> class	211
14.5.2 Comparison of the Approach Described in the PyTorch Tutorial With <code>spotPython</code>	211
14.5.3 The Search Space: Hyperparameters	212
14.5.4 Configuring the Search Space With Ray Tune	212
14.5.5 Configuring the Search Space With <code>spotPython</code>	213
14.6 Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	215
14.6.1 Optimizers	216
14.7 Step 7: Selection of the Objective (Loss) Function	218
14.7.1 Evaluation: Data Splitting	218
14.7.2 Hold-out Data Split	218
14.7.3 Cross-Validation	219
14.7.4 Overview of the Evaluation Settings	219
14.7.5 Evaluation: Loss Functions and Metrics	221
14.8 Step 8: Calling the SPOT Function	222
14.8.1 Preparing the SPOT Call	222
14.8.2 The Objective Function <code>fun_torch</code>	223
14.8.3 Using Default Hyperparameters or Results from Previous Runs	223
14.8.4 Starting the Hyperparameter Tuning	223
14.9 Step 9: Tensorboard	234
14.9.1 Tensorboard: Start Tensorboard	235
14.9.2 Saving the State of the Notebook	236
14.10 Step 10: Results	236
14.10.1 Get the Tuned Architecture (SPOT Results)	238
14.10.2 Get Default Hyperparameters	239
14.10.3 Evaluation of the Default Architecture	239
14.10.4 Evaluation of the Tuned Architecture	241
14.10.5 Detailed Hyperparameter Plots	243
14.11 Summary and Outlook	244
14.12 Appendix	244
14.12.1 Sample Output From Ray Tune’s Run	244
15 HPT: sklearn RandomForestClassifier VBDP Data	246
15.1 Step 1: Setup	246
15.2 Step 2: Initialization of the Empty <code>fun_control</code> Dictionary	247

15.3	Step 3: PyTorch Data Loading	248
15.3.1	Load Data: Classification VBDP	248
15.3.2	Holdout Train and Test Data	248
15.4	Step 4: Specification of the Preprocessing Model	249
15.5	Step 5: Select Model (<code>algorithm</code>) and <code>core_model_hyper_dict</code>	250
15.6	Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	252
15.6.1	Modify hyperparameter of type numeric and integer (boolean)	252
15.6.2	Modify hyperparameter of type factor	252
15.6.3	Optimizers	253
15.6.4	Selection of the Objective: Metric and Loss Functions	253
15.7	Step 7: Selection of the Objective (Loss) Function	253
15.7.1	Metric Function	253
15.7.2	Evaluation on Hold-out Data	254
15.7.3	OOB Score	255
15.8	Step 8: Calling the SPOT Function	256
15.8.1	Preparing the SPOT Call	256
15.8.2	The Objective Function	256
15.8.3	Run the <code>Spot</code> Optimizer	257
15.9	Step 9: Tensorboard	259
15.10	Step 10: Results	260
15.10.1	Show variable importance	261
15.10.2	Get Default Hyperparameters	261
15.10.3	Get SPOT Results	262
15.10.4	Evaluate SPOT Results	263
15.10.5	Handling Non-deterministic Results	264
15.10.6	Evaluation of the Default Hyperparameters	264
15.10.7	Plot: Compare Predictions	265
15.10.8	Cross-validated Evaluations	267
15.10.9	Detailed Hyperparameter Plots	268
15.10.10	Parallel Coordinates Plot	272
15.10.11	Plot all Combinations of Hyperparameters	272
16	HPT: sklearn XGB Classifier VBDP Data	273
16.1	Step 1: Setup	273
16.2	Step 2: Initialization of the Empty <code>fun_control</code> Dictionary	274
16.3	Step 3: PyTorch Data Loading	275
16.3.1	1. Load Data: Classification VBDP	275
16.3.2	Holdout Train and Test Data	275
16.4	Step 4: Specification of the Preprocessing Model	276
16.5	Step 5: Select Model (<code>algorithm</code>) and <code>core_model_hyper_dict</code>	277

16.6	Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	279
16.6.1	Modify hyperparameter of type numeric and integer (boolean)	279
16.6.2	Modify hyperparameter of type factor	279
16.6.3	Optimizers	280
16.7	Step 7: Selection of the Objective (Loss) Function	280
16.7.1	Evaluation	280
16.7.2	Selection of the Objective: Metric and Loss Functions	280
16.7.3	Loss Function	280
16.7.4	Metric Function	280
16.7.5	Evaluation on Hold-out Data	282
16.8	Step 8: Calling the SPOT Function	282
16.8.1	Preparing the SPOT Call	282
16.8.2	The Objective Function	283
16.8.3	Run the Spot Optimizer	283
16.9	Step 9: Tensorboard	285
16.10	Step 10: Results	285
16.10.1	Show variable importance	286
16.10.2	Get Default Hyperparameters	287
16.10.3	Get SPOT Results	287
16.10.4	Evaluate SPOT Results	288
16.10.5	Handling Non-deterministic Results	289
16.10.6	Evaluation of the Default Hyperparameters	290
16.10.7	Plot: Compare Predictions	290
16.10.8	Cross-validated Evaluations	292
16.10.9	Detailed Hyperparameter Plots	293
16.10.10	Parallel Coordinates Plot	297
16.10.11	Plot all Combinations of Hyperparameters	297
17	HPT: sklearn SVC VBDP Data	298
17.1	Step 1: Setup	298
17.2	Step 2: Initialization of the Empty <code>fun_control</code> Dictionary	299
17.3	Step 3: PyTorch Data Loading	300
17.3.1	1. Load Data: Classification VBDP	300
17.3.2	Holdout Train and Test Data	300
17.4	Step 4: Specification of the Preprocessing Model	301
17.5	Step 5: Select Model (<code>algorithm</code>) and <code>core_model_hyper_dict</code>	302
17.6	Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	304
17.6.1	Modify hyperparameter of type numeric and integer (boolean)	304
17.6.2	Modify hyperparameter of type factor	304
17.6.3	Optimizers	304
17.6.4	Selection of the Objective: Metric and Loss Functions	305

17.7	Step 7: Selection of the Objective (Loss) Function	305
17.7.1	Metric Function	305
17.7.2	Evaluation on Hold-out Data	306
17.8	Step 8: Calling the SPOT Function	307
17.8.1	Preparing the SPOT Call	307
17.8.2	The Objective Function	308
17.8.3	Run the Spot Optimizer	308
17.9	Step 9: Tensorboard	312
17.10	Step 10: Results	312
17.10.1	Show variable importance	313
17.10.2	Get Default Hyperparameters	314
17.10.3	Get SPOT Results	315
17.10.4	Evaluate SPOT Results	315
17.10.5	Handling Non-deterministic Results	316
17.10.6	Evaluation of the Default Hyperparameters	317
17.10.7	Plot: Compare Predictions	317
17.10.8	Cross-validated Evaluations	319
17.10.9	Detailed Hyperparameter Plots	320
17.10.10	Parallel Coordinates Plot	321
17.10.11	Plot all Combinations of Hyperparameters	321

18 HPT: sklearn KNN Classifier VBDP Data 323

18.1	Step 1: Setup	323
18.2	Step 2: Initialization of the Empty <code>fun_control</code> Dictionary	324
18.2.1	Load Data: Classification VBDP	324
18.2.2	Holdout Train and Test Data	325
18.3	Step 4: Specification of the Preprocessing Model	326
18.4	Step 5: Select Model (<code>algorithm</code>) and <code>core_model_hyper_dict</code>	327
18.5	Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	328
18.5.1	Modify hyperparameter of type numeric and integer (boolean)	328
18.5.2	Modify hyperparameter of type factor	329
18.5.3	Optimizers	329
18.5.4	Selection of the Objective: Metric and Loss Functions	329
18.6	Step 7: Selection of the Objective (Loss) Function	329
18.6.1	Metric Function	330
18.6.2	Evaluation on Hold-out Data	331
18.7	Step 8: Calling the SPOT Function	331
18.7.1	Preparing the SPOT Call	331
18.7.2	The Objective Function	332
18.7.3	Run the Spot Optimizer	332
18.8	Step 9: Tensorboard	336

18.9	Step 10: Results	337
18.9.1	Show variable importance	337
18.9.2	Get Default Hyperparameters	338
18.9.3	Get SPOT Results	339
18.9.4	Evaluate SPOT Results	339
18.9.5	Handling Non-deterministic Results	340
18.9.6	Evaluation of the Default Hyperparameters	341
18.9.7	Plot: Compare Predictions	341
18.9.8	Cross-validated Evaluations	343
18.9.9	Detailed Hyperparameter Plots	344
18.9.10	Parallel Coordinates Plot	345
18.9.11	Plot all Combinations of Hyperparameters	345
19	HPT PyTorch: Regression	346
19.1	Step 1: Setup	346
19.2	Step 2: Initialization of the <code>fun_control</code> Dictionary	348
19.3	Step 3: PyTorch Data Loading	348
19.4	Step 4: Specification of the Preprocessing Model	350
19.5	Step 5: Select Model (<code>algorithm</code>) and <code>core_model_hyper_dict</code>	350
19.5.1	Implementing a Configurable Neural Network With <code>spotPython</code>	350
19.5.2	The Search Space	352
19.5.3	Configuring the Search Space With <code>spotPython</code>	352
19.6	Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	354
19.6.1	Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	354
19.6.2	Optimizers	354
19.7	Step 7: Selection of the Objective (Loss) Function	355
19.7.1	Evaluation	355
19.7.2	Loss Functions and Metrics	355
19.7.3	Metric	355
19.8	Step 8: Calling the SPOT Function	355
19.8.1	Preparing the SPOT Call	355
19.8.2	The Objective Function <code>fun_torch</code>	356
19.8.3	Starting the Hyperparameter Tuning	356
19.9	Step 9: Tensorboard	409
19.10	Step 10: Results	410
19.10.1	Get the Tuned Architecture (SPOT Results)	411
19.10.2	Evaluation of the Tuned Architecture	412
19.10.3	Cross-validated Evaluations	413
19.10.4	Detailed Hyperparameter Plots	431
19.10.5	Parallel Coordinates Plot	436
19.11	Summary and Outlook	436

20 HPT: PyTorch With VBDP	438
20.1 Step 1: Setup	439
20.2 Step 2: Initialization of the <code>fun_control</code> Dictionary	440
20.3 Step 3: PyTorch Data Loading	440
20.3.1 1. Load VBDP Data	440
20.3.2 Check content of the target column	442
20.4 Step 4: Specification of the Preprocessing Model	443
20.5 Step 5: Select <code>algorithm</code> and <code>core_model_hyper_dict</code>	443
20.5.1 Implementing a Configurable Neural Network With <code>spotPython</code>	443
20.5.2 Add the NN Model to the <code>fun_control</code> Dictionary	443
20.6 Step 6: Modify <code>hyper_dict</code> Hyperparameters for the Selected Algorithm aka <code>core_model</code>	445
20.6.1 Optimizers	446
20.7 Step 7: Selection of the Objective (Loss) Function	446
20.7.1 Evaluation	446
20.7.2 Loss Functions and Metrics	446
20.7.3 Metric	447
20.8 Step 8: Calling the SPOT Function	447
20.8.1 Preparing the SPOT Call	447
20.8.2 The Objective Function <code>fun_torch</code>	448
20.8.3 Starting the Hyperparameter Tuning	449
20.9 Step 9: Tensorboard	456
20.10 Step 10: Results	457
20.10.1 Get the Tuned Architecture	458
20.10.2 Evaluation of the Tuned Architecture	459
20.10.3 Cross-validated Evaluations	460
20.10.4 Detailed Hyperparameter Plots	463
20.10.5 Parallel Coordinates Plot	465
20.10.6 Plot all Combinations of Hyperparameters	465
21 Documentation of the Sequential Parameter Optimization	466
21.1 Example: <code>spot</code>	466
21.1.1 The Objective Function	466
21.1.2 External Parameters	468
21.2 The <code>fun_control</code> Dictionary	471
21.3 The <code>design_control</code> Dictionary	471
21.4 The <code>surrogate_control</code> Dictionary	472
21.5 The <code>optimizer_control</code> Dictionary	472
21.6 Run	473
21.7 Print the Results	475
21.8 Show the Progress	475
21.9 Visualize the Surrogate	475
21.10 Init: Build Initial Design	476

21.11	Replicability	477
21.12	Surrogates	478
21.12.1	A Simple Predictor	478
21.13	Demo/Test: Objective Function Fails	478
21.14	PyTorch: Detailed Description of the Data Splitting	480
21.14.1	Description of the "train_hold_out" Setting	480

References **491**

Preface

The goal of hyperparameter tuning (or hyperparameter optimization) is to optimize the hyperparameters to improve the performance of the machine or deep learning model.

spotPython (“Sequential Parameter Optimization Toolbox in Python”) is the Python version of the well-known hyperparameter tuner SPOT, which has been developed in the R programming environment for statistical analysis for over a decade. The related open-access book is available here: [Hyperparameter Tuning for Machine and Deep Learning with R—A Practical Guide](#).

[scikit-learn](#) is a Python module for machine learning built on top of SciPy and is distributed under the 3-Clause BSD license. The project was started in 2007 by David Cournapeau as a Google Summer of Code project, and since then many volunteers have contributed.

[PyTorch](#) is an optimized tensor library for deep learning using GPUs and CPUs.

[River](#) is a Python library for online machine learning. It is designed to be used in real-world environments, where not all data is available at once, but streaming in.

! Important: This book is still under development.

Citation

If this document has been useful to you and you wish to cite it in a scientific publication, please refer to the following paper, which can be found on arXiv: <https://arxiv.org/abs/2305.11930>.

```
@ARTICLE{bart23earxiv,  
  author = {{Bartz-Beielstein}, Thomas},  
  title = "{PyTorch Hyperparameter Tuning -- A Tutorial for spotPython}",  
  journal = {arXiv e-prints},  
  keywords = {Computer Science - Machine Learning, Computer Science - Artificial Intell  
  year = 2023,  
  month = may,  
  eid = {arXiv:2305.11930},
```

```
    pages = {arXiv:2305.11930},
    doi = {10.48550/arXiv.2305.11930},
archivePrefix = {arXiv},
  eprint = {2305.11930},
primaryClass = {cs.LG},
  adsurl = {https://ui.adsabs.harvard.edu/abs/2023arXiv230511930B},
  adsnote = {Provided by the SAO/NASA Astrophysics Data System}
}
```

1 Introduction: Hyperparameter Tuning

Hyperparameter tuning is an important, but often difficult and computationally intensive task. Changing the architecture of a neural network or the learning rate of an optimizer can have a significant impact on the performance.

The goal of hyperparameter tuning is to optimize the hyperparameters in a way that improves the performance of the machine learning or deep learning model. The simplest, but also most computationally expensive, approach uses manual search (or trial-and-error (Meignan et al. 2015)). Commonly encountered is simple random search, i.e., random and repeated selection of hyperparameters for evaluation, and lattice search (“grid search”). In addition, methods that perform directed search and other model-free algorithms, i.e., algorithms that do not explicitly rely on a model, e.g., evolution strategies (Bartz-Beielstein et al. 2014) or pattern search (Lewis, Torczon, and Trosset 2000) play an important role. Also, “hyperband”, i.e., a multi-armed bandit strategy that dynamically allocates resources to a set of random configurations and uses successive bisections to stop configurations with poor performance (Li et al. 2016), is very common in hyperparameter tuning. The most sophisticated and efficient approaches are the Bayesian optimization and surrogate model based optimization methods, which are based on the optimization of cost functions determined by simulations or experiments.

We consider below a surrogate model based optimization-based hyperparameter tuning approach based on the Python version of the SPOT (“Sequential Parameter Optimization Toolbox”) (Bartz-Beielstein, Lasarczyk, and Preuss 2005), which is suitable for situations where only limited resources are available. This may be due to limited availability and cost of hardware, or due to the fact that confidential data may only be processed locally, e.g., due to legal requirements. Furthermore, in our approach, the understanding of algorithms is seen as a key tool for enabling transparency and explainability. This can be enabled, for example, by quantifying the contribution of machine learning and deep learning components (nodes, layers, split decisions, activation functions, etc.). Understanding the importance of hyperparameters and the interactions between multiple hyperparameters plays a major role in the interpretability and explainability of machine learning models. SPOT provides statistical tools for understanding hyperparameters and their interactions. Last but not least, it should be noted that the SPOT software code is available in the open source `spotPython` package on github¹, allowing replicability of the results. This tutorial describes the Python variant of SPOT, which is called

¹<https://github.com/sequential-parameter-optimization>

`spotPython`. The R implementation is described in Bartz et al. (2022). SPOT is an established open source software that has been maintained for more than 15 years (Bartz-Beielstein, Lasarczyk, and Preuss 2005) (Bartz et al. 2022).

This tutorial is structured as follows. The concept of the hyperparameter tuning software `spotPython` is described in Section 1.1. Chapter 14 describes the execution of the example from the tutorial “Hyperparameter Tuning with Ray Tune” (PyTorch 2023a). The integration of `spotPython` into the `PyTorch` training workflow is described in detail in the following sections. Section 14.1 describes the setup of the tuners. Section 14.3 describes the data loading. Section 14.5 describes the model to be tuned. The search space is introduced in Section 14.5.3. Optimizers are presented in Section 14.6.1. How to split the data in train, validation, and test sets is described in Section 14.7.1. The selection of the loss function and metrics is described in Section 14.7.5. Section 14.8.1 describes the preparation of the `spotPython` call. The objective function is described in Section 14.8.2. How to use results from previous runs and default hyperparameter configurations is described in Section 14.8.3. Starting the tuner is shown in Section 14.8.4. TensorBoard can be used to visualize the results as shown in Section 14.9. Results are discussed and explained in Section 14.10. Finally, Section 14.11 presents a summary and an outlook.

i Note

The corresponding `.ipynb` notebook (Bartz-Beielstein 2023) is updated regularly and reflects updates and changes in the `spotPython` package. It can be downloaded from https://github.com/sequential-parameter-optimization/spotPython/blob/main/notebooks/14_spot_ray_hpt_torch_cifar10.ipynb.

1.1 The Hyperparameter Tuning Software SPOT

Surrogate model based optimization methods are common approaches in simulation and optimization. SPOT was developed because there is a great need for sound statistical analysis of simulation and optimization algorithms. SPOT includes methods for tuning based on classical regression and analysis of variance techniques. It presents tree-based models such as classification and regression trees and random forests as well as Bayesian optimization (Gaussian process models, also known as Kriging). Combinations of different meta-modeling approaches are possible. SPOT comes with a sophisticated surrogate model based optimization method, that can handle discrete and continuous inputs. Furthermore, any model implemented in `scikit-learn` can be used out-of-the-box as a surrogate in `spotPython`.

SPOT implements key techniques such as exploratory fitness landscape analysis and sensitivity analysis. It can be used to understand the performance of various algorithms, while simultaneously giving insights into their algorithmic behavior. In addition, SPOT can be used as an

optimizer and for automatic and interactive tuning. Details on SPOT and its use in practice are given by Bartz et al. (2022).

A typical hyperparameter tuning process with `spotPython` consists of the following steps:

1. Loading the data (training and test datasets), see Section 14.3.
2. Specification of the preprocessing model, see Section 14.4. This model is called `prep_model` (“preparation” or pre-processing). The information required for the hyperparameter tuning is stored in the dictionary `fun_control`. Thus, the information needed for the execution of the hyperparameter tuning is available in a readable form.
3. Selection of the machine learning or deep learning model to be tuned, see Section 14.5. This is called the `core_model`. Once the `core_model` is defined, then the associated hyperparameters are stored in the `fun_control` dictionary. First, the hyperparameters of the `core_model` are initialized with the default values of the `core_model`. As default values we use the default values contained in the `spotPython` package for the algorithms of the `torch` package.
4. Modification of the default values for the hyperparameters used in `core_model`, see Section 14.6.0.1. This step is optional.
 1. numeric parameters are modified by changing the bounds.
 2. categorical parameters are modified by changing the categories (“levels”).
5. Selection of target function (loss function) for the optimizer, see Section 14.7.5.
6. Calling SPOT with the corresponding parameters, see Section 14.8.4. The results are stored in a dictionary and are available for further analysis.
7. Presentation, visualization and interpretation of the results, see Section 14.10.

1.2 Spot as an Optimizer

The `spot` loop consists of the following steps:

1. Init: Build initial design X
2. Evaluate initial design on real objective f : $y = f(X)$
3. Build surrogate: $S = S(X, y)$
4. Optimize on surrogate: $X_0 = \text{optimize}(S)$
5. Evaluate on real objective: $y_0 = f(X_0)$
6. Impute (Infill) new points: $X = X \cup X_0, y = y \cup y_0$.
7. Got 3.

Central Idea: Evaluation of the surrogate model S is much cheaper (or / and much faster) than running the real-world experiment f . We start with a small example.

1.3 Example: Spot and the Sphere Function

```
import numpy as np
from math import inf
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from scipy.optimize import shgo
from scipy.optimize import direct
from scipy.optimize import differential_evolution
import matplotlib.pyplot as plt
```

1.3.1 The Objective Function: Sphere

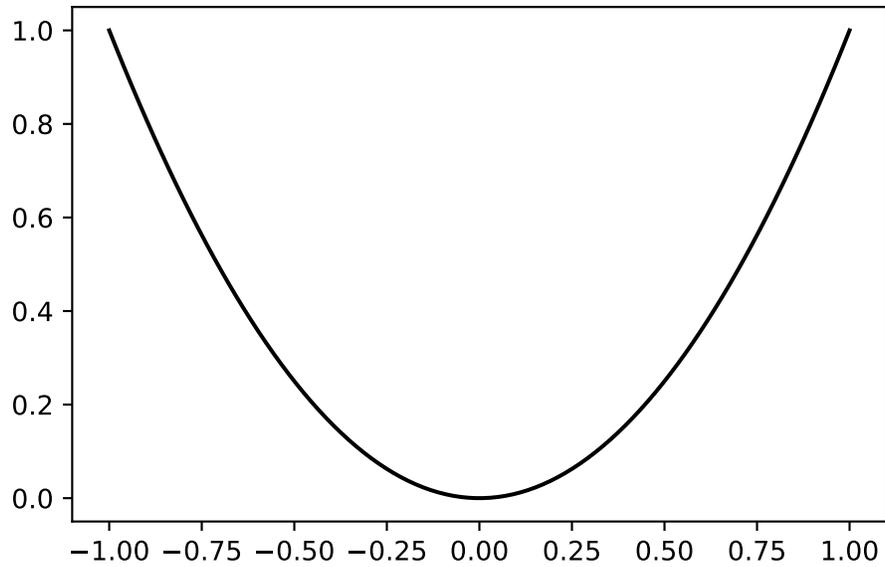
The `spotPython` package provides several classes of objective functions. We will use an analytical objective function, i.e., a function that can be described by a (closed) formula:

$$f(x) = x^2$$

```
fun = analytical().fun_sphere
```

We can apply the function `fun` to input values and plot the result:

```
x = np.linspace(-1,1,100).reshape(-1,1)
y = fun(x)
plt.figure()
plt.plot(x, y, "k")
plt.show()
```



```
spot_0 = spot.Spot(fun=fun,  
                  lower = np.array([-1]),  
                  upper = np.array([1]))
```

```
spot_0.run()
```

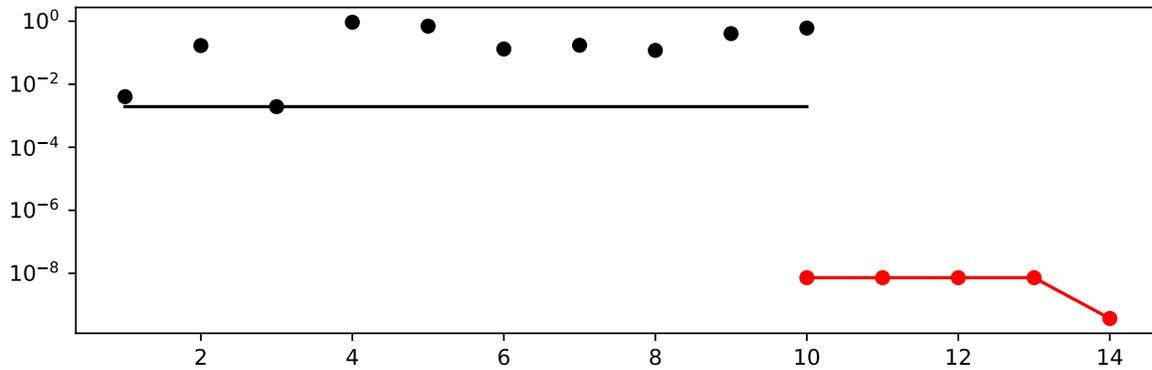
```
<spotPython.spot.spot.Spot at 0x107a5dc00>
```

```
spot_0.print_results()
```

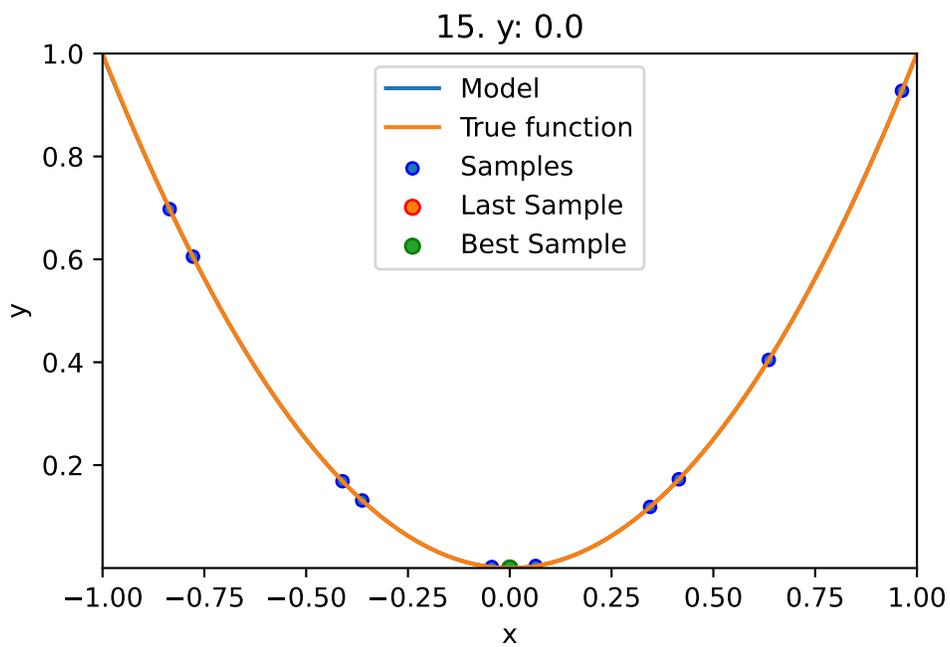
```
min y: 3.696886711914087e-10  
x0: 1.922728975158508e-05
```

```
[['x0', 1.922728975158508e-05]]
```

```
spot_0.plot_progress(log_y=True)
```



```
spot_0.plot_model()
```



1.4 Spot Parameters: fun_evals, init_size and show_models

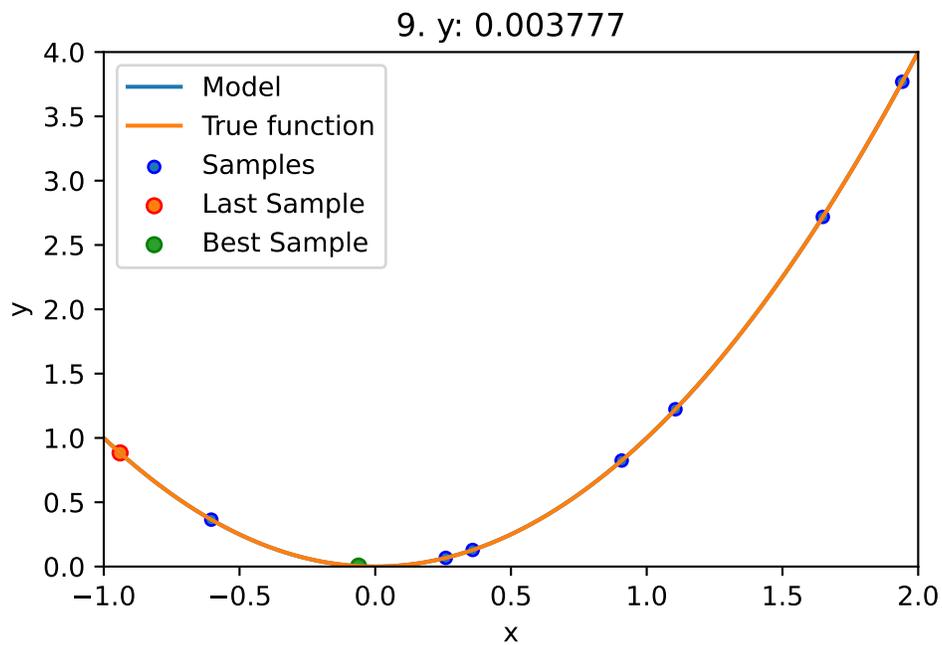
We will modify three parameters:

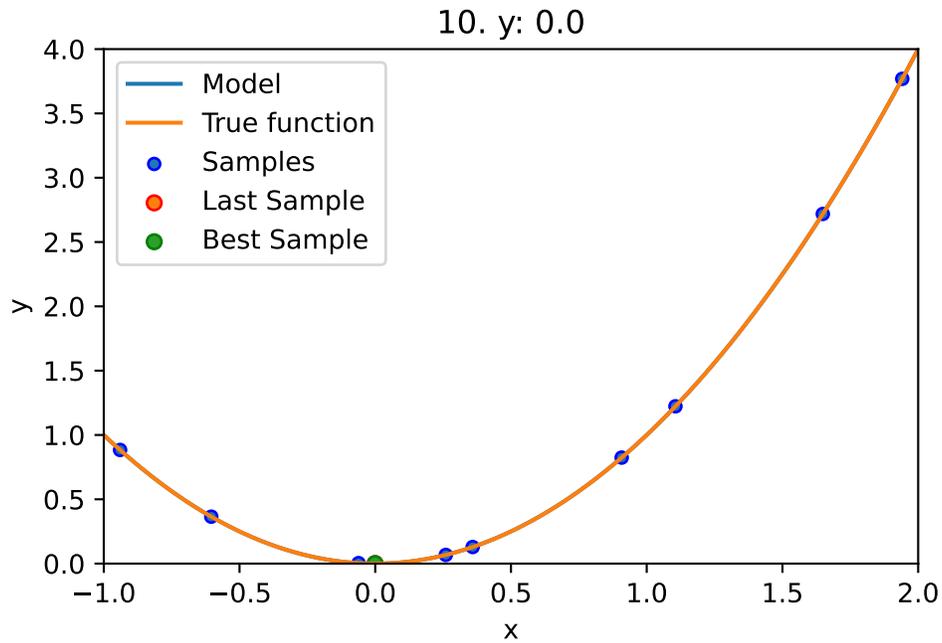
1. The number of function evaluations (`fun_evals`)
2. The size of the initial design (`init_size`)

3. The parameter `show_models`, which visualizes the search process for 1-dim functions.

The full list of the `Spot` parameters is shown in the Help System and in the notebook `spot_doc.ipynb`.

```
spot_1 = spot.Spot(fun=fun,  
                  lower = np.array([-1]),  
                  upper = np.array([2]),  
                  fun_evals= 10,  
                  seed=123,  
                  show_models=True,  
                  design_control={"init_size": 9})  
  
spot_1.run()
```





<spotPython.spot.spot.Spot at 0x2b00a39d0>

1.5 Print the Results

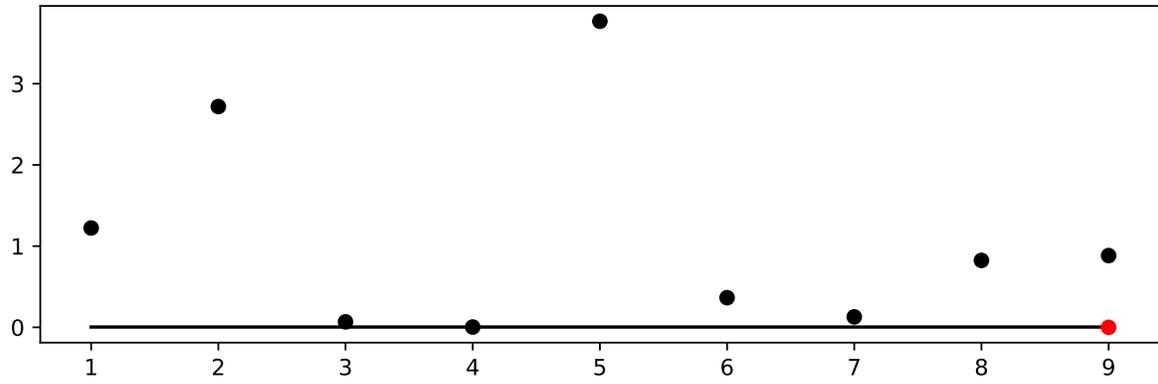
```
spot_1.print_results()
```

```
min y: 3.6779240309761575e-07
x0: -0.0006064589047063418
```

```
[['x0', -0.0006064589047063418]]
```

1.6 Show the Progress

```
spot_1.plot_progress()
```



2 Multi-dimensional Functions

This notebook illustrates how high-dimensional functions can be analyzed.

2.1 Example: Spot and the 3-dim Sphere Function

```
import numpy as np
from math import inf
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from scipy.optimize import shgo
from scipy.optimize import direct
from scipy.optimize import differential_evolution
import matplotlib.pyplot as plt
import pylab
from numpy import append, ndarray, multiply, isinf, linspace, meshgrid, ravel
from numpy import array
```

2.1.1 The Objective Function: 3-dim Sphere

- The spotPython package provides several classes of objective functions.
- We will use an analytical objective function, i.e., a function that can be described by a (closed) formula:

$$f(x) = \sum_i^n x_i^2$$

- Here we will use $n = 3$.

```
fun = analytical().fun_sphere
```

- The size of the lower bound vector determines the problem dimension.
- Here we will use `np.array([-1, -1, -1])`, i.e., a three-dim function.

- We will use three different `theta` values (one for each dimension), i.e., we set `surrogate_control={"n_theta": 3}`.

```
spot_3 = spot.Spot(fun=fun,
                  lower = -1.0*np.ones(3),
                  upper = np.ones(3),
                  var_name=["Pressure", "Temp", "Lambda"],
                  show_progress=True,
                  surrogate_control={"n_theta": 3})

spot_3.run()
```

```
spotPython tuning: 0.03443344056467332 [#####---] 73.33%
```

```
spotPython tuning: 0.03134865993507926 [#####--] 80.00%
```

```
spotPython tuning: 0.0009629342967936851 [#####-] 86.67%
```

```
spotPython tuning: 8.541951463966474e-05 [#####-] 93.33%
```

```
spotPython tuning: 6.285135731399678e-05 [#####] 100.00% Done...
```

```
<spotPython.spot.spot.Spot at 0x10875d9f0>
```

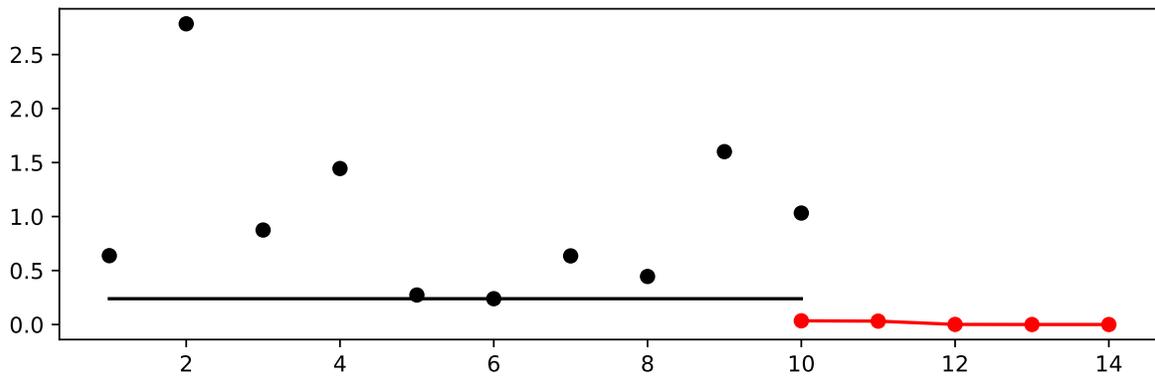
2.1.2 Results

```
spot_3.print_results()
```

```
min y: 6.285135731399678e-05
Pressure: 0.005236109709736696
Temp: 0.0019572552655686714
Lambda: 0.005621713639718905
```

```
[['Pressure', 0.005236109709736696],
 ['Temp', 0.0019572552655686714],
 ['Lambda', 0.005621713639718905]]
```

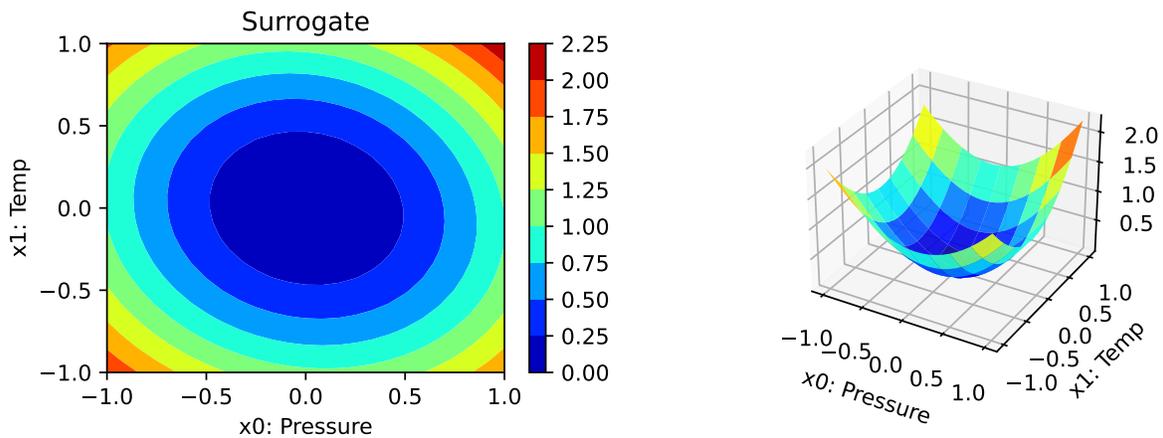
```
spot_3.plot_progress()
```



2.1.3 A Contour Plot


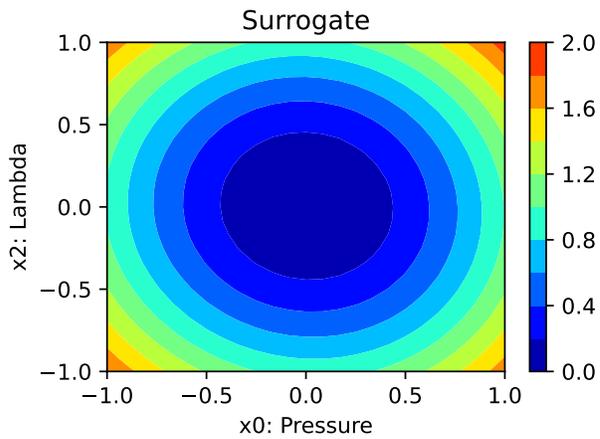
- We can select two dimensions, say $i = 0$ and $j = 1$, and generate a contour plot as follows.
 - Note: We have specified identical `min_z` and `max_z` values to generate comparable plots!

```
spot_3.plot_contour(i=0, j=1, min_z=0, max_z=2.25)
```





- In a similar manner, we can plot dimension $i = 0$ and $j = 2$:

```
spot_3.plot_contour(i=0, j=2, min_z=0, max_z=2.25)
```



- The final combination is $i = 1$ and $j = 2$:

```
spot_3.plot_contour(i=1, j=2, min_z=0, max_z=2.25)
```



- The three plots look very similar, because the `fun_sphere` is symmetric.
- This can also be seen from the variable importance:

```
spot_3.print_importance()
```

```
Pressure: 99.35185545837122
```

```
Temp: 99.99999999999999
```

Lambda: 94.31627052007231

```
[['Pressure', 99.35185545837122],  
 ['Temp', 99.99999999999999],  
 ['Lambda', 94.31627052007231]]
```

2.2 Conclusion

Based on this quick analysis, we can conclude that all three dimensions are equally important (as expected, because the analytical function is known).

2.3 Exercises

- Important:
 - Results from these exercises should be added to this document, i.e., you should submit an updated version of this notebook.
 - Please combine your results using this notebook.
 - Only one notebook from each group!
 - Presentation is based on this notebook. No additional slides are required!
 - spotPython version 0.16.11 (or greater) is required

2.3.1 The Three Dimensional `fun_cubed`

- The input dimension is 3. The search range is $-1 \leq x \leq 1$ for all dimensions.
- Generate contour plots
- Calculate the variable importance.
- Discuss the variable importance:
 - Are all variables equally important?
 - If not:
 - * Which is the most important variable?
 - * Which is the least important variable?

2.3.2 The Ten Dimensional `fun_wing_wt`

- The input dimension is 10. The search range is $0 \leq x \leq 1$ for all dimensions.
- Calculate the variable importance.
- Discuss the variable importance:
 - Are all variables equally important?
 - If not:
 - * Which is the most important variable?
 - * Which is the least important variable?
 - Generate contour plots for the three most important variables. Do they confirm your selection?

2.3.3 The Three Dimensional `fun_runge`

- The input dimension is 3. The search range is $-5 \leq x \leq 5$ for all dimensions.
- Generate contour plots
- Calculate the variable importance.
- Discuss the variable importance:
 - Are all variables equally important?
 - If not:
 - * Which is the most important variable?
 - * Which is the least important variable?

2.3.4 The Three Dimensional `fun_linear`

- The input dimension is 3. The search range is $-5 \leq x \leq 5$ for all dimensions.
- Generate contour plots
- Calculate the variable importance.
- Discuss the variable importance:
 - Are all variables equally important?
 - If not:
 - * Which is the most important variable?
 - * Which is the least important variable?

3 Isotropic and Anisotropic Kriging

3.1 Example: Isotropic Spot Surrogate and the 2-dim Sphere Function

```
import numpy as np
from math import inf
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from scipy.optimize import shgo
from scipy.optimize import direct
from scipy.optimize import differential_evolution
import matplotlib.pyplot as plt
```

3.1.1 The Objective Function: 2-dim Sphere

- The `spotPython` package provides several classes of objective functions.
- We will use an analytical objective function, i.e., a function that can be described by a (closed) formula:

$$f(x, y) = x^2 + y^2$$

```
fun = analytical().fun_sphere
fun_control = {"sigma": 0,
              "seed": 123}
```

- The size of the `lower` bound vector determines the problem dimension.
- Here we will use `np.array([-1, -1])`, i.e., a two-dim function.

```
spot_2 = spot.Spot(fun=fun,
                  lower = np.array([-1, -1]),
                  upper = np.array([1, 1]))

spot_2.run()
```

```
<spotPython.spot.spot.Spot at 0x148c23880>
```

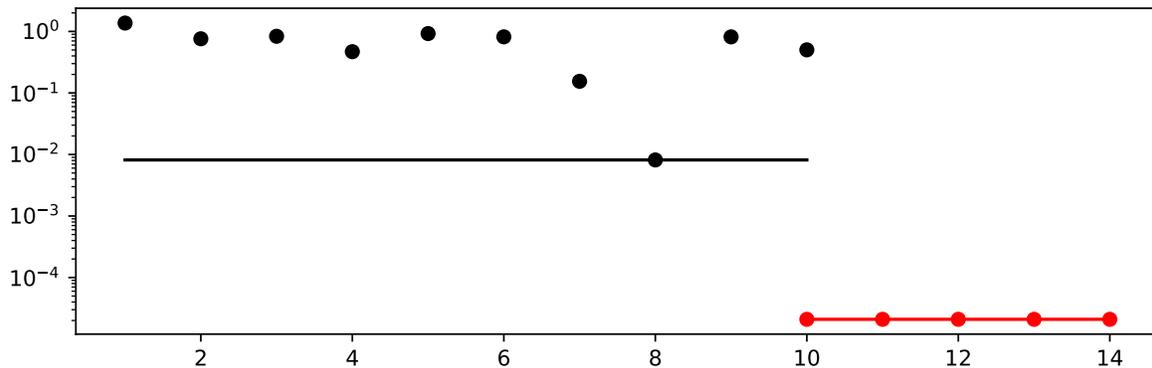
3.1.2 Results

```
spot_2.print_results()
```

```
min y: 2.093282610941807e-05  
x0: 0.0016055267473267492  
x1: 0.00428428640184529
```

```
[['x0', 0.0016055267473267492], ['x1', 0.00428428640184529]]
```

```
spot_2.plot_progress(log_y=True)
```



3.2 Example With Anisotropic Kriging

- The default parameter setting of `spotPython`'s Kriging surrogate uses the same `theta` value for every dimension.
- This is referred to as “using an isotropic kernel”.
- If different `theta` values are used for each dimension, then an anisotropic kernel is used
- To enable anisotropic models in `spotPython`, the number of `theta` values should be larger than one.
- We can use `surrogate_control={"n_theta": 2}` to enable this behavior (2 is the problem dimension).

```

spot_2_anisotropic = spot.Spot(fun=fun,
                               lower = np.array([-1, -1]),
                               upper = np.array([1, 1]),
                               surrogate_control={"n_theta": 2})
spot_2_anisotropic.run()

```

<spotPython.spot.spot.Spot at 0x14ffaf6d0>

3.2.1 Taking a Look at the theta Values

- We can check, whether one or several `theta` values were used.
- The `theta` values from the surrogate can be printed as follows:

```
spot_2_anisotropic.surrogate.theta
```

```
array([0.19447342, 0.30813872])
```


- Since the surrogate from the isotropic setting was stored as `spot_2`, we can also take a look at the `theta` value from this model:

```
spot_2.surrogate.theta
```

```
array([0.26287447])
```

- Next, the search progress of the optimization with the anisotropic model can be visualized:

```
spot_2_anisotropic.plot_progress(log_y=True)
```




```
spot_2_anisotropic.print_results()
```

```
min y: 7.77061191821505e-06  
x0: -0.0024488252797500764  
x1: -0.0013318658594137815
```

```
[['x0', -0.0024488252797500764], ['x1', -0.0013318658594137815]]
```

```
spot_2_anisotropic.surrogate.plot()
```



3.3 Exercises

3.3.1 fun_branin

- Describe the function.
 - The input dimension is 2. The search range is $-5 \leq x_1 \leq 10$ and $0 \leq x_2 \leq 15$.

- Compare the results from `spotPython` run a) with isotropic and b) anisotropic surrogate models.
- Modify the termination criterion: instead of the number of evaluations (which is specified via `fun_evals`), the time should be used as the termination criterion. This can be done as follows (`max_time=1` specifies a run time of one minute):

```
fun_evals=inf,
max_time=1,
```

3.3.2 `fun_sin_cos`

- Describe the function.
 - The input dimension is 2. The search range is $-2\pi \leq x_1 \leq 2\pi$ and $-2\pi \leq x_2 \leq 2\pi$.
- Compare the results from `spotPython` run a) with isotropic and b) anisotropic surrogate models.
- Modify the termination criterion (`max_time` instead of `fun_evals`) as described for `fun_branin`.

3.3.3 `fun_runge`

- Describe the function.
 - The input dimension is 2. The search range is $-5 \leq x_1 \leq 5$ and $-5 \leq x_2 \leq 5$.
- Compare the results from `spotPython` run a) with isotropic and b) anisotropic surrogate models.
- Modify the termination criterion (`max_time` instead of `fun_evals`) as described for `fun_branin`.

3.3.4 `fun_wingwt`

- Describe the function.
 - The input dimension is 10. The search ranges are between 0 and 1 (values are mapped internally to their natural bounds).
- Compare the results from `spotPython` run a) with isotropic and b) anisotropic surrogate models.
- Modify the termination criterion (`max_time` instead of `fun_evals`) as described for `fun_branin`.

4 Using sklearn Surrogates in spotPython

This notebook explains how different surrogate models from `scikit-learn` can be used as surrogates in `spotPython` optimization runs.

```
import numpy as np
from math import inf
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from scipy.optimize import shgo
from scipy.optimize import direct
from scipy.optimize import differential_evolution
import matplotlib.pyplot as plt
```

4.1 Example: Branin Function with spotPython's Internal Kriging Surrogate

4.1.1 The Objective Function Branin

- The `spotPython` package provides several classes of objective functions.
- We will use an analytical objective function, i.e., a function that can be described by a (closed) formula.
- Here we will use the Branin function:

$$y = a * (x_2 - b * x_1^2 + c * x_1 - r) ** 2 + s * (1 - t) * \text{np.cos}(x_1) + s,$$
where values of a , b , c , r , s and t are: $a = 1$, $b = 5.1 / (4 * \text{pi}^2)$,
 $c = 5 / \text{pi}$, $r = 6$, $s = 10$ and $t = 1 / (8 * \text{pi})$.

- It has three global minima:

$$f(x) = 0.397887 \text{ at } (-\text{pi}, 12.275), (\text{pi}, 2.275), \text{ and } (9.42478, 2.475).$$

```
from spotPython.fun.objectivefunctions import analytical
lower = np.array([-5,-0])
```

```
upper = np.array([10,15])

fun = analytical().fun_branin
```

4.1.2 Running the surrogate model based optimizer Spot:

```
spot_2 = spot.Spot(fun=fun,
                  lower = lower,
                  upper = upper,
                  fun_evals = 20,
                  max_time = inf,
                  seed=123,
                  design_control={"init_size": 10})
```

```
spot_2.run()
```

<spotPython.spot.spot.Spot at 0x104c72050>

4.1.3 Print the Results

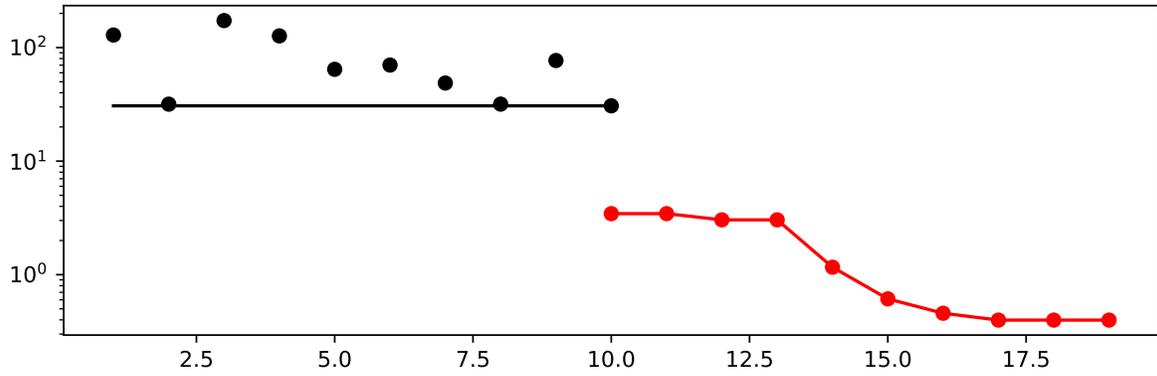
```
spot_2.print_results()
```

```
min y: 0.3982295132785083
x0: 3.135528626303215
x1: 2.2926027772585886
```

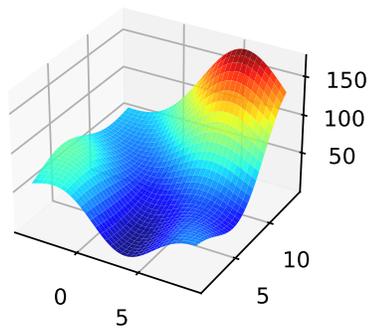
```
[['x0', 3.135528626303215], ['x1', 2.2926027772585886]]
```

4.1.4 Show the Progress and the Surrogate

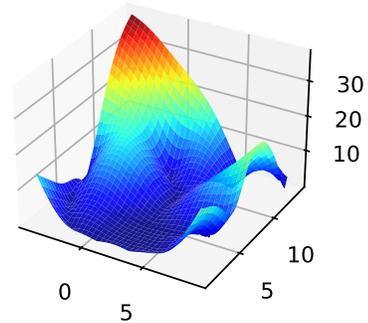
```
spot_2.plot_progress(log_y=True)
```



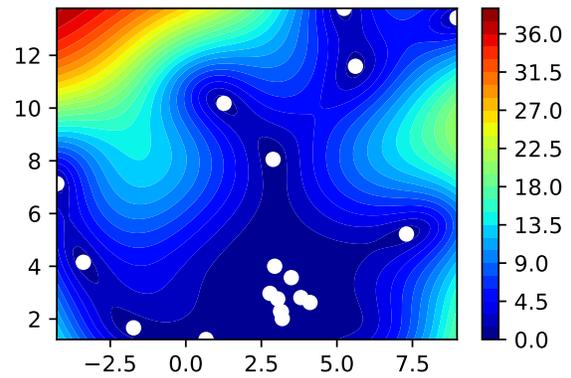
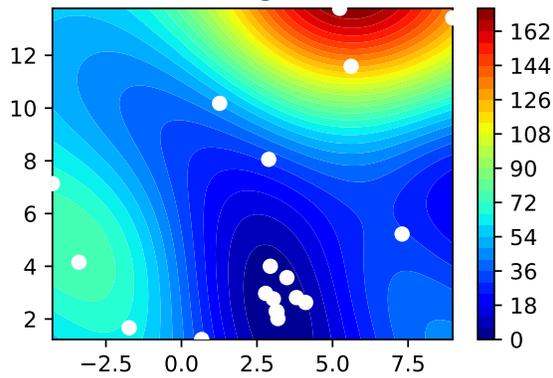
```
spot_2.surrogate.plot()
```



Surrogate



Error



4.2 Example: Using Surrogates From scikit-learn

- Default is the `spotPython` (i.e., the internal) `kriging` surrogate.

- It can be called explicitly and passed to `Spot`.

```
from spotPython.build.kriging import Kriging
S_0 = Kriging(name='kriging', seed=123)
```

- Alternatively, models from `scikit-learn` can be selected, e.g., Gaussian Process, RBFs, Regression Trees, etc.

```
# Needed for the sklearn surrogates:
from sklearn.gaussian_process import GaussianProcessRegressor
from sklearn.gaussian_process.kernels import RBF
from sklearn.tree import DecisionTreeRegressor
from sklearn.ensemble import RandomForestRegressor
from sklearn import linear_model
from sklearn import tree
import pandas as pd
```

- Here are some additional models that might be useful later:

```
S_Tree = DecisionTreeRegressor(random_state=0)
S_LM = linear_model.LinearRegression()
S_Ridge = linear_model.Ridge()
S_RF = RandomForestRegressor(max_depth=2, random_state=0)
```

4.2.1 GaussianProcessRegressor as a Surrogate

- To use a Gaussian Process model from `sklearn`, that is similar to `spotPython`'s `Kriging`, we can proceed as follows:

```
kernel = 1 * RBF(length_scale=1.0, length_scale_bounds=(1e-2, 1e2))
S_GP = GaussianProcessRegressor(kernel=kernel, n_restarts_optimizer=9)
```

- The `scikit-learn` GP model `S_GP` is selected for `Spot` as follows:

```
surrogate = S_GP
```

- We can check the kind of surrogate model with the command `isinstance`:

```
isinstance(S_GP, GaussianProcessRegressor)
```

True

```
isinstance(S_0, Kriging)
```

True


- Similar to the Spot run with the internal Kriging model, we can call the run with the scikit-learn surrogate:

```
fun = analytical(seed=123).fun_branin
spot_2_GP = spot.Spot(fun=fun,
                      lower = lower,
                      upper = upper,
                      fun_evals = 20,
                      seed=123,
                      design_control={"init_size": 10},
                      surrogate = S_GP)

spot_2_GP.run()
```

<spotPython.spot.spot.Spot at 0x164d637c0>

```
spot_2_GP.plot_progress()
```



```
spot_2_GP.print_results()
```

```
min y: 0.39821939809704077
x0: 3.14976322931894
x1: 2.272032569093181
```

```
[['x0', 3.14976322931894], ['x1', 2.272032569093181]]
```

4.3 Example: One-dimensional Sphere Function With spotPython's Kriging

- In this example, we will use an one-dimensional function, which allows us to visualize the optimization process.

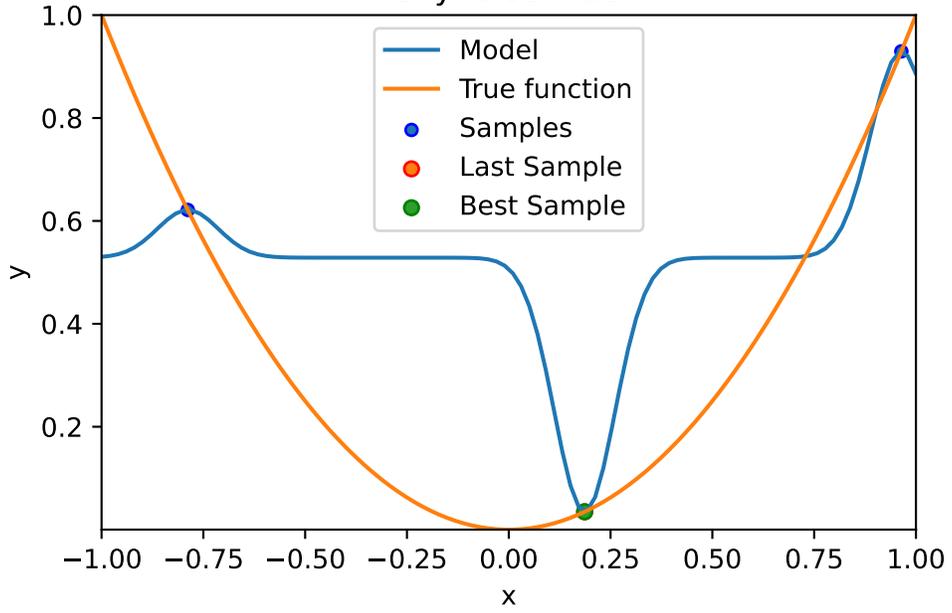
– `show_models= True` is added to the argument list.

```
from spotPython.fun.objectivefunctions import analytical
lower = np.array([-1])
upper = np.array([1])
fun = analytical(seed=123).fun_sphere

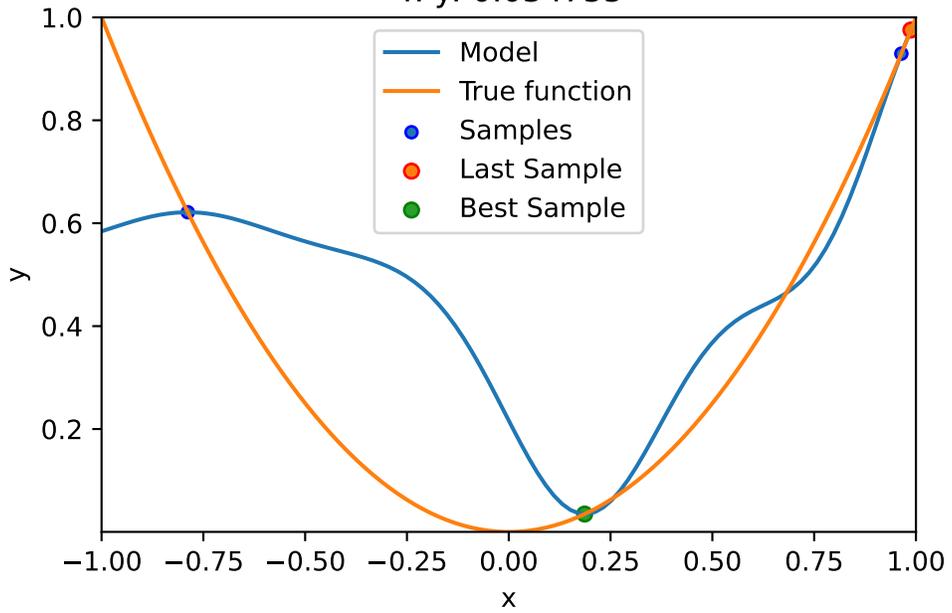
spot_1 = spot.Spot(fun=fun,
                  lower = lower,
                  upper = upper,
                  fun_evals = 10,
                  max_time = inf,
                  seed=123,
                  show_models= True,
                  tolerance_x = np.sqrt(np.spacing(1)),
                  design_control={"init_size": 3},)

spot_1.run()
```


3. $y: 0.034755$




4. $y: 0.034755$




5. $y: 0.034755$




6. $y: 0.014959$





7. y: 0.000215



8. y: 0.0





<spotPython.spot.spot.Spot at 0x164d63010>

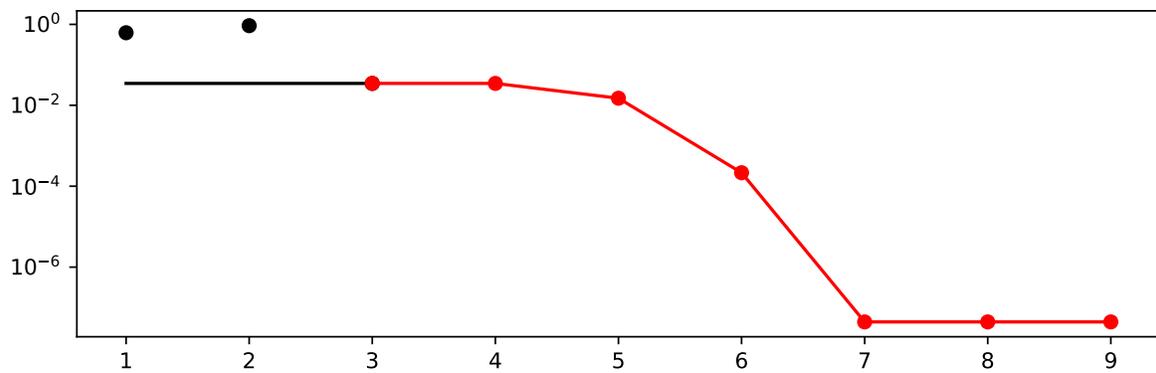
4.3.1 Results

```
spot_1.print_results()
```

```
min y: 4.41925228274096e-08  
x0: -0.00021022017702259125
```

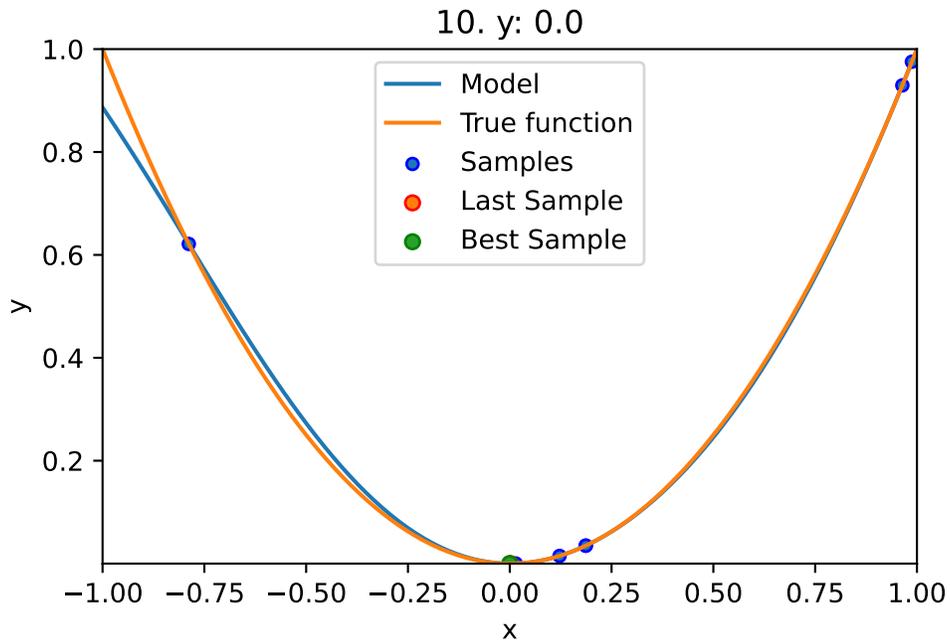
```
[['x0', -0.00021022017702259125]]
```

```
spot_1.plot_progress(log_y=True)
```



- The method `plot_model` plots the final surrogate:

```
spot_1.plot_model()
```



4.4 Example: Sklearn Model GaussianProcess

- This example visualizes the search process on the `GaussianProcessRegression` surrogate from `sklearn`.
- Therefore `surrogate = S_GP` is added to the argument list.


```

fun = analytical(seed=123).fun_sphere
spot_1_GP = spot.Spot(fun=fun,
                      lower = lower,
                      upper = upper,
                      fun_evals = 10,
                      max_time = inf,
                      seed=123,
                      show_models= True,
                      design_control={"init_size": 3},
                      surrogate = S_GP)


spot_1_GP.run()

```


3. y: 0.034755




4. y: 0.004926







5. y: 0.002612



6. y: 0.0








<spotPython.spot.spot.Spot at 0x2a05ddf00>

```
spot_1_GP.print_results()
```


```
min y: 3.1635674949756653e-11  
x0: -5.62455997832334e-06
```

```
[['x0', -5.62455997832334e-06]]
```

```
spot_1_GP.plot_progress(log_y=True)
```



```
spot_1_GP.plot_model()
```



4.5 Exercises

4.5.1 DecisionTreeRegressor

- Describe the surrogate model.
- Use the surrogate as the model for optimization.

4.5.2 RandomForestRegressor

- Describe the surrogate model.
- Use the surrogate as the model for optimization.

4.5.3 linear_model.LinearRegression

- Describe the surrogate model.
- Use the surrogate as the model for optimization.

4.5.4 `linear_model.Ridge`

- Describe the surrogate model.
- Use the surrogate as the model for optimization.

4.6 Exercise 2

- Compare the performance of the five different surrogates on both objective functions:
 - `spotPython`'s internal Kriging
 - `DecisionTreeRegressor`
 - `RandomForestRegressor`
 - `linear_model.LinearRegression`
 - `linear_model.Ridge`

5 Sequential Parameter Optimization: Using scipy Optimizers

This notebook describes how different optimizers from the `scipy optimize` package can be used on the surrogate. The optimization algorithms are available from <https://docs.scipy.org/doc/scipy/reference/optimize.html>

```
import numpy as np
from math import inf
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from scipy.optimize import shgo
from scipy.optimize import direct
from scipy.optimize import differential_evolution
from scipy.optimize import dual_annealing
from scipy.optimize import basinhopping
import matplotlib.pyplot as plt
```

5.1 The Objective Function Branin

- The `spotPython` package provides several classes of objective functions.
- We will use an analytical objective function, i.e., a function that can be described by a (closed) formula.
- Here we will use the Branin function. The 2-dim Branin function is

$$y = a * (x_2 - b * x_1^2 + c * x_1 - r)^2 + s * (1 - t) * \cos(x_1) + s,$$

where values of a , b , c , r , s and t are: $a = 1$, $b = 5.1/(4 * \pi^2)$, $c = 5/\pi$, $r = 6$, $s = 10$ and $t = 1/(8 * \pi)$.

- It has three global minima:

$$f(x) = 0.397887 \text{ at } (-\pi, 12.275), (\pi, 2.275), \text{ and } (9.42478, 2.475).$$

- Input Domain: This function is usually evaluated on the square x_1 in $[-5, 10]$ x x_2 in $[0, 15]$.

```
from spotPython.fun.objectivefunctions import analytical
lower = np.array([-5,-0])
upper = np.array([10,15])

fun = analytical(seed=123).fun_branin
```

5.2 The Optimizer

- Differential Evolution from the `scikit.optimize` package, see https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.differential_evolution.html#scipy.optimize.differential_evolution is the default optimizer for the search on the surrogate.

- Other optimizers that are available in `spotPython`:

- `dual_annealing`
- `direct`
- `shgo`
- `basinhopping`, see <https://docs.scipy.org/doc/scipy/reference/optimize.html#global-optimization>.

- These can be selected as follows:

```
surrogate_control = "model_optimizer": differential_evolution
```

- We will use `differential_evolution`.
- The optimizer can use 1000 evaluations. This value will be passed to the `differential_evolution` method, which has the argument `maxiter` (int). It defines the maximum number of generations over which the entire differential evolution population is evolved, see https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.differential_evolution.html#scipy.optimize.differential_evolution

```
spot_de = spot.Spot(fun=fun,
                   lower = lower,
                   upper = upper,
                   fun_evals = 20,
                   max_time = inf,
                   seed=125,
                   noise=False,
```

```
show_models= False,  
design_control={"init_size": 10},  
surrogate_control={"n_theta": 2,  
                   "model_optimizer": differential_evolution,  
                   "model_fun_evals": 1000,  
                   })  
  
spot_de.run()
```

<spotPython.spot.spot.Spot at 0x107a70880>

5.3 Print the Results

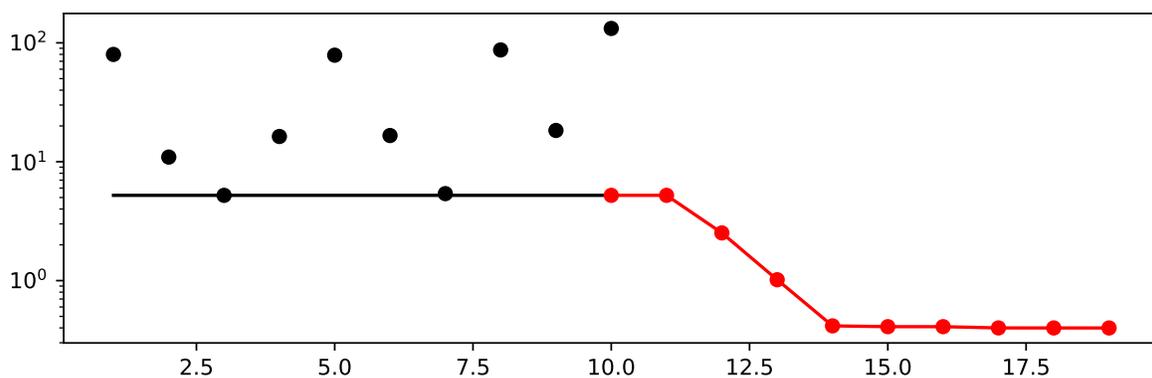
```
spot_de.print_results()
```

```
min y: 0.39951958110619046  
x0: -3.1570201165683587  
x1: 12.289980569430284
```

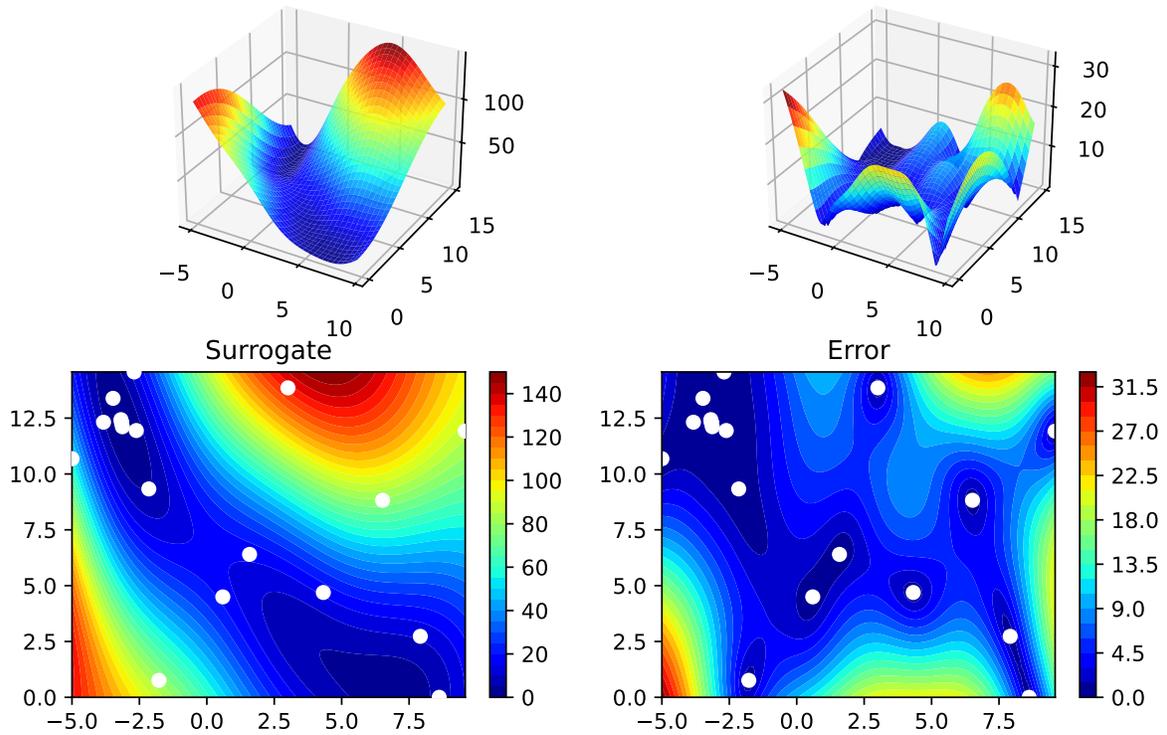
```
[['x0', -3.1570201165683587], ['x1', 12.289980569430284]]
```

5.4 Show the Progress

```
spot_de.plot_progress(log_y=True)
```



```
spot_de.surrogate.plot()
```



5.5 Exercises

5.5.1 dual_annealing

- Describe the optimization algorithm
- Use the algorithm as an optimizer on the surrogate

5.5.2 direct

- Describe the optimization algorithm
- Use the algorithm as an optimizer on the surrogate

5.5.3 shgo

- Describe the optimization algorithm
- Use the algorithm as an optimizer on the surrogate

5.5.4 basinhopping

- Describe the optimization algorithm
- Use the algorithm as an optimizer on the surrogate

5.5.5 Performance Comparison

Compare the performance and run time of the 5 different optimizers:

```
* `differential_evolution`  
* `dual_annealing`  
* `direct`  
* `shgo`  
* `basinhopping`.
```

The Branin function has three global minima:

- $f(x) = 0.397887$ at
 - $(-\pi, 12.275)$,
 - $(\pi, 2.275)$, and
 - $(9.42478, 2.475)$.
- Which optima are found by the optimizers? Does the `seed` change this behavior?

6 Sequential Parameter Optimization: Gaussian Process Models

- This notebook analyzes differences between
 - the Kriging implementation in `spotPython` and
 - the `GaussianProcessRegressor` in `scikit-learn`.

```
import numpy as np
from math import inf
from spotPython.fun.objectivefunctions import analytical
from spotPython.design.spacefilling import spacefilling
from spotPython.spot import spot
from spotPython.build.kriging import Kriging
from scipy.optimize import shgo
from scipy.optimize import direct
from scipy.optimize import differential_evolution
import matplotlib.pyplot as plt
import math as m
from sklearn.gaussian_process import GaussianProcessRegressor
from sklearn.gaussian_process.kernels import RBF
```

6.1 Gaussian Processes Regression: Basic Introductory `scikit-learn` Example

- This is the example from `scikit-learn`: https://scikit-learn.org/stable/auto_examples/gaussian_process/pl
- After fitting our model, we see that the hyperparameters of the kernel have been optimized.
- Now, we will use our kernel to compute the mean prediction of the full dataset and plot the 95% confidence interval.

6.1.1 Train and Test Data

```
X = np.linspace(start=0, stop=10, num=1_000).reshape(-1, 1)
y = np.squeeze(X * np.sin(X))
rng = np.random.RandomState(1)
training_indices = rng.choice(np.arange(y.size), size=6, replace=False)
X_train, y_train = X[training_indices], y[training_indices]
```

6.1.2 Building the Surrogate With Sklearn


- The model building with `sklearn` consists of three steps:
 1. Instantiating the model, then
 2. fitting the model (using `fit`), and
 3. making predictions (using `predict`)

```
kernel = 1 * RBF(length_scale=1.0, length_scale_bounds=(1e-2, 1e2))
gaussian_process = GaussianProcessRegressor(kernel=kernel, n_restarts_optimizer=9)
gaussian_process.fit(X_train, y_train)
mean_prediction, std_prediction = gaussian_process.predict(X, return_std=True)
```

6.1.3 Plotting the SklearnModel

```
plt.plot(X, y, label=r"$f(x) = x \sin(x)$", linestyle="dotted")
plt.scatter(X_train, y_train, label="Observations")
plt.plot(X, mean_prediction, label="Mean prediction")
plt.fill_between(
    X.ravel(),
    mean_prediction - 1.96 * std_prediction,
    mean_prediction + 1.96 * std_prediction,
    alpha=0.5,
    label=r"95% confidence interval",
)
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("sk-learn Version: Gaussian process regression on noise-free dataset")
```

sk-learn Version: Gaussian process regression on noise-free dataset



6.1.4 The spotPython Version

- The spotPython version is very similar:
 1. Instantiating the model, then
 2. fitting the model and
 3. making predictions (using `predict`).

```
S = Kriging(name='kriging', seed=123, log_level=50, cod_type="norm")
S.fit(X_train, y_train)
S_mean_prediction, S_std_prediction, S_ei = S.predict(X, return_val="all")
```

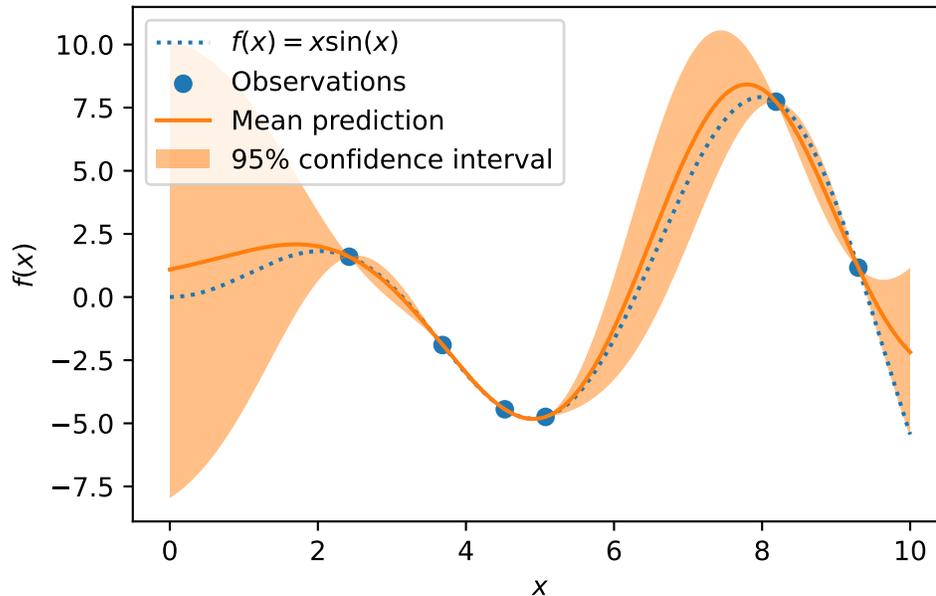
```
plt.plot(X, y, label=r"$f(x) = x \sin(x)$", linestyle="dotted")
plt.scatter(X_train, y_train, label="Observations")
plt.plot(X, S_mean_prediction, label="Mean prediction")
plt.fill_between(
    X.ravel(),
    S_mean_prediction - 1.96 * S_std_prediction,
    S_mean_prediction + 1.96 * S_std_prediction,
    alpha=0.5,
    label=r"95% confidence interval",
```

```

)
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("spotPython Version: Gaussian process regression on noise-free dataset")

```

spotPython Version: Gaussian process regression on noise-free dataset




6.1.5 Visualizing the Differences Between the spotPython and the sklearn Model Fits

```

plt.plot(X, y, label=r"$f(x) = x \sin(x)$", linestyle="dotted")
plt.scatter(X_train, y_train, label="Observations")
plt.plot(X, S_mean_prediction, label="spotPython Mean prediction")
plt.plot(X, mean_prediction, label="Sklearn Mean Prediction")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Comparing Mean Predictions")

```



6.2 Exercises

6.2.1 Schonlau Example Function

- The Schonlau Example Function is based on sample points only (there is no analytical function description available):

```
X = np.linspace(start=0, stop=13, num=1_000).reshape(-1, 1)
X_train = np.array([1., 2., 3., 4., 12.]).reshape(-1,1)
y_train = np.array([0., -1.75, -2, -0.5, 5.])
```

- Describe the function.
- Compare the two models that were build using the `spotPython` and the `sklearn` surrogate.
- Note: Since there is no analytical function available, you might be interested in adding some points and describe the effects.

6.2.2 Forrester Example Function

- The Forrester Example Function is defined as follows:

$f(x) = (6x - 2)^2 \sin(12x - 4)$ for x in $[0, 1]$.

- Data points are generated as follows:

```
X = np.linspace(start=-0.5, stop=1.5, num=1_000).reshape(-1, 1)
X_train = np.array([0.0, 0.175, 0.225, 0.3, 0.35, 0.375, 0.5, 1]).reshape(-1, 1)
fun = analytical().fun_forrester
fun_control = {"sigma": 0.1,
               "seed": 123}
y = fun(X, fun_control=fun_control)
y_train = fun(X_train, fun_control=fun_control)
```

- Describe the function.
- Compare the two models that were build using the `spotPython` and the `sklearn` surrogate.
- Note: Modify the noise level ("`sigma`"), e.g., use a value of 0.2, and compare the two models.

```
fun_control = {"sigma": 0.2}
```

6.2.3 fun_runge Function (1-dim)

- The Runge function is defined as follows:

$f(x) = 1 / (1 + \sum(x_i))^2$

- Data points are generated as follows:

```
gen = spacefilling(1)
rng = np.random.RandomState(1)
lower = np.array([-10])
upper = np.array([10])
fun = analytical().fun_runge
fun_control = {"sigma": 0.025,
               "seed": 123}
X_train = gen.scipy_lhd(10, lower=lower, upper = upper).reshape(-1, 1)
y_train = fun(X, fun_control=fun_control)
X = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
y = fun(X, fun_control=fun_control)
```

- Describe the function.

- Compare the two models that were build using the `spotPython` and the `sklearn` surrogate.
- Note: Modify the noise level ("`sigma`"), e.g., use a value of 0.05, and compare the two models.

```
fun_control = {"sigma": 0.5}
```

6.2.4 fun_cubed (1-dim)

- The Cubed function is defined as follows:

```
np.sum(X[i]** 3)
```

- Data points are generated as follows:

```
gen = spacefilling(1)
rng = np.random.RandomState(1)
lower = np.array([-10])
upper = np.array([10])
fun = analytical().fun_cubed
fun_control = {"sigma": 0.025,
              "seed": 123}
X_train = gen.scipy_lhd(10, lower=lower, upper = upper).reshape(-1,1)
y_train = fun(X, fun_control=fun_control)
X = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
y = fun(X, fun_control=fun_control)
```

- Describe the function.
- Compare the two models that were build using the `spotPython` and the `sklearn` surrogate.
- Note: Modify the noise level ("`sigma`"), e.g., use a value of 0.05, and compare the two models.

```
fun_control = {"sigma": 0.05}
```

6.2.5 The Effect of Noise

How does the behavior of the `spotPython` fit changes when the argument `noise` is set to `True`, i.e.,

```
S = Kriging(name='kriging', seed=123, n_theta=1, noise=True)
```

is used?

7 Expected Improvement

7.1 Example: Spot and the 1-dim Sphere Function

```
import numpy as np
from math import inf
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from scipy.optimize import shgo
from scipy.optimize import direct
from scipy.optimize import differential_evolution
import matplotlib.pyplot as plt
```

7.1.1 The Objective Function: 1-dim Sphere

- The spotPython package provides several classes of objective functions.
- We will use an analytical objective function, i.e., a function that can be described by a (closed) formula:

$$f(x) = x^2$$

```
fun = analytical().fun_sphere
```

```
fun = analytical().fun_sphere
fun_control = {"sigma": 0,
              "seed": 123}
```

- The size of the lower bound vector determines the problem dimension.
- Here we will use `np.array([-1])`, i.e., a one-dim function.

```
spot_1 = spot.Spot(fun=fun,
                  lower = np.array([-1]),
                  upper = np.array([1]))
```

```
spot_1.run()
```

```
<spotPython.spot.spot.Spot at 0x15a763b50>
```

7.1.2 Results

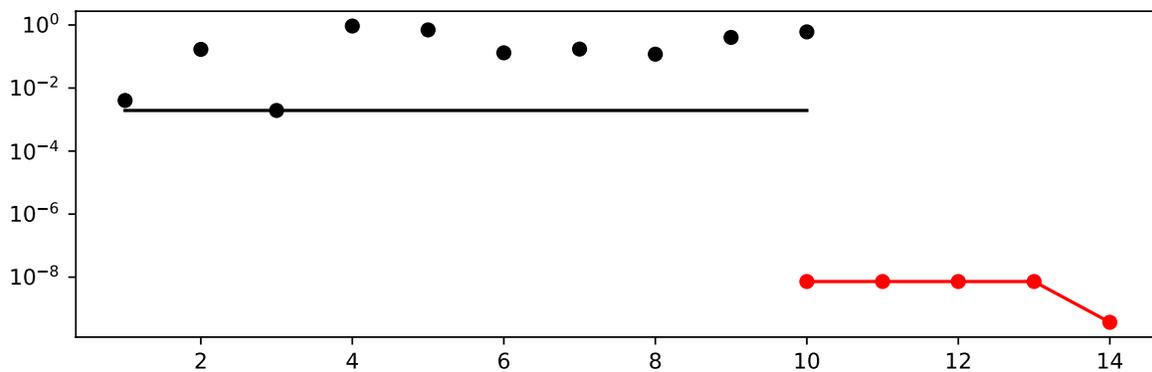
```
spot_1.print_results()
```

```
min y: 3.696886711914087e-10
```

```
x0: 1.922728975158508e-05
```

```
[['x0', 1.922728975158508e-05]]
```

```
spot_1.plot_progress(log_y=True)
```

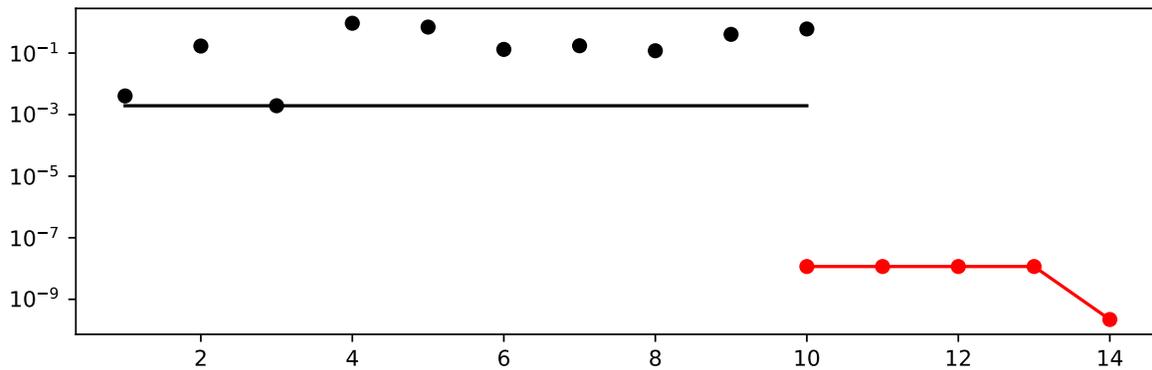


7.2 Same, but with EI as infill_criterion

```
spot_1_ei = spot.Spot(fun=fun,  
                      lower = np.array([-1]),  
                      upper = np.array([1]),  
                      infill_criterion = "ei")  
spot_1_ei.run()
```

```
<spotPython.spot.spot.Spot at 0x15fee73d0>
```

```
spot_1_ei.plot_progress(log_y=True)
```



```
spot_1_ei.print_results()
```

```
min y: 2.207887258868953e-10  
x0: 1.4858961130809088e-05
```

```
[['x0', 1.4858961130809088e-05]]
```

7.3 Non-isotropic Kriging

```
spot_2_ei_noniso = spot.Spot(fun=fun,  
                             lower = np.array([-1, -1]),  
                             upper = np.array([1, 1]),  
                             fun_evals = 20,  
                             fun_repeats = 1,  
                             max_time = inf,  
                             noise = False,  
                             tolerance_x = np.sqrt(np.spacing(1)),  
                             var_type=["num"],  
                             infill_criterion = "ei",  
                             n_points = 1,  
                             seed=123,  
                             log_level = 50,  
                             show_models=True,  
                             fun_control = fun_control,
```

```

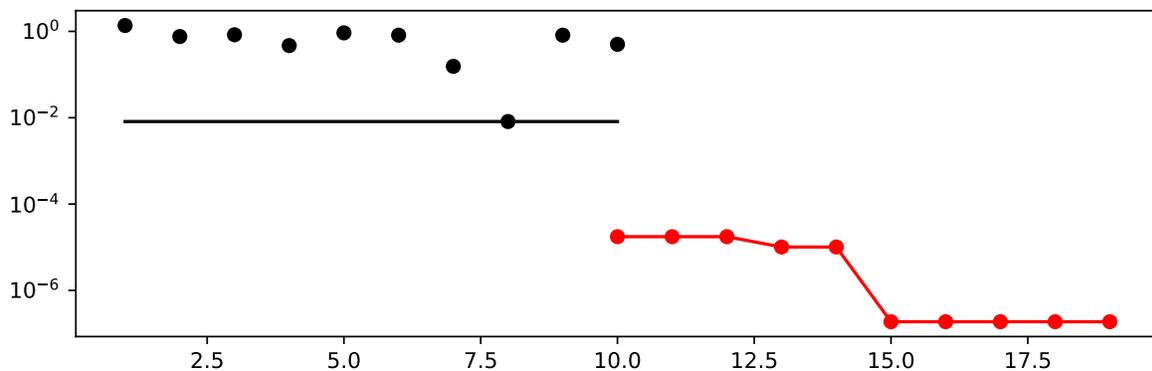
design_control={"init_size": 10,
               "repeats": 1},
surrogate_control={"noise": False,
                  "cod_type": "norm",
                  "min_theta": -4,
                  "max_theta": 3,
                  "n_theta": 2,
                  "model_optimizer": differential_evolution,
                  "model_fun_evals": 1000,
                  })

spot_2_ei_noniso.run()

```

<spotPython.spot.spot.Spot at 0x1694c5fc0>

```
spot_2_ei_noniso.plot_progress(log_y=True)
```



```
spot_2_ei_noniso.print_results()
```

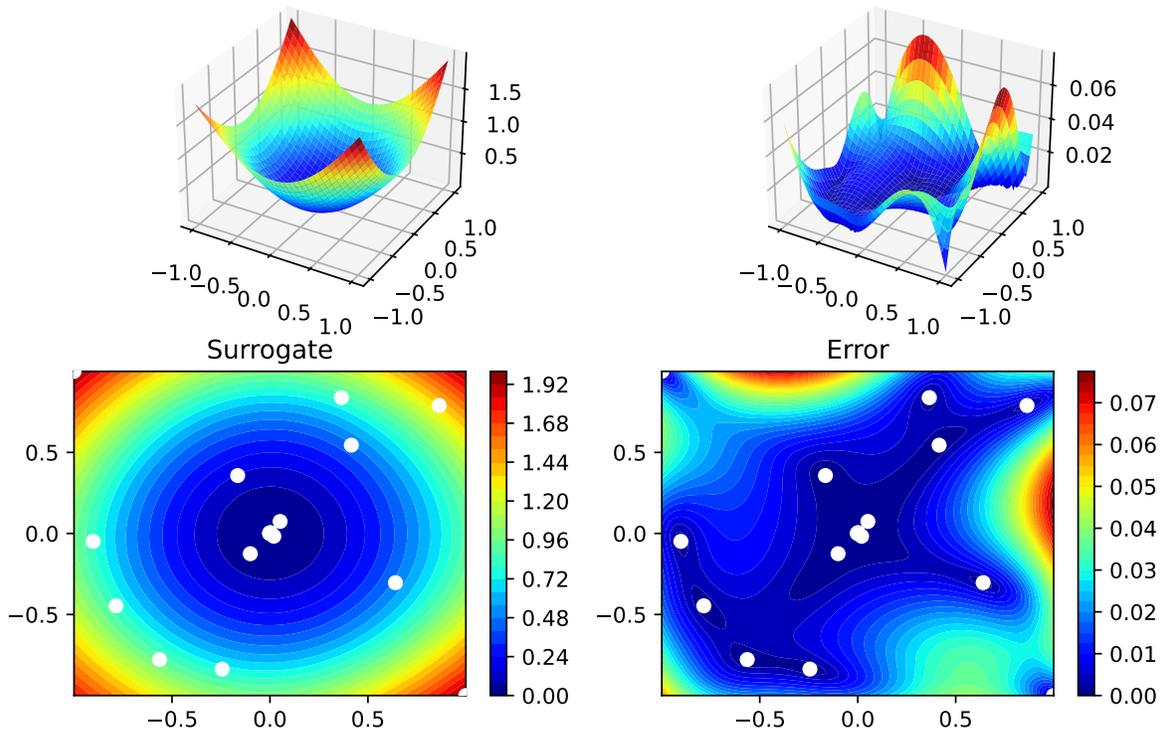
```

min y: 1.8779971830281702e-07
x0: -0.0002783721390529846
x1: 0.0003321274913371111

```

```
[['x0', -0.0002783721390529846], ['x1', 0.0003321274913371111]]
```

```
spot_2_ei_noniso.surrogate.plot()
```



7.4 Using sklearn Surrogates

7.4.1 The spot Loop

The spot loop consists of the following steps:

1. Init: Build initial design X
2. Evaluate initial design on real objective f : $y = f(X)$
3. Build surrogate: $S = S(X, y)$
4. Optimize on surrogate: $X_0 = \text{optimize}(S)$
5. Evaluate on real objective: $y_0 = f(X_0)$
6. Impute (Infill) new points: $X = X \cup X_0$, $y = y \cup y_0$.
7. Got 3.

The spot loop is implemented in R as follows:

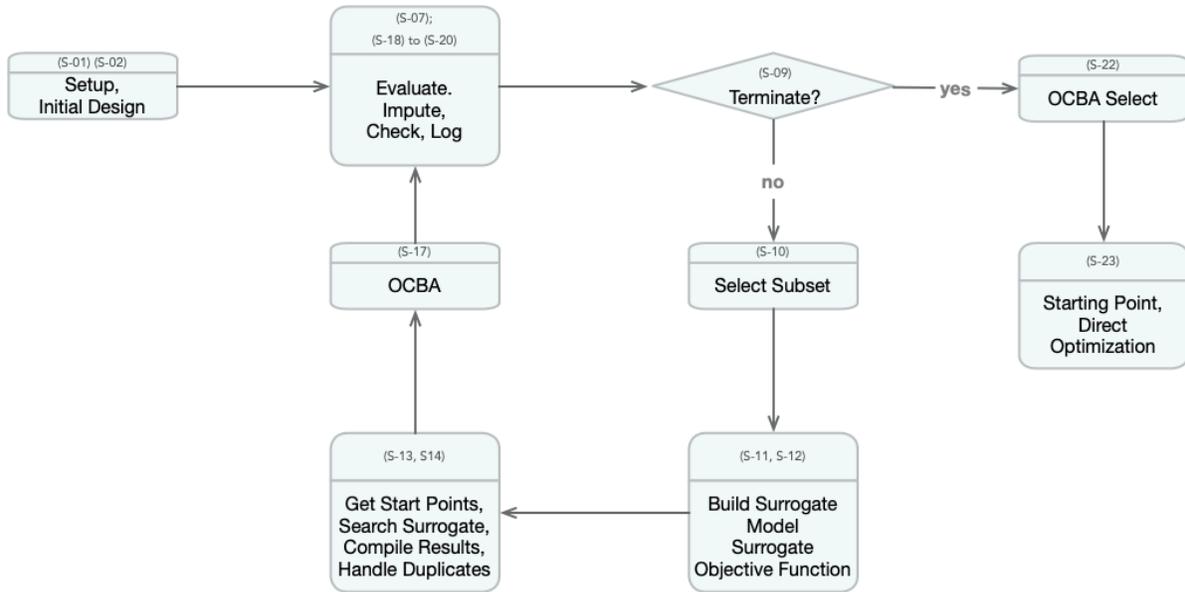


Figure 7.1: Visual representation of the model based search with SPOT. Taken from: Bartz-Beielstein, T., and Zaefferer, M. Hyperparameter tuning approaches. In Hyperparameter Tuning for Machine and Deep Learning with R - A Practical Guide, E. Bartz, T. Bartz-Beielstein, M. Zaefferer, and O. Mersmann, Eds. Springer, 2022, ch. 4, pp. 67–114.

7.4.2 spot: The Initial Model


7.4.2.1 Example: Modifying the initial design size

This is the “Example: Modifying the initial design size” from Chapter 4.5.1 in [bart21i].

```
spot_ei = spot.Spot(fun=fun,  
                   lower = np.array([-1,-1]),  
                   upper= np.array([1,1]),  
                   design_control={"init_size": 5})  
spot_ei.run()
```

<spotPython.spot.spot.Spot at 0x16b5c6d10>

```
spot_ei.plot_progress()
```



```
np.min(spot_1.y), np.min(spot_ei.y)
```

(3.696886711914087e-10, 1.7928640814182596e-05)

7.4.3 Init: Build Initial Design

```
from spotPython.design.spacefilling import spacefilling  
from spotPython.build.kriging import Kriging  
from spotPython.fun.objectivefunctions import analytical  
gen = spacefilling(2)
```

```

rng = np.random.RandomState(1)
lower = np.array([-5,-0])
upper = np.array([10,15])
fun = analytical().fun_branin
fun_control = {"sigma": 0,
               "seed": 123}

X = gen.scipy_lhd(10, lower=lower, upper = upper)
print(X)
y = fun(X, fun_control=fun_control)
print(y)

```

```

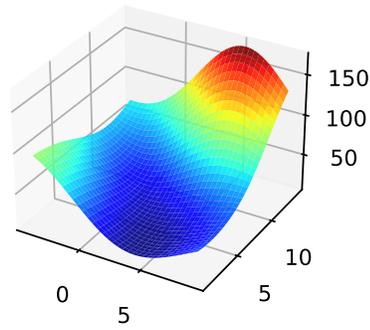
[[ 8.97647221 13.41926847]
 [ 0.66946019  1.22344228]
 [ 5.23614115 13.78185824]
 [ 5.6149825  11.5851384 ]
 [-1.72963184  1.66516096]
 [-4.26945568  7.1325531 ]
 [ 1.26363761 10.17935555]
 [ 2.88779942  8.05508969]
 [-3.39111089  4.15213772]
 [ 7.30131231  5.22275244]]
[128.95676449  31.73474356 172.89678121 126.71295908  64.34349975
 70.16178611  48.71407916  31.77322887  76.91788181  30.69410529]

```

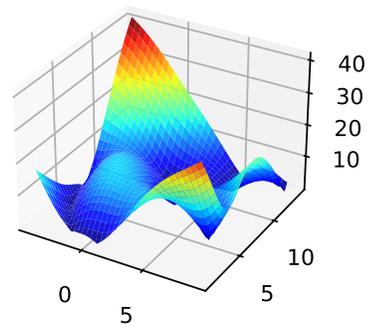
```

S = Kriging(name='kriging', seed=123)
S.fit(X, y)
S.plot()

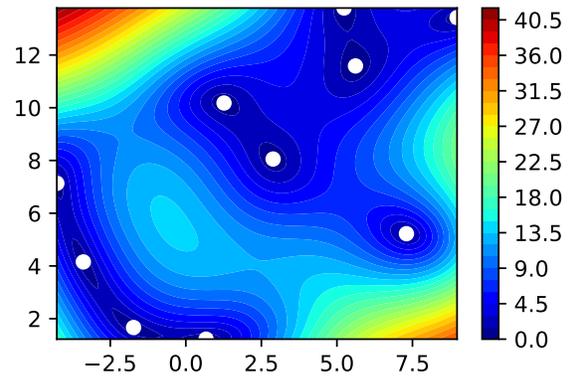
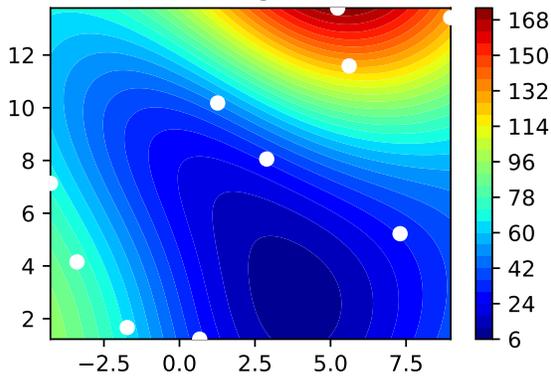
```



Surrogate



Error



```

gen = spacefilling(2, seed=123)
X0 = gen.scipy_lhd(3)
gen = spacefilling(2, seed=345)
X1 = gen.scipy_lhd(3)
X2 = gen.scipy_lhd(3)
gen = spacefilling(2, seed=123)
X3 = gen.scipy_lhd(3)
X0, X1, X2, X3

```

```

(array([[0.77254938, 0.31539299],
       [0.59321338, 0.93854273],
       [0.27469803, 0.3959685 ]]),
 array([[0.78373509, 0.86811887],
       [0.06692621, 0.6058029 ],
       [0.41374778, 0.00525456]]),
 array([[0.121357 , 0.69043832],
       [0.41906219, 0.32838498],
       [0.86742658, 0.52910374]]),

```

```
array([[0.77254938, 0.31539299],
       [0.59321338, 0.93854273],
       [0.27469803, 0.3959685 ]])
```

7.4.4 Evaluate

7.4.5 Build Surrogate

7.4.6 A Simple Predictor

The code below shows how to use a simple model for prediction.

- Assume that only two (very costly) measurements are available:
 1. $f(0) = 0.5$
 2. $f(2) = 2.5$
- We are interested in the value at $x_0 = 1$, i.e., $f(x_0 = 1)$, but cannot run an additional, third experiment.

```
from sklearn import linear_model
X = np.array([[0], [2]])
y = np.array([0.5, 2.5])
S_lm = linear_model.LinearRegression()
S_lm = S_lm.fit(X, y)
X0 = np.array([[1]])
y0 = S_lm.predict(X0)
print(y0)
```

[1.5]

- Central Idea:
 - Evaluation of the surrogate model `S_lm` is much cheaper (or / and much faster) than running the real-world experiment f .

7.5 Gaussian Processes regression: basic introductory example

This example was taken from [scikit-learn](#). After fitting our model, we see that the hyperparameters of the kernel have been optimized. Now, we will use our kernel to compute the mean prediction of the full dataset and plot the 95% confidence interval.

```

import numpy as np
import matplotlib.pyplot as plt
import math as m
from sklearn.gaussian_process import GaussianProcessRegressor
from sklearn.gaussian_process.kernels import RBF

X = np.linspace(start=0, stop=10, num=1_000).reshape(-1, 1)
y = np.squeeze(X * np.sin(X))
rng = np.random.RandomState(1)
training_indices = rng.choice(np.arange(y.size), size=6, replace=False)
X_train, y_train = X[training_indices], y[training_indices]


kernel = 1 * RBF(length_scale=1.0, length_scale_bounds=(1e-2, 1e2))
gaussian_process = GaussianProcessRegressor(kernel=kernel, n_restarts_optimizer=9)
gaussian_process.fit(X_train, y_train)
gaussian_process.kernel_

mean_prediction, std_prediction = gaussian_process.predict(X, return_std=True)

plt.plot(X, y, label=r"$f(x) = x \sin(x)$", linestyle="dotted")
plt.scatter(X_train, y_train, label="Observations")
plt.plot(X, mean_prediction, label="Mean prediction")
plt.fill_between(
    X.ravel(),
    mean_prediction - 1.96 * std_prediction,
    mean_prediction + 1.96 * std_prediction,
    alpha=0.5,
    label=r"95% confidence interval",
)
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("sk-learn Version: Gaussian process regression on noise-free dataset")

```

sk-learn Version: Gaussian process regression on noise-free dataset



```
from spotPython.build.kriging import Kriging
import numpy as np
import matplotlib.pyplot as plt
rng = np.random.RandomState(1)
X = np.linspace(start=0, stop=10, num=1_000).reshape(-1, 1)
y = np.squeeze(X * np.sin(X))
training_indices = rng.choice(np.arange(y.size), size=6, replace=False)
X_train, y_train = X[training_indices], y[training_indices]

S = Kriging(name='kriging', seed=123, log_level=50, cod_type="norm")
S.fit(X_train, y_train)

mean_prediction, std_prediction, ei = S.predict(X, return_val="all")

std_prediction

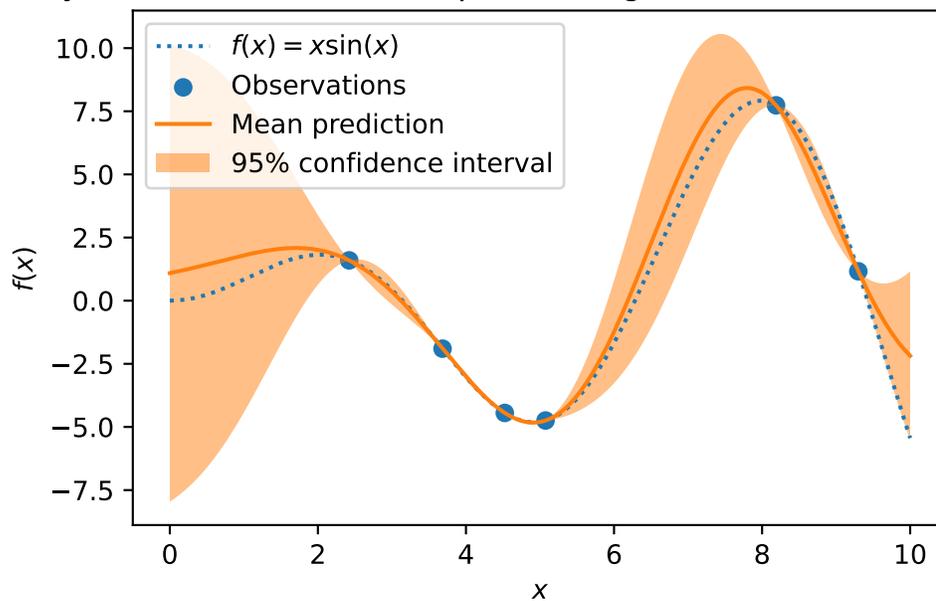
plt.plot(X, y, label=r"$f(x) = x \sin(x)$", linestyle="dotted")
plt.scatter(X_train, y_train, label="Observations")
plt.plot(X, mean_prediction, label="Mean prediction")
plt.fill_between(
```

```

X.ravel(),
mean_prediction - 1.96 * std_prediction,
mean_prediction + 1.96 * std_prediction,
alpha=0.5,
label=r"95% confidence interval",
)
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("spotPython Version: Gaussian process regression on noise-free dataset")

```

spotPython Version: Gaussian process regression on noise-free dataset



7.6 The Surrogate: Using scikit-learn models

Default is the internal kriging surrogate.

```
S_0 = Kriging(name='kriging', seed=123)
```

Models from `scikit-learn` can be selected, e.g., Gaussian Process:

```

# Needed for the sklearn surrogates:
from sklearn.gaussian_process import GaussianProcessRegressor
from sklearn.gaussian_process.kernels import RBF
from sklearn.tree import DecisionTreeRegressor
from sklearn.ensemble import RandomForestRegressor
from sklearn import linear_model
from sklearn import tree
import pandas as pd

kernel = 1 * RBF(length_scale=1.0, length_scale_bounds=(1e-2, 1e2))
S_GP = GaussianProcessRegressor(kernel=kernel, n_restarts_optimizer=9)

```

- and many more:

```

S_Tree = DecisionTreeRegressor(random_state=0)
S_LM = linear_model.LinearRegression()
S_Ridge = linear_model.Ridge()
S_RF = RandomForestRegressor(max_depth=2, random_state=0)

```

- The scikit-learn GP model S_GP is selected.

```
S = S_GP
```

```
isinstance(S, GaussianProcessRegressor)
```

True

```

from spotPython.fun.objectivefunctions import analytical
fun = analytical().fun_branin
lower = np.array([-5,-0])
upper = np.array([10,15])
design_control={"init_size": 5}
surrogate_control={
    "infill_criterion": None,
    "n_points": 1,
}
spot_GP = spot.Spot(fun=fun, lower = lower, upper= upper, surrogate=S,
    fun_evals = 15, noise = False, log_level = 50,
    design_control=design_control,
    surrogate_control=surrogate_control)

```


```
spot_GP.run()
```

```
<spotPython.spot.spot.Spot at 0x16b716710>
```

```
spot_GP.y
```

```
array([ 69.32459936, 152.38491454, 107.92560483,  24.51465459,  
       76.73500031,  86.3042577 , 128.1584069 ,   2.0848684 ,  
        5.65023665,   2.59664684,   1.94316346,   1.94350862,  
       12.83580716,  21.55925114,   1.94238966])
```

```
spot_GP.plot_progress()
```



```
spot_GP.print_results()
```

```
min y: 1.9423896629538913  
x0: 9.999794212305758  
x1: 2.984757626104978
```

```
[['x0', 9.999794212305758], ['x1', 2.984757626104978]]
```

7.7 Additional Examples

```

# Needed for the sklearn surrogates:
from sklearn.gaussian_process import GaussianProcessRegressor
from sklearn.gaussian_process.kernels import RBF
from sklearn.tree import DecisionTreeRegressor
from sklearn.ensemble import RandomForestRegressor
from sklearn import linear_model
from sklearn import tree
import pandas as pd

kernel = 1 * RBF(length_scale=1.0, length_scale_bounds=(1e-2, 1e2))
S_GP = GaussianProcessRegressor(kernel=kernel, n_restarts_optimizer=9)

from spotPython.build.kriging import Kriging
import numpy as np
import spotPython
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot

S_K = Kriging(name='kriging',
              seed=123,
              log_level=50,
              infill_criterion = "y",
              n_theta=1,
              noise=False,
              cod_type="norm")
fun = analytical().fun_sphere
lower = np.array([-1,-1])
upper = np.array([1,1])

design_control={"init_size": 10}
surrogate_control={
    "n_points": 1,
}
spot_S_K = spot.Spot(fun=fun,
                    lower = lower,
                    upper= upper,
                    surrogate=S_K,
                    fun_evals = 25,
                    noise = False,
                    log_level = 50,

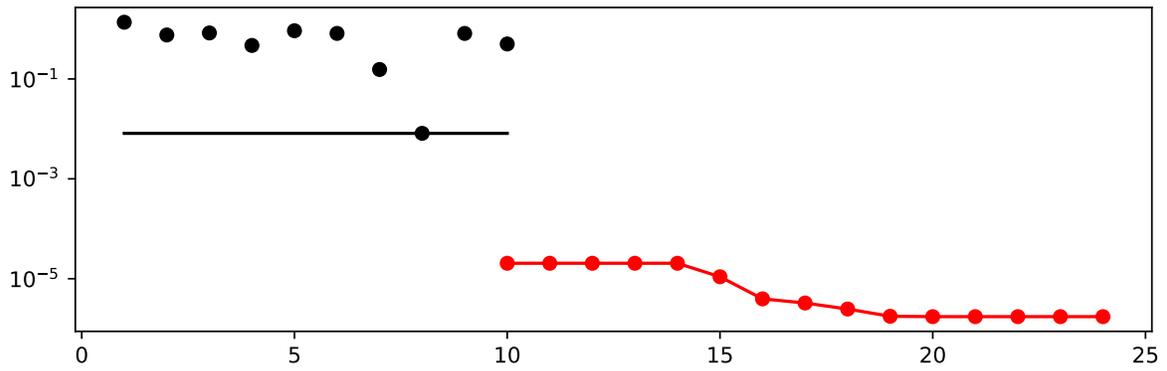
```

```
design_control=design_control,  
surrogate_control=surrogate_control)
```

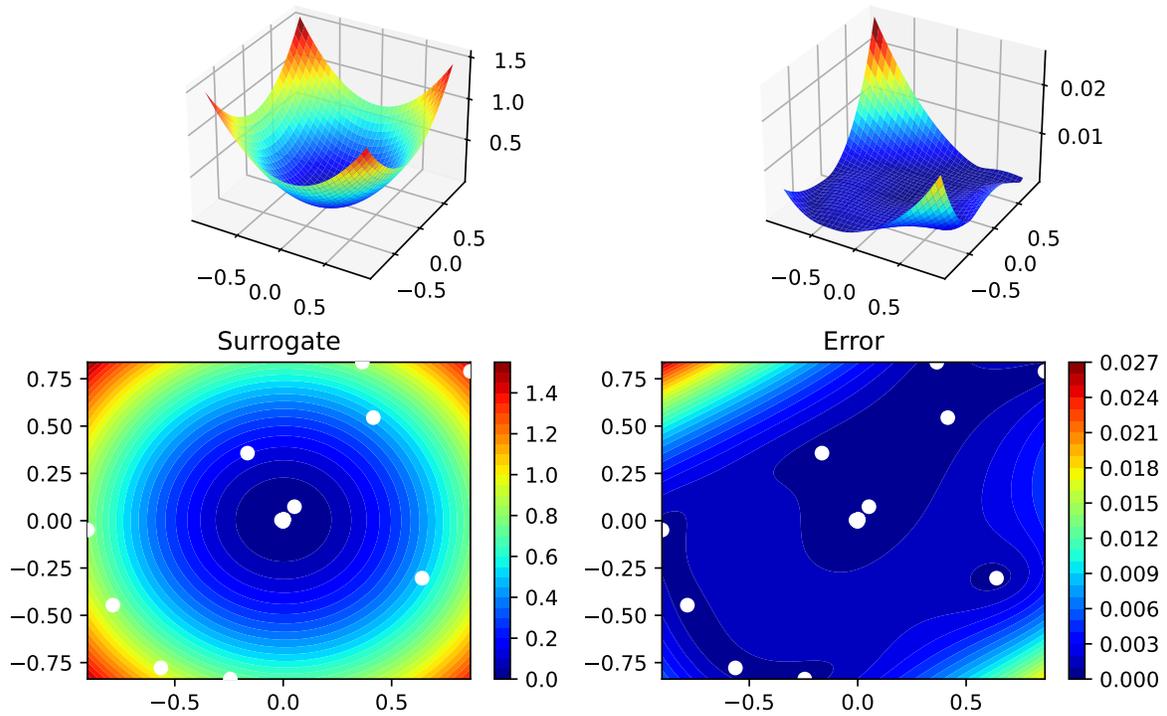
```
spot_S_K.run()
```

```
<spotPython.spot.spot.Spot at 0x16d33f700>
```

```
spot_S_K.plot_progress(log_y=True)
```



```
spot_S_K.surrogate.plot()
```



```
spot_S_K.print_results()
```

```
min y: 1.7395335905335862e-06
x0: -0.0013044072412622557
x1: 0.0001950777780173277
```

```
[['x0', -0.0013044072412622557], ['x1', 0.0001950777780173277]]
```

7.7.1 Optimize on Surrogate

7.7.2 Evaluate on Real Objective

7.7.3 Impute / Infill new Points

7.8 Tests

```
import numpy as np
from spotPython.spot import spot
from spotPython.fun.objectivefunctions import analytical

fun_sphere = analytical().fun_sphere
spot_1 = spot.Spot(
    fun=fun_sphere,
    lower=np.array([-1, -1]),
    upper=np.array([1, 1]),
    n_points = 2
)

# (S-2) Initial Design:
spot_1.X = spot_1.design.scipy_lhd(
    spot_1.design_control["init_size"], lower=spot_1.lower, upper=spot_1.upper
)
print(spot_1.X)

# (S-3): Eval initial design:
spot_1.y = spot_1.fun(spot_1.X)
print(spot_1.y)

spot_1.surrogate.fit(spot_1.X, spot_1.y)
X0 = spot_1.suggest_new_X()
print(X0)
assert X0.size == spot_1.n_points * spot_1.k
```

```
[[ 0.86352963  0.7892358 ]
 [-0.24407197 -0.83687436]
 [ 0.36481882  0.8375811 ]
 [ 0.415331    0.54468512]
 [-0.56395091 -0.77797854]
 [-0.90259409 -0.04899292]]
```

```

[-0.16484832  0.35724741]
[ 0.05170659  0.07401196]
[-0.78548145 -0.44638164]
[ 0.64017497 -0.30363301]]
[1.36857656  0.75992983  0.83463487  0.46918172  0.92329124  0.8170764
 0.15480068  0.00815134  0.81623768  0.502017  ]
[[0.00160553  0.00428429]
 [0.00160553  0.00428429]]

```

7.9 EI: The Famous Schonlau Example

```

X_train0 = np.array([1, 2, 3, 4, 12]).reshape(-1,1)
X_train = np.linspace(start=0, stop=10, num=5).reshape(-1, 1)

from spotPython.build.kriging import Kriging
import numpy as np
import matplotlib.pyplot as plt

X_train = np.array([1., 2., 3., 4., 12.]).reshape(-1,1)
y_train = np.array([0., -1.75, -2, -0.5, 5.])

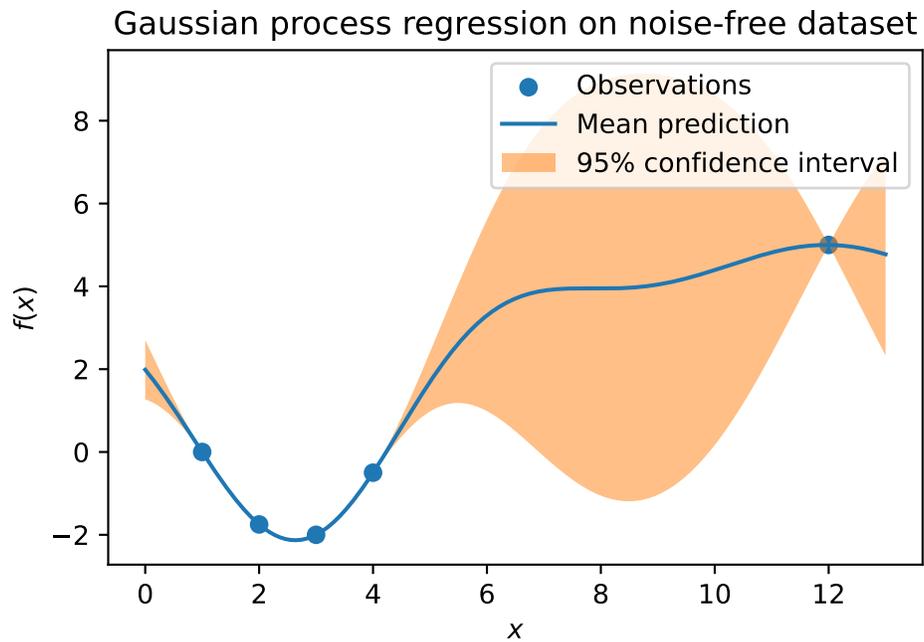
S = Kriging(name='kriging', seed=123, log_level=50, n_theta=1, noise=False, cod_type="non")
S.fit(X_train, y_train)

X = np.linspace(start=0, stop=13, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S.predict(X, return_val="all")

plt.scatter(X_train, y_train, label="Observations")
plt.plot(X, mean_prediction, label="Mean prediction")
if True:
    plt.fill_between(
        X.ravel(),
        mean_prediction - 2 * std_prediction,
        mean_prediction + 2 * std_prediction,
        alpha=0.5,
        label=r"95% confidence interval",
    )
plt.legend()
plt.xlabel("$x$")


```

```
plt.ylabel("$f(x)$")
_ = plt.title("Gaussian process regression on noise-free dataset")
```



```
#plt.plot(X, y, label=r"$f(x) = x \sin(x)$", linestyle="dotted")
# plt.scatter(X_train, y_train, label="Observations")
plt.plot(X, -ei, label="Expected Improvement")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Gaussian process regression on noise-free dataset")
```

Gaussian process regression on noise-free dataset



S.log

```
{'negLnLike': array([1.20788205]),  
 'theta': array([1.09276]),  
 'p': array([2.]),  
 'Lambda': array([None], dtype=object)}
```

7.10 EI: The Forrester Example

```
from spotPython.build.kriging import Kriging  
import numpy as np  
import matplotlib.pyplot as plt  
import spotPython  
from spotPython.fun.objectivefunctions import analytical  
from spotPython.spot import spot  
  
# exact x locations are unknown:  
X_train = np.array([0.0, 0.175, 0.225, 0.3, 0.35, 0.375, 0.5, 1]).reshape(-1,1)
```

```

fun = analytical().fun_forrester
fun_control = {"sigma": 1.0,
               "seed": 123}
y_train = fun(X_train, fun_control=fun_control)

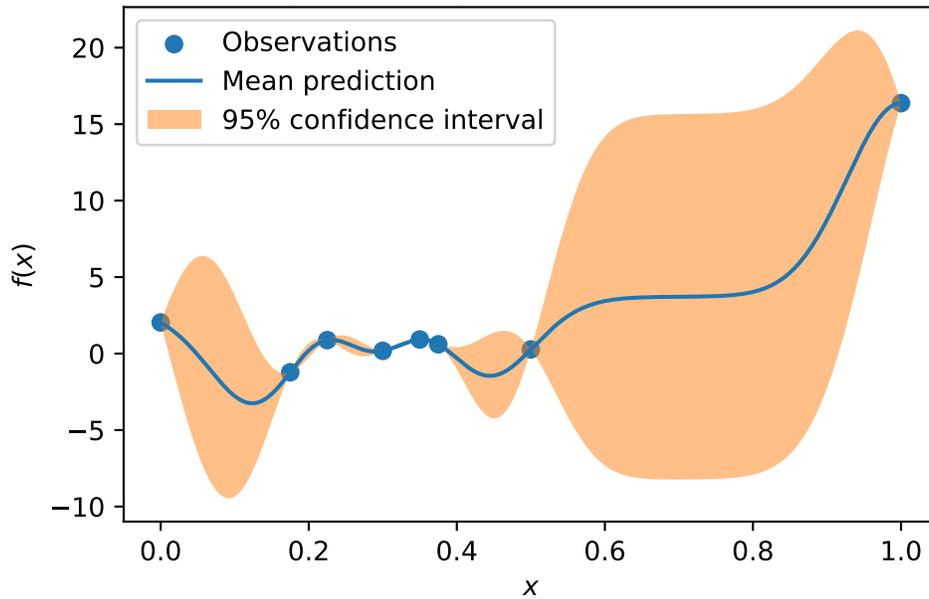
S = Kriging(name='kriging', seed=123, log_level=50, n_theta=1, noise=False, cod_type="non")
S.fit(X_train, y_train)

X = np.linspace(start=0, stop=1, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S.predict(X, return_val="all")

plt.scatter(X_train, y_train, label="Observations")
plt.plot(X, mean_prediction, label="Mean prediction")
if True:
    plt.fill_between(
        X.ravel(),
        mean_prediction - 2 * std_prediction,
        mean_prediction + 2 * std_prediction,
        alpha=0.5,
        label=r"95% confidence interval",
    )
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Gaussian process regression on noise-free dataset")

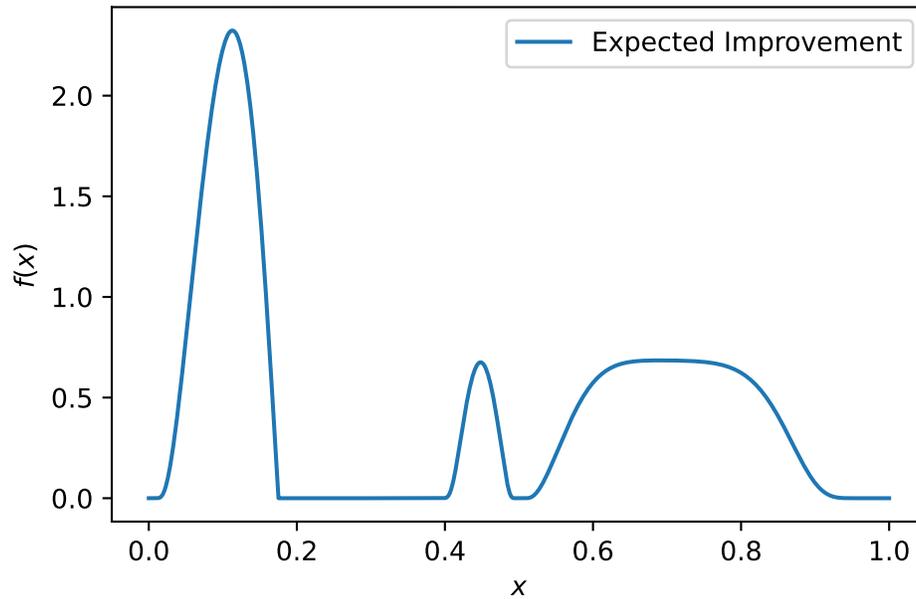
```

Gaussian process regression on noise-free dataset



```
#plt.plot(X, y, label=r"$f(x) = x \sin(x)$", linestyle="dotted")
# plt.scatter(X_train, y_train, label="Observations")
plt.plot(X, -ei, label="Expected Improvement")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Gaussian process regression on noise-free dataset")
```

Gaussian process regression on noise-free dataset



7.11 Noise

```
import numpy as np
import spotPython
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from spotPython.design.spacefilling import spacefilling
from spotPython.build.kriging import Kriging
import matplotlib.pyplot as plt

gen = spacefilling(1)
rng = np.random.RandomState(1)
lower = np.array([-10])
upper = np.array([10])
fun = analytical().fun_sphere
fun_control = {"sigma": 2,
               "seed": 125}
X = gen.scipy_lhd(10, lower=lower, upper = upper)
print(X)
y = fun(X, fun_control=fun_control)
```

```

print(y)
y.shape
X_train = X.reshape(-1,1)
y_train = y

S = Kriging(name='kriging',
            seed=123,
            log_level=50,
            n_theta=1,
            noise=False)
S.fit(X_train, y_train)

X_axis = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S.predict(X_axis, return_val="all")

#plt.plot(X, y, label=r"$f(x) = x \sin(x)$", linestyle="dotted")
plt.scatter(X_train, y_train, label="Observations")
#plt.plot(X, ei, label="Expected Improvement")
plt.plot(X_axis, mean_prediction, label="mue")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Sphere: Gaussian process regression on noisy dataset")

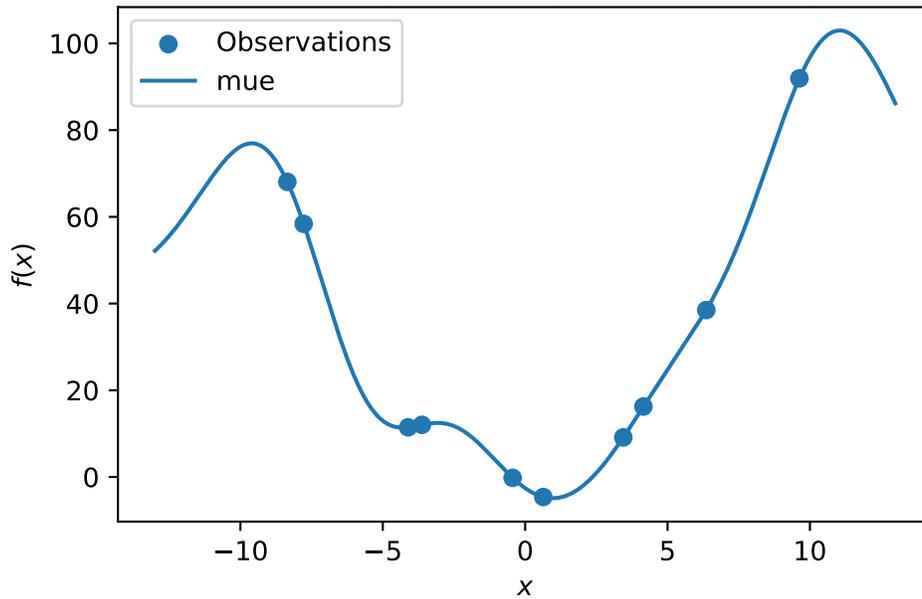
```

```

[[ 0.63529627]
 [-4.10764204]
 [-0.44071975]
 [ 9.63125638]
 [-8.3518118 ]
 [-3.62418901]
 [ 4.15331  ]
 [ 3.4468512 ]
 [ 6.36049088]
 [-7.77978539]]
[-4.61635371 11.44873209 -0.19988024 91.92791676 68.05926244 12.02926818
 16.2470957  9.12729929 38.4987029 58.38469104]

```

Sphere: Gaussian process regression on noisy dataset



S.log

```
{'negLnLike': array([24.69806131]),  
'theta': array([1.31023943]),  
'p': array([2.]),  
'Lambda': array([None], dtype=object)}
```

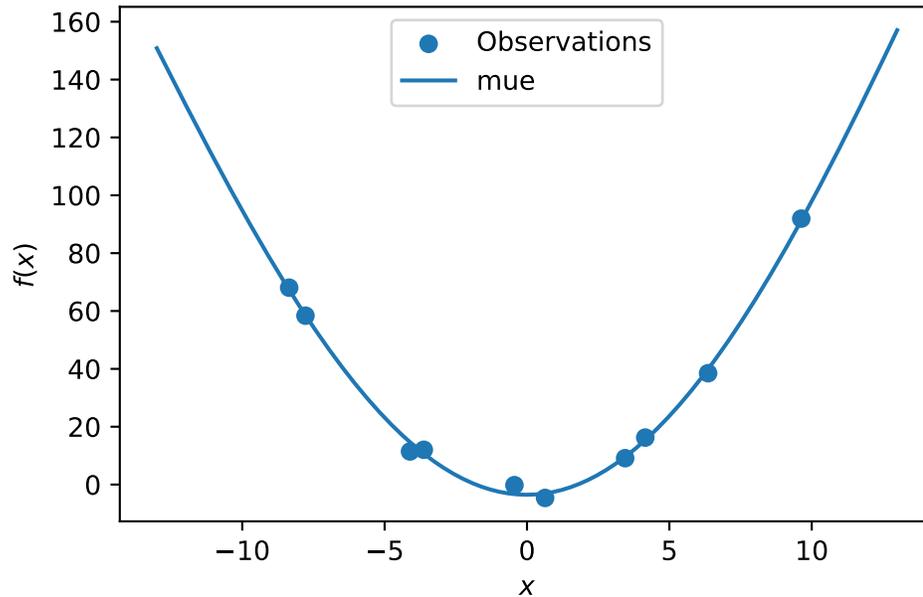
```
S = Kriging(name='kriging',  
           seed=123,  
           log_level=50,  
           n_theta=1,  
           noise=True)  
S.fit(X_train, y_train)
```

```
X_axis = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)  
mean_prediction, std_prediction, ei = S.predict(X_axis, return_val="all")
```

```
#plt.plot(X, y, label=r"$f(x) = x \sin(x)$", linestyle="dotted")  
plt.scatter(X_train, y_train, label="Observations")  
#plt.plot(X, ei, label="Expected Improvement")  
plt.plot(X_axis, mean_prediction, label="mue")
```

```
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Sphere: Gaussian process regression with nugget on noisy dataset")
```

Sphere: Gaussian process regression with nugget on noisy dataset



S.log

```
{'negLnLike': array([22.14095646]),
 'theta': array([-0.32527397]),
 'p': array([2.]),
 'Lambda': array([9.08815007e-05])}
```

7.12 Cubic Function

```
import numpy as np
import spotPython
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from spotPython.design.spacefilling import spacefilling
```

```

from spotPython.build.kriging import Kriging
import matplotlib.pyplot as plt

gen = spacefilling(1)
rng = np.random.RandomState(1)
lower = np.array([-10])
upper = np.array([10])
fun = analytical().fun_cubed
fun_control = {"sigma": 10,
               "seed": 123}

X = gen.scipy_lhd(10, lower=lower, upper = upper)
print(X)
y = fun(X, fun_control=fun_control)
print(y)
y.shape
X_train = X.reshape(-1,1)
y_train = y

S = Kriging(name='kriging', seed=123, log_level=50, n_theta=1, noise=False)
S.fit(X_train, y_train)

X_axis = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S.predict(X_axis, return_val="all")

plt.scatter(X_train, y_train, label="Observations")
#plt.plot(X, ei, label="Expected Improvement")
plt.plot(X_axis, mean_prediction, label="mue")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Cubed: Gaussian process regression on noisy dataset")


```

```

[[ 0.63529627]
 [-4.10764204]
 [-0.44071975]
 [ 9.63125638]
 [-8.3518118 ]
 [-3.62418901]
 [ 4.15331  ]
 [ 3.4468512 ]
 [ 6.36049088]

```

```
[-7.77978539]]
[ -9.63480707 -72.98497325  12.7936499  895.34567477 -573.35961837
 -41.83176425  65.27989461  46.37081417  254.1530734 -474.09587355]
```

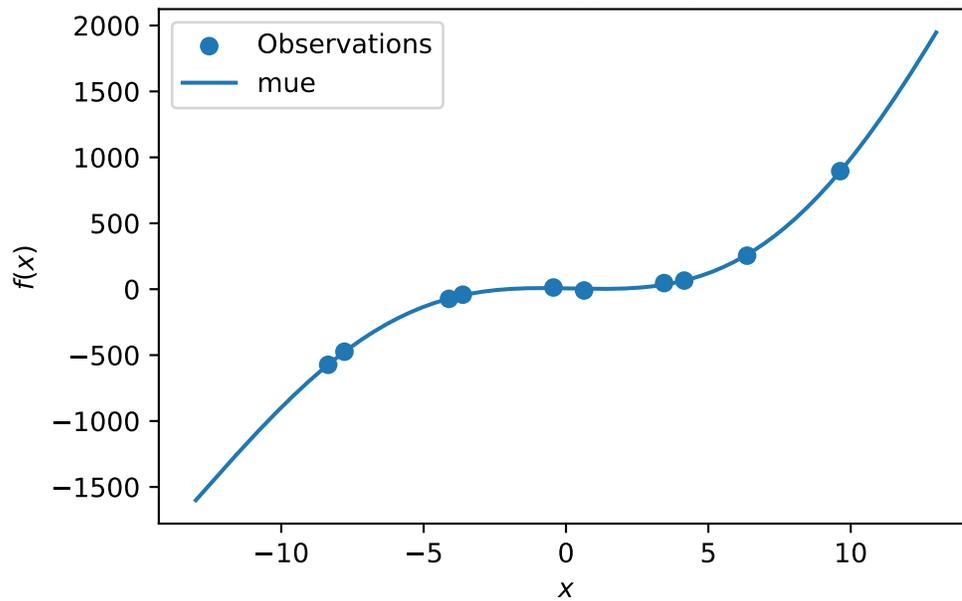


```
S = Kriging(name='kriging', seed=123, log_level=0, n_theta=1, noise=True)
S.fit(X_train, y_train)

X_axis = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S.predict(X_axis, return_val="all")

plt.scatter(X_train, y_train, label="Observations")
#plt.plot(X, ei, label="Expected Improvement")
plt.plot(X_axis, mean_prediction, label="mue")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Cubed: Gaussian process with nugget regression on noisy dataset")
```

Cubed: Gaussian process with nugget regression on noisy dataset



```
import numpy as np
import spotPython
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from spotPython.design.spacefilling import spacefilling
from spotPython.build.kriging import Kriging
import matplotlib.pyplot as plt

gen = spacefilling(1)
rng = np.random.RandomState(1)
lower = np.array([-10])
upper = np.array([10])
fun = analytical().fun_runge
fun_control = {"sigma": 0.25,
              "seed": 123}

X = gen.scipy_lhd(10, lower=lower, upper = upper)
print(X)
y = fun(X, fun_control=fun_control)
print(y)
y.shape
```

```

X_train = X.reshape(-1,1)
y_train = y

S = Kriging(name='kriging', seed=123, log_level=50, n_theta=1, noise=False)
S.fit(X_train, y_train)

X_axis = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S.predict(X_axis, return_val="all")

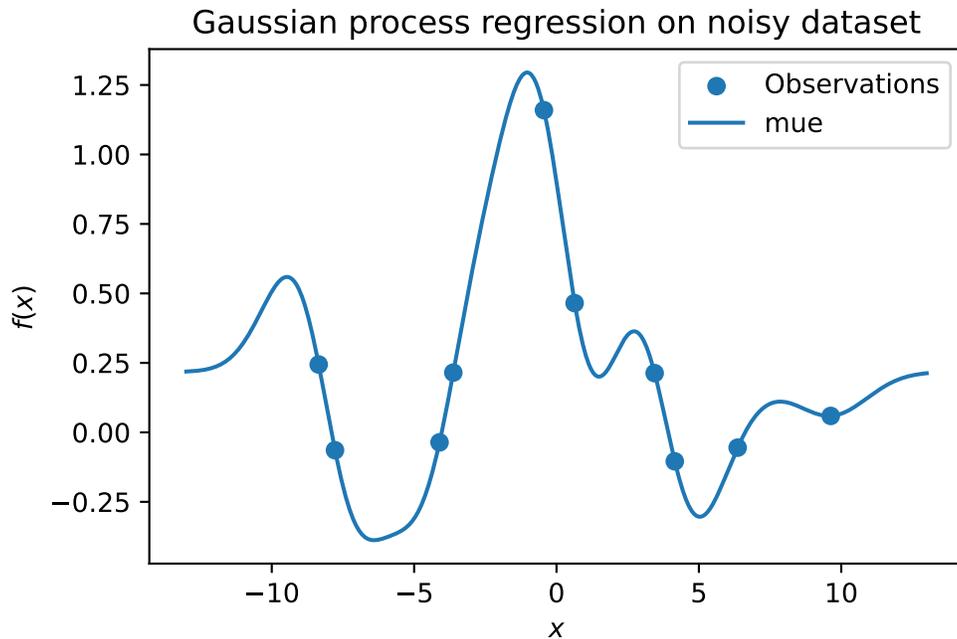
plt.scatter(X_train, y_train, label="Observations")
#plt.plot(X, ei, label="Expected Improvement")
plt.plot(X_axis, mean_prediction, label="mue")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Gaussian process regression on noisy dataset")

```

```

[[ 0.63529627]
 [-4.10764204]
 [-0.44071975]
 [ 9.63125638]
 [-8.3518118 ]
 [-3.62418901]
 [ 4.15331  ]
 [ 3.4468512 ]
 [ 6.36049088]
 [-7.77978539]]
[ 0.46517267 -0.03599548  1.15933822  0.05915901  0.24419145  0.21502359
 -0.10432134  0.21312309 -0.05502681 -0.06434374]

```



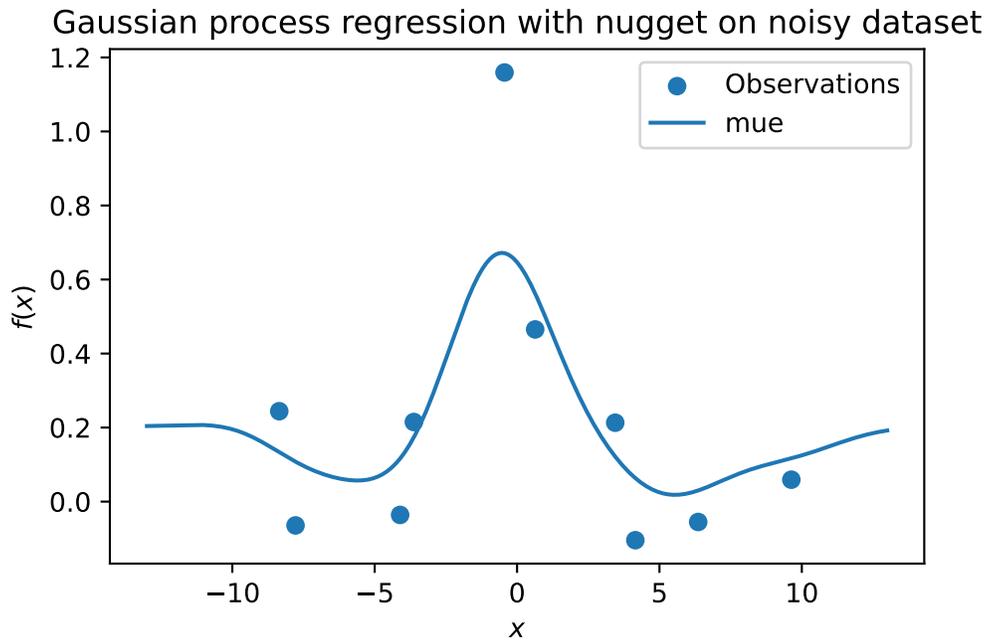
```

S = Kriging(name='kriging',
            seed=123,
            log_level=50,
            n_theta=1,
            noise=True)
S.fit(X_train, y_train)

X_axis = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S.predict(X_axis, return_val="all")

plt.scatter(X_train, y_train, label="Observations")
#plt.plot(X, ei, label="Expected Improvement")
plt.plot(X_axis, mean_prediction, label="mue")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Gaussian process regression with nugget on noisy dataset")

```



7.13 Factors

```
["num"] * 3
```

```
['num', 'num', 'num']
```

```
from spotPython.design.spacefilling import spacefilling
from spotPython.build.kriging import Kriging
from spotPython.fun.objectivefunctions import analytical
import numpy as np
```

```
gen = spacefilling(2)
n = 30
rng = np.random.RandomState(1)
lower = np.array([-5,-0])
upper = np.array([10,15])
fun = analytical().fun_branin_factor
#fun = analytical(sigma=0).fun_sphere
```

```

X0 = gen.scipy_lhd(n, lower=lower, upper = upper)
X1 = np.random.randint(low=1, high=3, size=(n,))
X = np.c_[X0, X1]
y = fun(X)
S = Kriging(name='kriging', seed=123, log_level=50, n_theta=3, noise=False, var_type=["nu
S.fit(X, y)
Sf = Kriging(name='kriging', seed=123, log_level=50, n_theta=3, noise=False, var_type=["n
Sf.fit(X, y)
n = 50
X0 = gen.scipy_lhd(n, lower=lower, upper = upper)
X1 = np.random.randint(low=1, high=3, size=(n,))
X = np.c_[X0, X1]
y = fun(X)
s=np.sum(np.abs(S.predict(X)[0] - y))
sf=np.sum(np.abs(Sf.predict(X)[0] - y))
sf - s

```

-193.01351282220276

```
# vars(S)
```

```
# vars(Sf)
```

8 Hyperparameter Tuning and Noise

This chapter demonstrates how noisy functions can be handled by Spot.

8.1 Example: Spot and the Noisy Sphere Function

```
import numpy as np
from math import inf
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from scipy.optimize import shgo
from scipy.optimize import direct
from scipy.optimize import differential_evolution
import matplotlib.pyplot as plt
import os
import copy
import socket
from datetime import datetime
from dateutil.tz import tzlocal

start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '10-sklearn' + "_" + HOSTNAME + "_" + str(start_time).split(".", 1)[0].r
experiment_name = experiment_name.replace(':', '-')
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')
```

10-sklearn_bartz09_2023-06-19_02-25-15

8.1.1 The Objective Function: Noisy Sphere

- The spotPython package provides several classes of objective functions.

- We will use an analytical objective function with noise, i.e., a function that can be described by a (closed) formula:

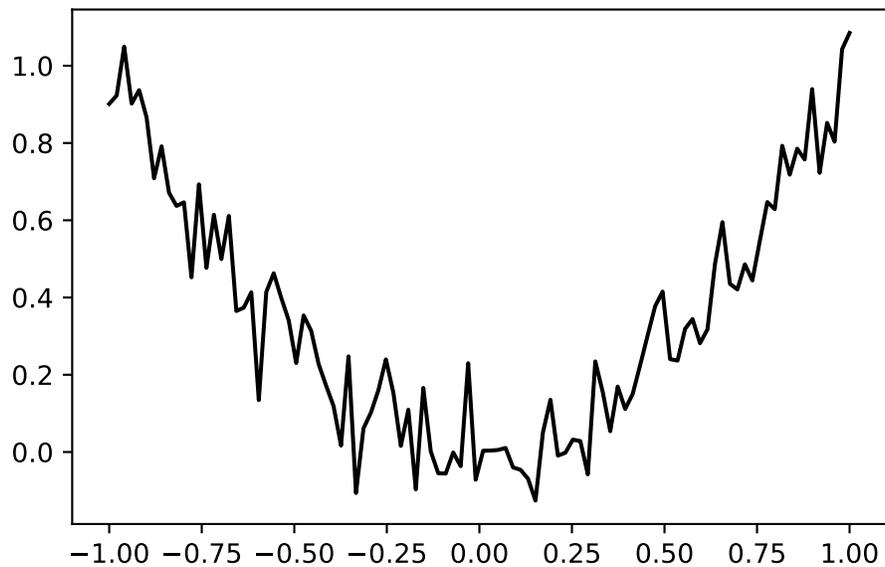
$$f(x) = x^2 + \epsilon$$

- Since `sigma` is set to 0.1, noise is added to the function:

```
fun = analytical().fun_sphere
fun_control = {"sigma": 0.1,
              "seed": 123}
```

- A plot illustrates the noise:

```
x = np.linspace(-1,1,100).reshape(-1,1)
y = fun(x, fun_control=fun_control)
plt.figure()
plt.plot(x,y, "k")
plt.show()
```



Spot is adopted as follows to cope with noisy functions:

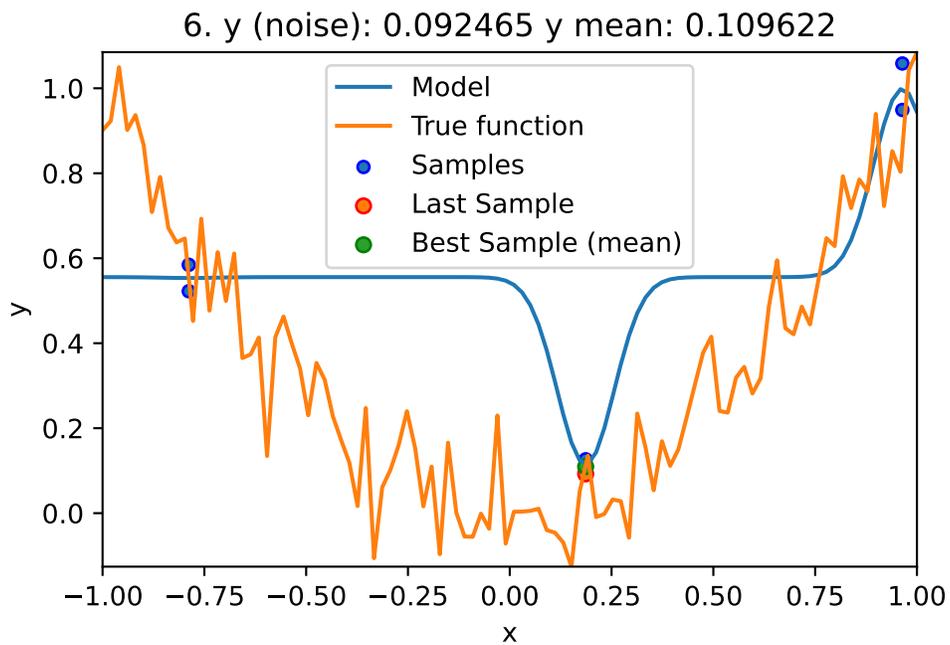
1. `fun_repeats` is set to a value larger than 1 (here: 2)
2. `noise` is set to `true`. Therefore, a nugget (`Lambda`) term is added to the correlation matrix
3. `init_size` (of the `design_control` dictionary) is set to a value larger than 1 (here: 2)

```

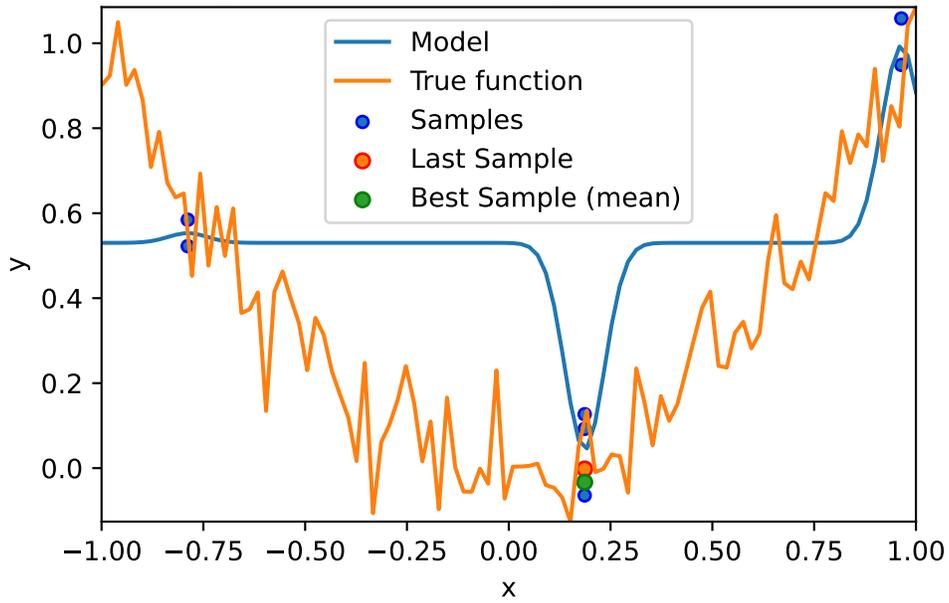
spot_1_noisy = spot.Spot(fun=fun,
    lower = np.array([-1]),
    upper = np.array([1]),
    fun_evals = 10,
    fun_repeats = 2,
    noise = True,
    seed=123,
    show_models=True,
    fun_control = fun_control,
    design_control={"init_size": 3,
        "repeats": 2},
    surrogate_control={"noise": True})

```

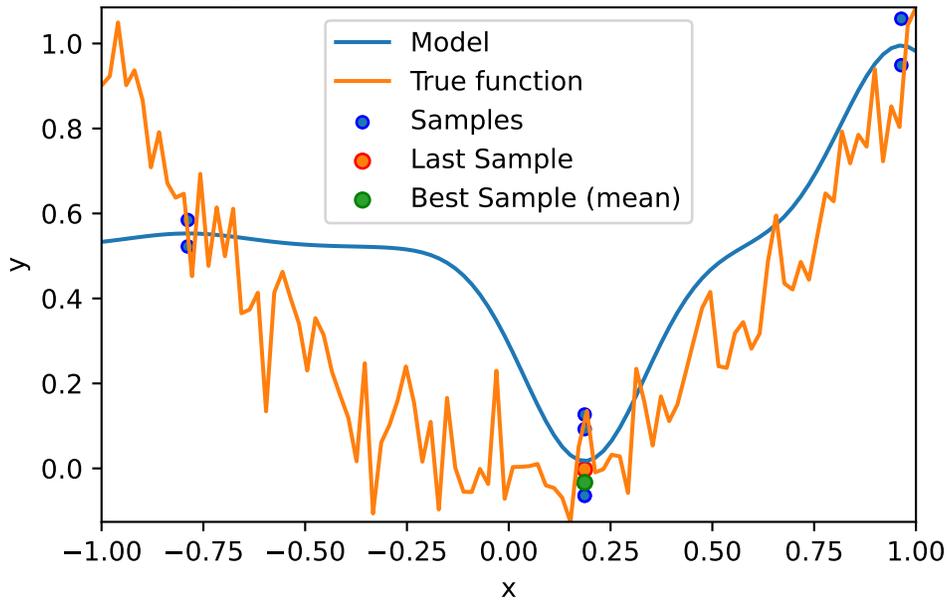
```
spot_1_noisy.run()
```



8. y (noise): -0.064157 y mean: -0.03309



10. y (noise): -0.064157 y mean: -0.03309



<spotPython.spot.spot.Spot at 0x16a68b880>

8.2 Print the Results

```
spot_1_noisy.print_results()
```

```
min y: -0.06415721594238855
x0: 0.18642671238960512
min mean y: -0.03309048099839016
x0: 0.18642671238960512
```

```
[['x0', 0.18642671238960512], ['x0', 0.18642671238960512]]
```

```
spot_1_noisy.plot_progress(log_y=False,
                             filename="./figures/" + experiment_name+"_progress.png")
```

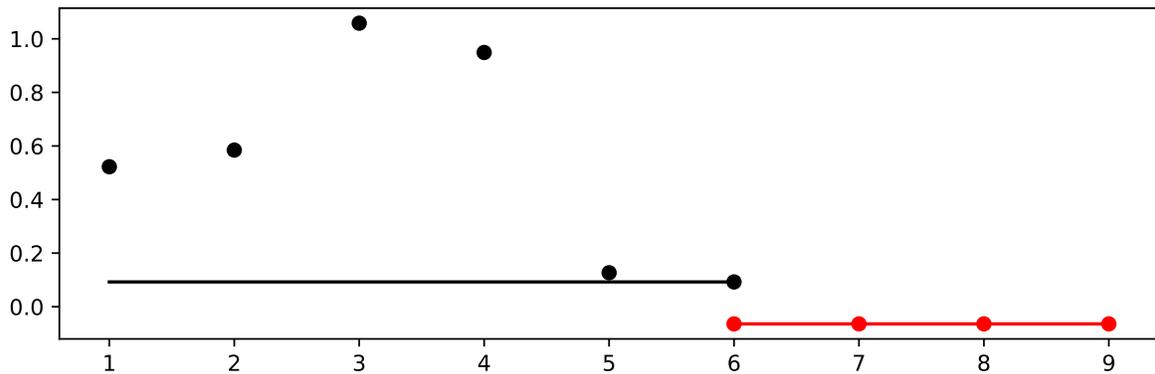


Figure 8.1: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

8.3 Noise and Surrogates: The Nugget Effect

8.3.1 The Noisy Sphere

8.3.1.1 The Data

- We prepare some data first:

```

import numpy as np
import spotPython
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from spotPython.design.spacefilling import spacefilling
from spotPython.build.kriging import Kriging
import matplotlib.pyplot as plt

gen = spacefilling(1)
rng = np.random.RandomState(1)
lower = np.array([-10])
upper = np.array([10])
fun = analytical().fun_sphere
fun_control = {"sigma": 2,
               "seed": 125}
X = gen.scipy_lhd(10, lower=lower, upper = upper)
y = fun(X, fun_control=fun_control)
X_train = X.reshape(-1,1)
y_train = y

```

- A surrogate without nugget is fitted to these data:

```

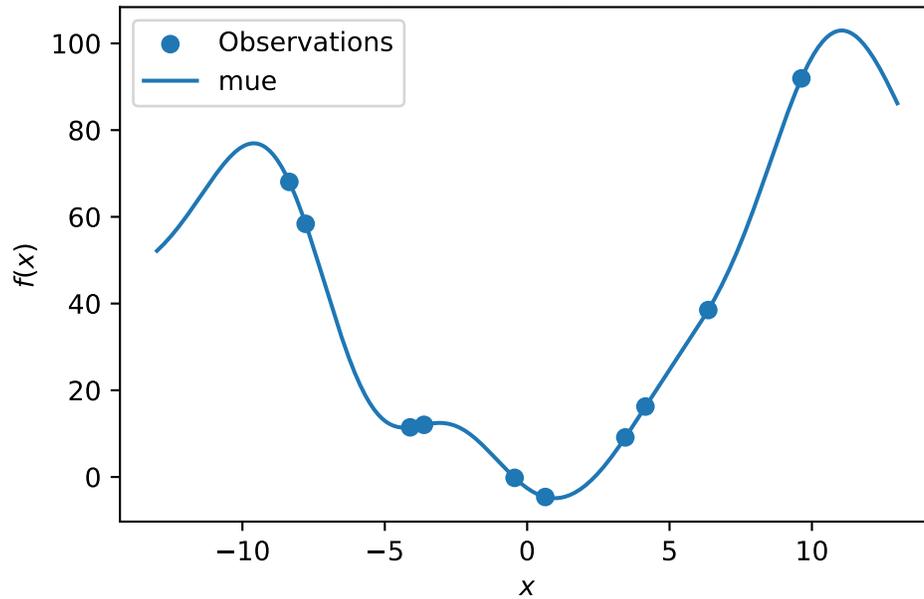
S = Kriging(name='kriging',
            seed=123,
            log_level=50,
            n_theta=1,
            noise=False)
S.fit(X_train, y_train)

X_axis = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S.predict(X_axis, return_val="all")

plt.scatter(X_train, y_train, label="Observations")
plt.plot(X_axis, mean_prediction, label="mue")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Sphere: Gaussian process regression on noisy dataset")

```

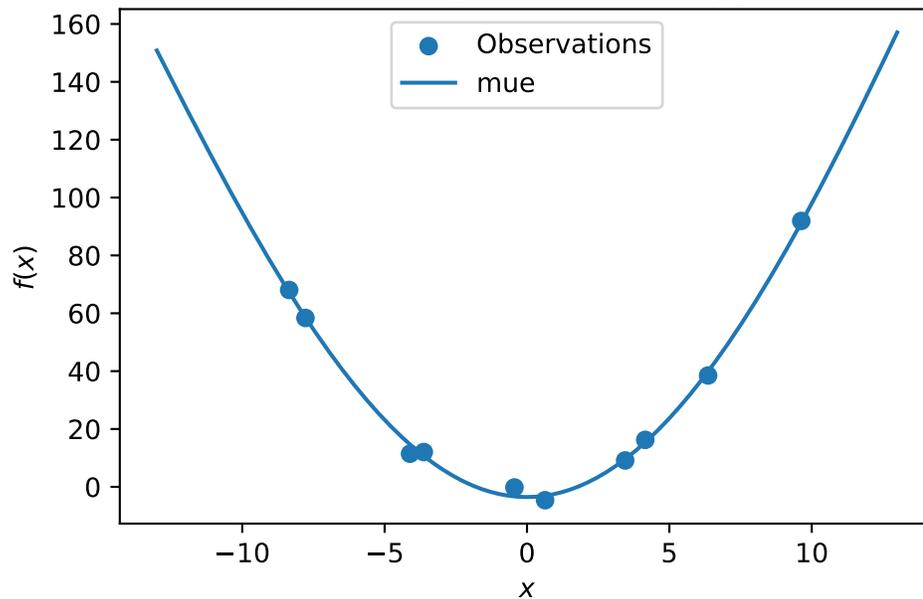
Sphere: Gaussian process regression on noisy dataset



- In comparison to the surrogate without nugget, we fit a surrogate with nugget to the data:

```
S_nug = Kriging(name='kriging',
                seed=123,
                log_level=50,
                n_theta=1,
                noise=True)
S_nug.fit(X_train, y_train)
X_axis = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S_nug.predict(X_axis, return_val="all")
plt.scatter(X_train, y_train, label="Observations")
plt.plot(X_axis, mean_prediction, label="mue")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Sphere: Gaussian process regression with nugget on noisy dataset")
```

Sphere: Gaussian process regression with nugget on noisy dataset



- The value of the nugget term can be extracted from the model as follows:

```
S.Lambda
```

```
S_nug.Lambda
```

```
9.088150066416743e-05
```

- We see:
 - the first model `S` has no nugget,
 - whereas the second model has a nugget value (`Lambda`) larger than zero.

8.4 Exercises

8.4.1 Noisy `fun_cubed`

- Analyse the effect of noise on the `fun_cubed` function with the following settings:

```
fun = analytical().fun_cubed  
fun_control = {"sigma": 10,
```

```
        "seed": 123}
lower = np.array([-10])
upper = np.array([10])
```

8.4.2 fun_runge

- Analyse the effect of noise on the `fun_runge` function with the following settings:

```
lower = np.array([-10])
upper = np.array([10])
fun = analytical().fun_runge
fun_control = {"sigma": 0.25,
               "seed": 123}
```

8.4.3 fun_forrester

- Analyse the effect of noise on the `fun_forrester` function with the following settings:

```
lower = np.array([0])
upper = np.array([1])
fun = analytical().fun_forrester
fun_control = {"sigma": 5,
               "seed": 123}
```

8.4.4 fun_xsin

- Analyse the effect of noise on the `fun_xsin` function with the following settings:

```
lower = np.array([-1.])
upper = np.array([1.])
fun = analytical().fun_xsin
fun_control = {"sigma": 0.5,
               "seed": 123}
```

9 Handling Noise: Optimal Computational Budget Allocation in Spot

This notebook demonstrates how noisy functions can be handled with OCBA by Spot.

9.1 Example: Spot, OCBA, and the Noisy Sphere Function

```
import numpy as np
from math import inf
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from scipy.optimize import shgo
from scipy.optimize import direct
from scipy.optimize import differential_evolution
import matplotlib.pyplot as plt
```

9.1.1 The Objective Function: Noisy Sphere

The `spotPython` package provides several classes of objective functions. We will use an analytical objective function with noise, i.e., a function that can be described by a (closed) formula:

$$f(x) = x^2 + \epsilon$$

Since `sigma` is set to 0.1, noise is added to the function:


```
fun = analytical().fun_sphere
fun_control = {"sigma": 0.1,
              "seed": 123}
```

A plot illustrates the noise:

```

x = np.linspace(-1,1,100).reshape(-1,1)
y = fun(x, fun_control=fun_control)
plt.figure()
plt.plot(x,y, "k")
plt.show()

```



Spot is adopted as follows to cope with noisy functions:

1. `fun_repeats` is set to a value larger than 1 (here: 2)
2. `noise` is set to `true`. Therefore, a nugget (`Lambda`) term is added to the correlation matrix
3. `init_size` (of the `design_control` dictionary) is set to a value larger than 1 (here: 2)

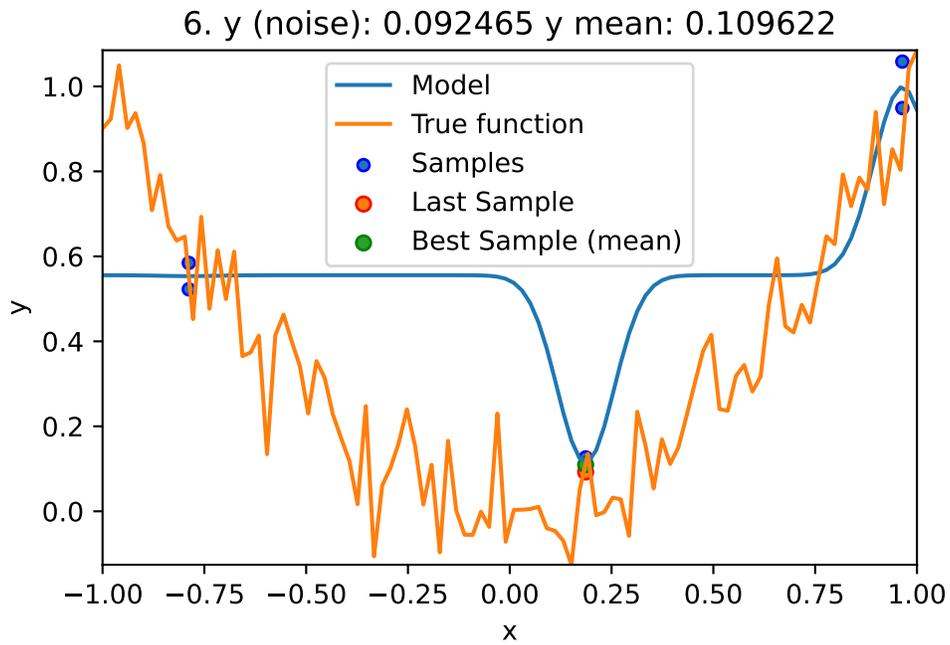
```

spot_1_noisy = spot.Spot(fun=fun,
                        lower = np.array([-1]),
                        upper = np.array([1]),
                        fun_evals = 50,
                        fun_repeats = 2,
                        infill_criterion="ei",
                        noise = True,
                        tolerance_x=0.0,
                        ocba_delta = 1,
                        seed=123,

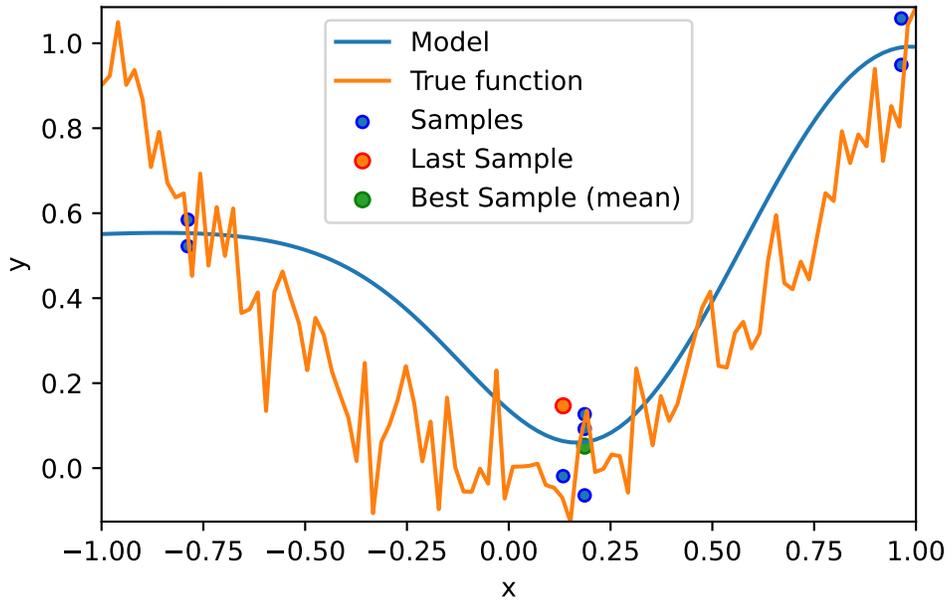
```

```
show_models=True,  
fun_control = fun_control,  
design_control={"init_size": 3,  
              "repeats": 2},  
surrogate_control={"noise": True})
```

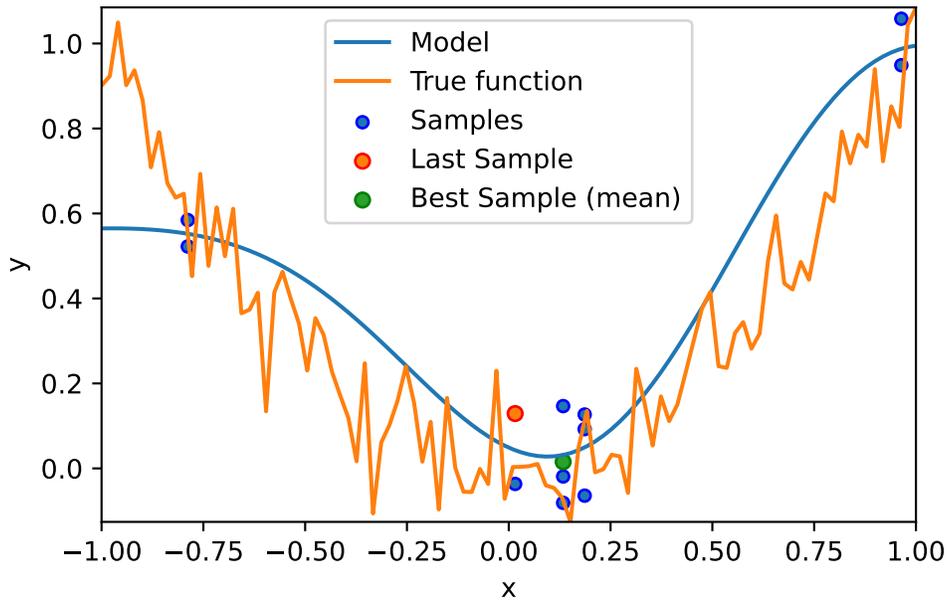
```
spot_1_noisy.run()
```



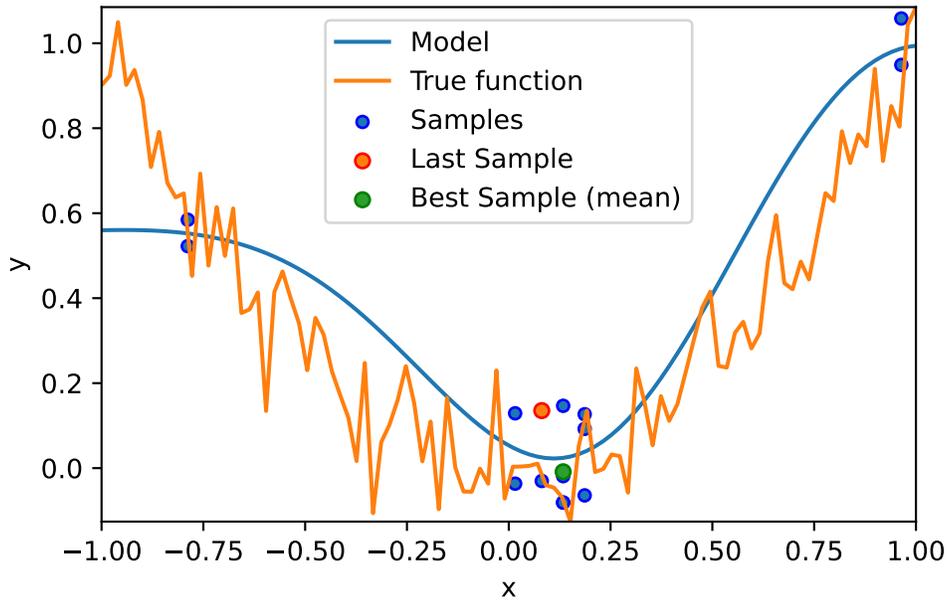
9. y (noise): -0.064157 y mean: 0.051695



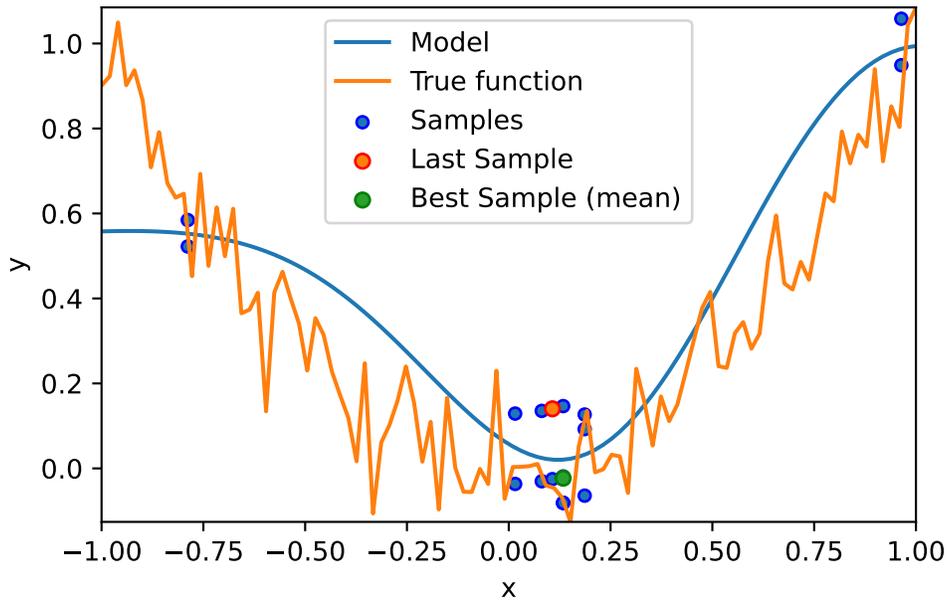
12. y (noise): -0.081063 y mean: 0.01555





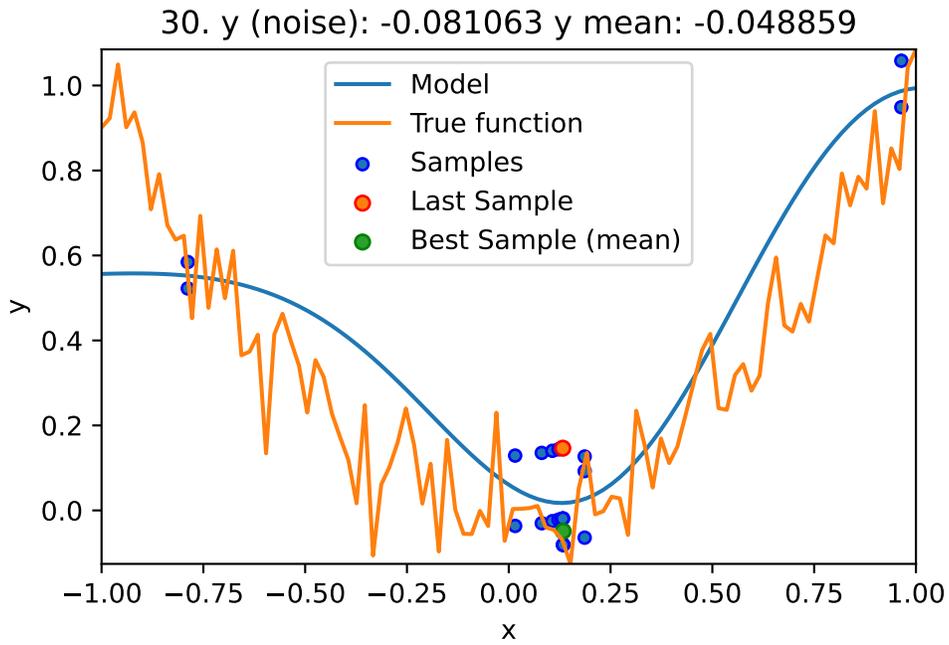
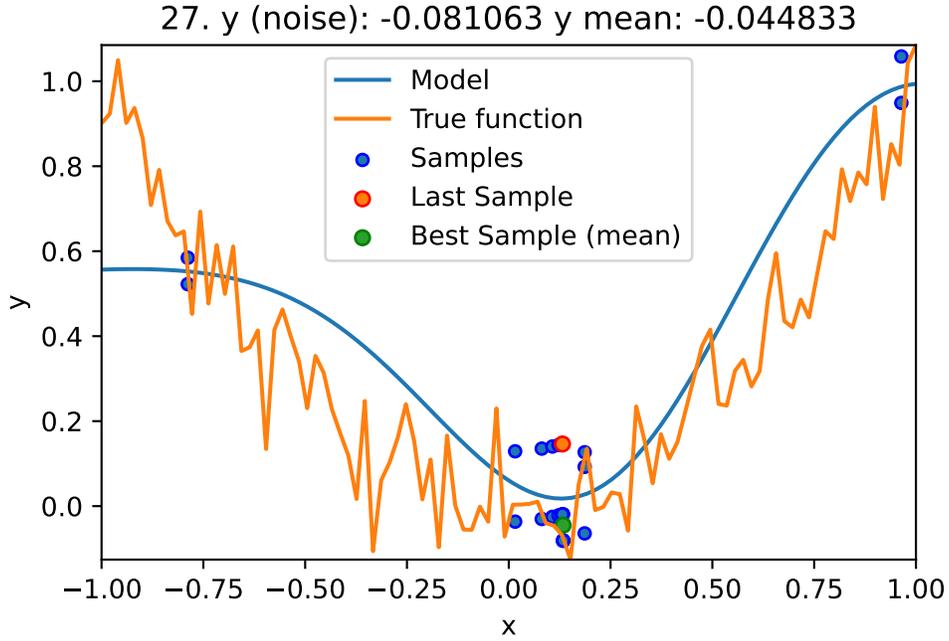
15. y (noise): -0.081063 y mean: -0.008604



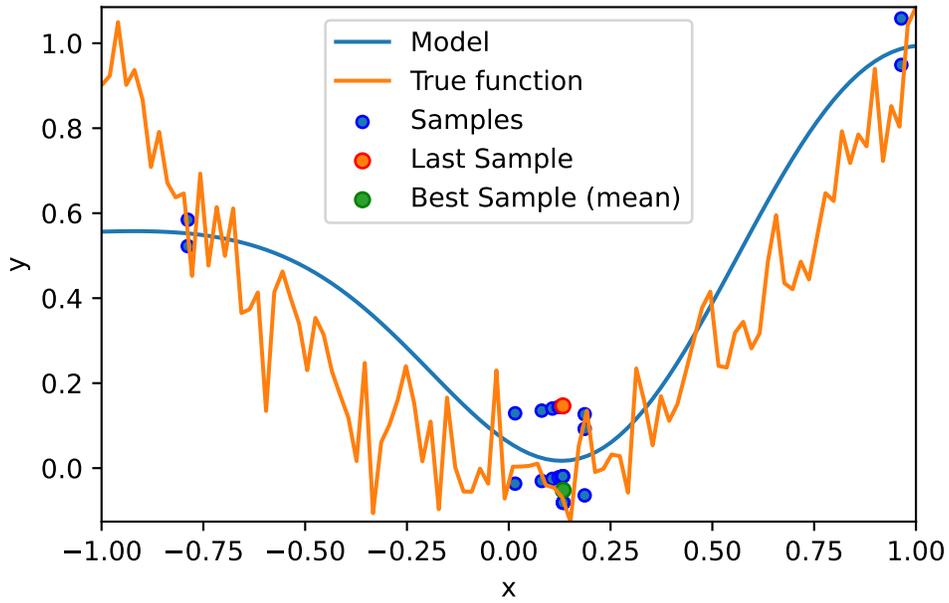
18. y (noise): -0.081063 y mean: -0.023096



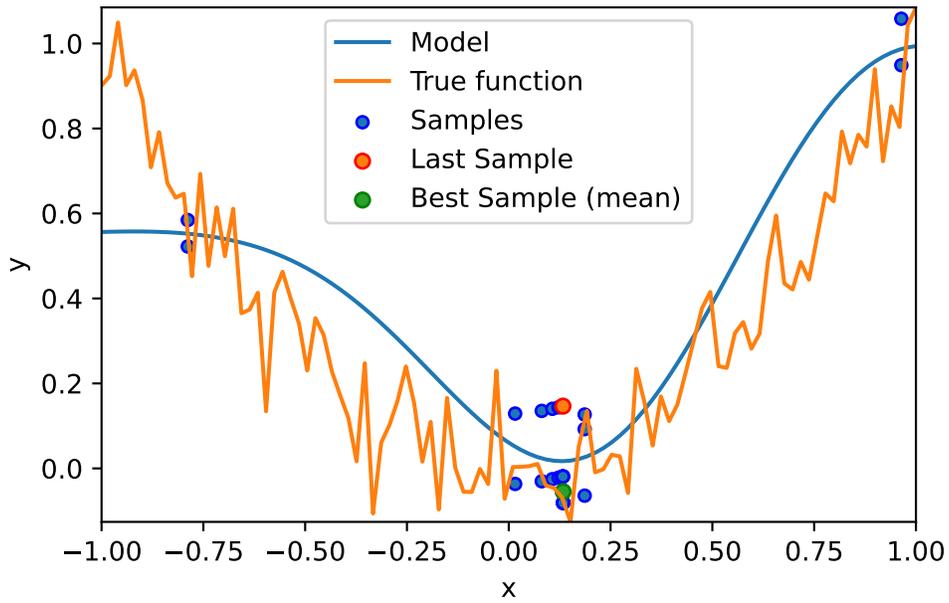




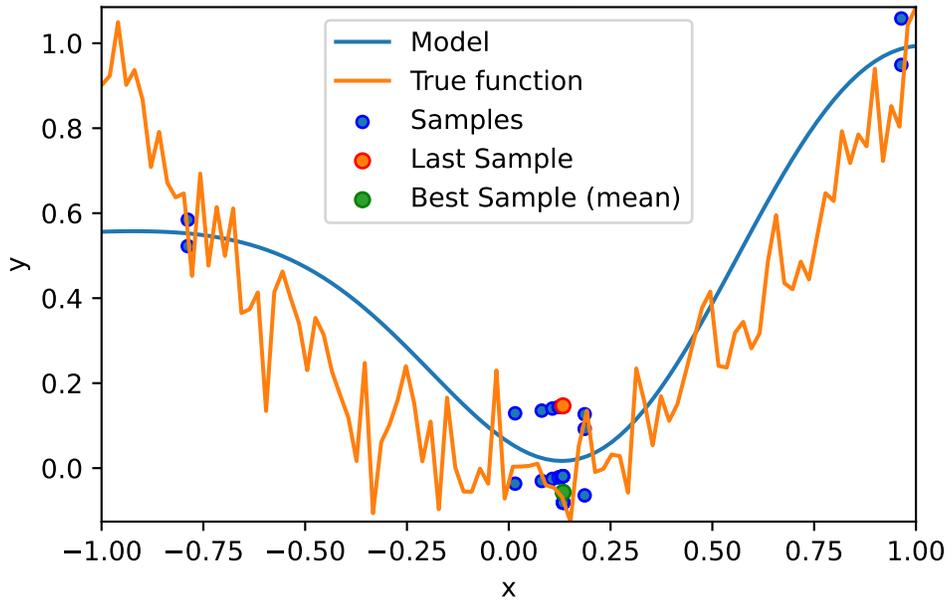
33. y (noise): -0.081063 y mean: -0.052079



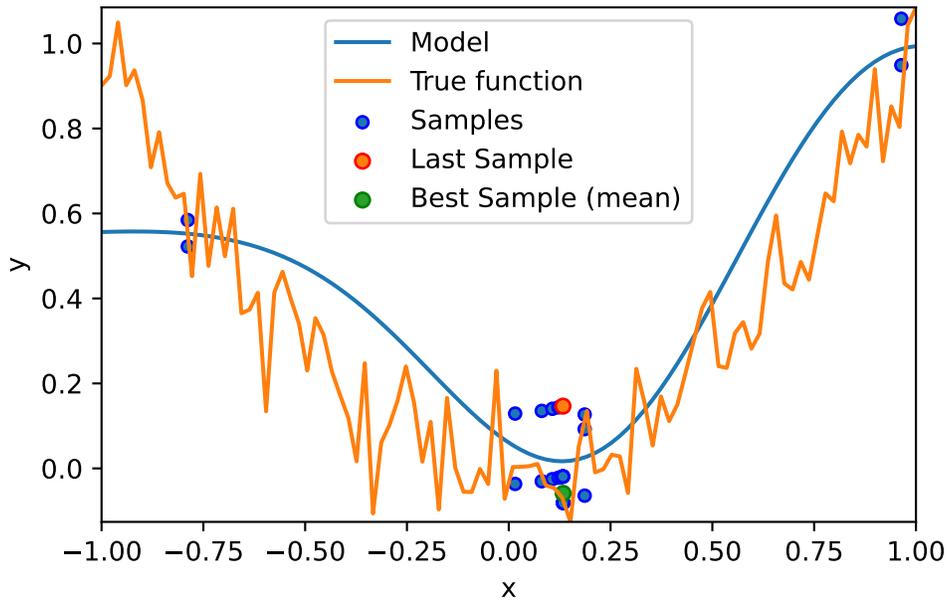
36. y (noise): -0.081063 y mean: -0.054714




39. y (noise): -0.081063 y mean: -0.05691




42. y (noise): -0.081063 y mean: -0.058768




45. y (noise): -0.081063 y mean: -0.06036



48. y (noise): -0.081063 y mean: -0.061741





```
<spotPython.spot.spot.Spot at 0x16c0eb700>
```

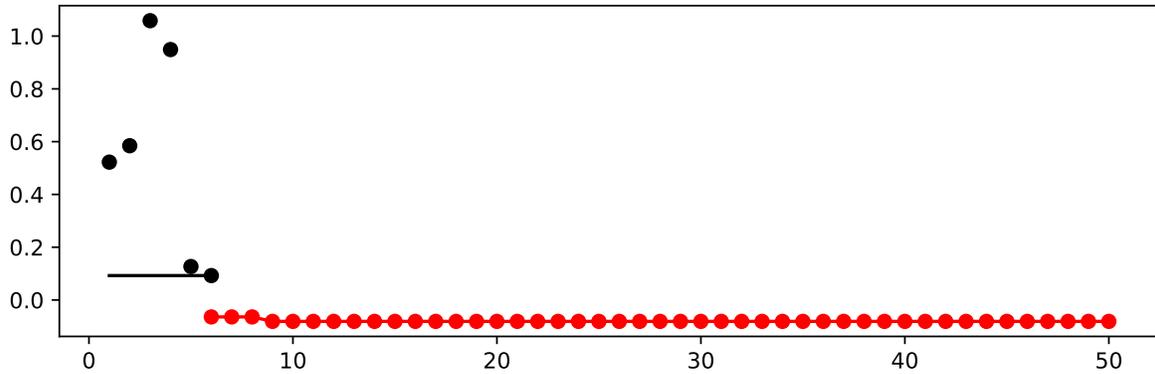
9.2 Print the Results

```
spot_1_noisy.print_results()
```

```
min y: -0.08106318979661208
x0: 0.1335999447536301
min mean y: -0.06294830660588041
x0: 0.1335999447536301
```

```
[['x0', 0.1335999447536301], ['x0', 0.1335999447536301]]
```

```
spot_1_noisy.plot_progress(log_y=False)
```



9.3 Noise and Surrogates: The Nugget Effect

9.3.1 The Noisy Sphere

9.3.1.1 The Data

We prepare some data first:

```
import numpy as np
import spotPython
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from spotPython.design.spacefilling import spacefilling
from spotPython.build.kriging import Kriging
import matplotlib.pyplot as plt

gen = spacefilling(1)
rng = np.random.RandomState(1)
lower = np.array([-10])
upper = np.array([10])
fun = analytical().fun_sphere
fun_control = {"sigma": 2,
               "seed": 125}
X = gen.scipy_lhd(10, lower=lower, upper = upper)
y = fun(X, fun_control=fun_control)
X_train = X.reshape(-1,1)
y_train = y
```

A surrogate without nugget is fitted to these data:


```

S = Kriging(name='kriging',
            seed=123,
            log_level=50,
            n_theta=1,
            noise=False)
S.fit(X_train, y_train)

X_axis = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S.predict(X_axis, return_val="all")

plt.scatter(X_train, y_train, label="Observations")
plt.plot(X_axis, mean_prediction, label="mue")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Sphere: Gaussian process regression on noisy dataset")

```




In comparison to the surrogate without nugget, we fit a surrogate with nugget to the data:

```

S_nug = Kriging(name='kriging',
                seed=123,
                log_level=50,
                n_theta=1,
                noise=True)
S_nug.fit(X_train, y_train)
X_axis = np.linspace(start=-13, stop=13, num=1000).reshape(-1, 1)
mean_prediction, std_prediction, ei = S_nug.predict(X_axis, return_val="all")
plt.scatter(X_train, y_train, label="Observations")
plt.plot(X_axis, mean_prediction, label="mue")
plt.legend()
plt.xlabel("$x$")
plt.ylabel("$f(x)$")
_ = plt.title("Sphere: Gaussian process regression with nugget on noisy dataset")

```

Sphere: Gaussian process regression with nugget on noisy dataset



The value of the nugget term can be extracted from the model as follows:

```
S.Lambda
```

```
S_nug.Lambda
```

9.088150066416743e-05

We see:

- the first model S has no nugget,
- whereas the second model has a nugget value (Lambda) larger than zero.

9.4 Exercises

9.4.1 Noisy fun_cubed

Analyse the effect of noise on the `fun_cubed` function with the following settings:

```
fun = analytical().fun_cubed
fun_control = {"sigma": 10,
              "seed": 123}
lower = np.array([-10])
upper = np.array([10])
```

9.4.2 fun_runge

Analyse the effect of noise on the `fun_runge` function with the following settings:

```
lower = np.array([-10])
upper = np.array([10])
fun = analytical().fun_runge
fun_control = {"sigma": 0.25,
              "seed": 123}
```

9.4.3 fun_forrester

Analyse the effect of noise on the `fun_forrester` function with the following settings:

```
lower = np.array([0])
upper = np.array([1])
fun = analytical().fun_forrester
fun_control = {"sigma": 5,
              "seed": 123}
```

9.4.4 fun_xsin

Analyse the effect of noise on the `fun_xsin` function with the following settings:

```
lower = np.array([-1.])
upper = np.array([1.])
fun = analytical().fun_xsin
fun_control = {"sigma": 0.5,
              "seed": 123}
```

```
spot_1_noisy.mean_y.shape[0]
```

18

10 HPT: sklearn SVC on Moons Data

This document refers to the following software versions:

- python: 3.10.10

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

spotPython can be installed via pip. Alternatively, the source code can be downloaded from gitHub: <https://github.com/sequential-parameter-optimization/spotPython>.

```
!pip install spotPython
```

- Uncomment the following lines if you want to for (re-)installation the latest version of spotPython from gitHub.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

10.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time and the initial design size.

```
MAX_TIME = 1
INIT_SIZE = 5
```

```

import os
import copy
import socket
from datetime import datetime
from dateutil.tz import tzlocal
start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '10-sklearn' + "_" + HOSTNAME + "_" + str(MAX_TIME) + "min_" + str(INIT_
experiment_name = experiment_name.replace(':', '-')
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')

```

10-sklearn_bartz09_1min_5init_2023-06-19_02-26-52

10.2 Step 2: Initialization of the Empty fun_control Dictionary

 Caution: Tensorboard does not work under Windows

- Since tensorboard does not work under Windows, we recommend setting the parameter `tensorboard_path` to `None` if you are working under Windows.

```

from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="classification",
    tensorboard_path="runs/10_spot_hpt_sklearn_classification")

```

10.3 Step 3: SKlearn Load Data (Classification)

Randomly generate classification data.

```

import pandas as pd
import numpy as np
from sklearn.model_selection import train_test_split
from sklearn.datasets import make_moons, make_circles, make_classification
n_features = 2
n_samples = 250
target_column = "y"

```

```

ds = make_moons(n_samples, noise=0.5, random_state=0)
X, y = ds
X_train, X_test, y_train, y_test = train_test_split(
    X, y, test_size=0.4, random_state=42
)
train = pd.DataFrame(np.hstack((X_train, y_train.reshape(-1, 1))))
test = pd.DataFrame(np.hstack((X_test, y_test.reshape(-1, 1))))
train.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
test.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
train.head()

```

	x1	x2	y
0	1.083978	-1.246111	1.0
1	0.074916	0.868104	0.0
2	-1.668535	0.751752	0.0
3	1.286597	1.454165	0.0
4	1.387021	0.448355	1.0

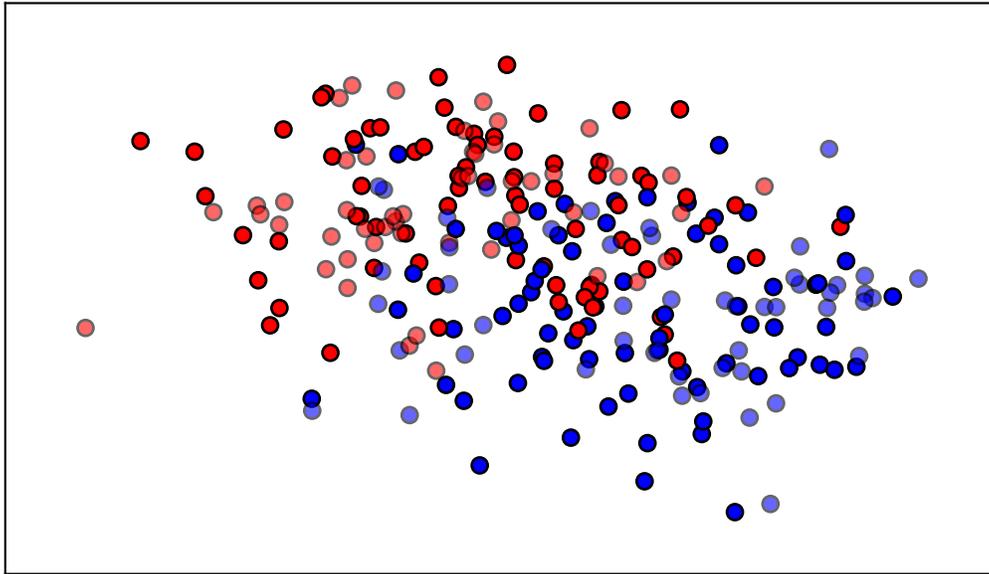
```

import matplotlib.pyplot as plt
from matplotlib.colors import ListedColormap

x_min, x_max = X[:, 0].min() - 0.5, X[:, 0].max() + 0.5
y_min, y_max = X[:, 1].min() - 0.5, X[:, 1].max() + 0.5
cm = plt.cm.RdBu
cm_bright = ListedColormap(["#FF0000", "#0000FF"])
ax = plt.subplot(1, 1, 1)
ax.set_title("Input data")
# Plot the training points
ax.scatter(X_train[:, 0], X_train[:, 1], c=y_train, cmap=cm_bright, edgecolors="k")
# Plot the testing points
ax.scatter(
    X_test[:, 0], X_test[:, 1], c=y_test, cmap=cm_bright, alpha=0.6, edgecolors="k"
)
ax.set_xlim(x_min, x_max)
ax.set_ylim(y_min, y_max)
ax.set_xticks(())
ax.set_yticks(())
plt.tight_layout()
plt.show()

```

Input data



```
n_samples = len(train)
# add the dataset to the fun_control
fun_control.update({"data": None, # dataset,
                  "train": train,
                  "test": test,
                  "n_samples": n_samples,
                  "target_column": target_column})
```

10.4 Step 4: Specification of the Preprocessing Model

Data preprocessing can be very simple, e.g., you can ignore it. Then you would choose the `prep_model` "None":

```
prep_model = None
fun_control.update({"prep_model": prep_model})
```

A default approach for numerical data is the `StandardScaler` (mean 0, variance 1). This can be selected as follows:

```

from sklearn.preprocessing import StandardScaler
prep_model = StandardScaler()
fun_control.update({"prep_model": prep_model})

```

Even more complicated pre-processing steps are possible, e.g., the following pipeline:

```

# categorical_columns = []
# one_hot_encoder = OneHotEncoder(handle_unknown="ignore", sparse_output=False)
# prep_model = ColumnTransformer(
#     transformers=[
#         ("categorical", one_hot_encoder, categorical_columns),
#     ],
#     remainder=StandardScaler(),
# )

```

10.5 Step 5: Select Model (algorithm) and core_model_hyper_dict

The selection of the algorithm (ML model) that should be tuned is done by specifying the its name from the `sklearn` implementation. For example, the SVC support vector machine classifier is selected as follows:

```

from sklearn.linear_model import RidgeCV
from sklearn.ensemble import RandomForestClassifier
from sklearn.svm import SVC
from sklearn.linear_model import LogisticRegression
from sklearn.neighbors import KNeighborsClassifier
from sklearn.ensemble import GradientBoostingClassifier
from sklearn.ensemble import GradientBoostingRegressor
from sklearn.linear_model import ElasticNet
from spotPython.hyperparameters.values import add_core_model_to_fun_control
from spotPython.data.sklearn_hyper_dict import SklearnHyperDict
from spotPython.fun.hypersklearn import HyperSklearn

# core_model = RidgeCV
# core_model = GradientBoostingRegressor
# core_model = ElasticNet
# core_model = RandomForestClassifier
core_model = SVC

```

```

# core_model = LogisticRegression
# core_model = KNeighborsClassifier
# core_model = GradientBoostingClassifier
fun_control = add_core_model_to_fun_control(core_model=core_model,
                                          fun_control=fun_control,
                                          hyper_dict=SklearnHyperDict,
                                          filename=None)

```

Now `fun_control` has the information from the JSON file. The corresponding entries for the `core_model` class are shown below.

```
fun_control['core_model_hyper_dict']
```

```

{'C': {'type': 'float',
      'default': 1.0,
      'transform': 'None',
      'lower': 0.1,
      'upper': 10.0},
 'kernel': {'levels': ['linear', 'poly', 'rbf', 'sigmoid'],
            'type': 'factor',
            'default': 'rbf',
            'transform': 'None',
            'core_model_parameter_type': 'str',
            'lower': 0,
            'upper': 3},
 'degree': {'type': 'int',
            'default': 3,
            'transform': 'None',
            'lower': 3,
            'upper': 3},
 'gamma': {'levels': ['scale', 'auto'],
           'type': 'factor',
           'default': 'scale',
           'transform': 'None',
           'core_model_parameter_type': 'str',
           'lower': 0,
           'upper': 1},
 'coef0': {'type': 'float',
           'default': 0.0,
           'transform': 'None',
           'lower': 0.0,
           'upper': 0.0},

```

```

'shrinking': {'levels': [0, 1],
'type': 'factor',
'default': 0,
'transform': 'None',
'core_model_parameter_type': 'bool',
'lower': 0,
'upper': 1},
'probability': {'levels': [0, 1],
'type': 'factor',
'default': 0,
'transform': 'None',
'core_model_parameter_type': 'bool',
'lower': 0,
'upper': 1},
'tol': {'type': 'float',
'default': 0.001,
'transform': 'None',
'lower': 0.0001,
'upper': 0.01},
'cache_size': {'type': 'float',
'default': 200,
'transform': 'None',
'lower': 100,
'upper': 400},
'break_ties': {'levels': [0, 1],
'type': 'factor',
'default': 0,
'transform': 'None',
'core_model_parameter_type': 'bool',
'lower': 0,
'upper': 1}}

```

10.6 Step 6: Modify hyper_dict Hyperparameters for the Selected Algorithm aka core_model

spotPython provides functions for modifying the hyperparameters, their bounds and factors as well as for activating and de-activating hyperparameters without re-compilation of the Python source code. These functions were described in Section [14.6](#).

10.6.1 Modify hyperparameter of type numeric and integer (boolean)

Numeric and boolean values can be modified using the `modify_hyper_parameter_bounds` method. For example, to change the `tol` hyperparameter of the SVC model to the interval `[1e-3, 1e-2]`, the following code can be used:

```
from spotPython.hyperparameters.values import modify_hyper_parameter_bounds
fun_control = modify_hyper_parameter_bounds(fun_control, "tol", bounds=[1e-3, 1e-2])
# fun_control = modify_hyper_parameter_bounds(fun_control, "min_samples_split", bounds=[3,
#fun_control = modify_hyper_parameter_bounds(fun_control, "merit_preprune", bounds=[0, 0])
fun_control["core_model_hyper_dict"]["tol"]
```

```
{'type': 'float',
 'default': 0.001,
 'transform': 'None',
 'lower': 0.001,
 'upper': 0.01}
```

10.6.2 Modify hyperparameter of type factor

Factors can be modified with the `modify_hyper_parameter_levels` function. For example, to exclude the `sigmoid` kernel from the tuning, the `kernel` hyperparameter of the SVC model can be modified as follows:

```
from spotPython.hyperparameters.values import modify_hyper_parameter_levels
fun_control = modify_hyper_parameter_levels(fun_control, "kernel", ["linear", "poly", "rbf", "sigmoid"])
fun_control["core_model_hyper_dict"]["kernel"]
```

```
{'levels': ['linear', 'poly', 'rbf'],
 'type': 'factor',
 'default': 'rbf',
 'transform': 'None',
 'core_model_parameter_type': 'str',
 'lower': 0,
 'upper': 2}
```

10.6.3 Optimizers

Optimizers are described in Section [14.6.1](#).

10.7 Step 7: Selection of the Objective (Loss) Function

There are two metrics:

1. `metric_river` is used for the river based evaluation via `eval_oml_iter_progressive`.
2. `metric_sklearn` is used for the sklearn based evaluation.

```
from sklearn.metrics import mean_absolute_error, accuracy_score, roc_curve, roc_auc_score,
fun_control.update({
    "metric_sklearn": log_loss,
})
```

10.7.1 Predict Classes or Class Probabilities

If the key `"predict_proba"` is set to `True`, the class probabilities are predicted. `False` is the default, i.e., the classes are predicted.

```
fun_control.update({
    "predict_proba": False,
})
```

10.8 Step 8: Calling the SPOT Function

10.8.1 Preparing the SPOT Call

The following code passes the information about the parameter ranges and bounds to `spot`.

```
# extract the variable types, names, and bounds
from spotPython.hyperparameters.values import (get_bound_values,
    get_var_name,
    get_var_type,)
var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
fun_control.update({"var_type": var_type,
    "var_name": var_name})
lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")
```

```
from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))
```

name	type	default	lower	upper	transform
C	float	1.0	0.1	10	None
kernel	factor	rbf	0	2	None
degree	int	3	3	3	None
gamma	factor	scale	0	1	None
coef0	float	0.0	0	0	None
shrinking	factor	0	0	1	None
probability	factor	0	0	1	None
tol	float	0.001	0.001	0.01	None
cache_size	float	200.0	100	400	None
break_ties	factor	0	0	1	None

10.8.2 The Objective Function

The objective function is selected next. It implements an interface from sklearn's training, validation, and testing methods to spotPython.

```
from spotPython.fun.hypersklearn import HyperSklearn
fun = HyperSklearn().fun_sklearn
```

10.8.3 Run the Spot Optimizer

- Run SPOT for approx. x mins (`max_time`).
- Note: the run takes longer, because the evaluation time of initial design (here: `inital_size`, 20 points) is not considered.

```
from spotPython.hyperparameters.values import get_default_hyperparameters_as_array
hyper_dict=SklearnHyperDict().load()
X_start = get_default_hyperparameters_as_array(fun_control, hyper_dict)
X_start
```

```
array([[1.e+00, 2.e+00, 3.e+00, 0.e+00, 0.e+00, 0.e+00, 0.e+00, 1.e-03,
        2.e+02, 0.e+00]])
```

10.8.4 Starting the Hyperparameter Tuning

```
import numpy as np
from spotPython.spot import spot
from math import inf
spot_tuner = spot.Spot(fun=fun,
                      lower = lower,
                      upper = upper,
                      fun_evals = inf,
                      fun_repeats = 1,
                      max_time = MAX_TIME,
                      noise = False,
                      tolerance_x = np.sqrt(np.spacing(1)),
                      var_type = var_type,
                      var_name = var_name,
                      infill_criterion = "y",
                      n_points = 1,
                      seed=123,
                      log_level = 50,
                      show_models= False,
                      show_progress= True,
                      fun_control = fun_control,
                      design_control={"init_size": INIT_SIZE,
                                     "repeats": 1},
                      surrogate_control={"noise": True,
                                       "cod_type": "norm",
                                       "min_theta": -4,
                                       "max_theta": 3,
                                       "n_theta": len(var_name),
                                       "model_fun_evals": 10_000,
                                       "log_level": 50
                                       })

spot_tuner.run(X_start=X_start)
```

spotPython tuning: 5.691103166702708 [-----] 2.78%

spotPython tuning: 5.691103166702708 [-----] 4.59%

spotPython tuning: 5.691103166702708 [#-----] 6.10%

spotPython tuning: 5.691103166702708 [#-----] 7.54%

```
spotPython tuning: 5.691103166702708 [#-----] 8.97%
spotPython tuning: 5.691103166702708 [#-----] 11.38%
spotPython tuning: 5.691103166702708 [#-----] 13.68%
spotPython tuning: 5.691103166702708 [##-----] 15.88%
spotPython tuning: 5.691103166702708 [##-----] 18.18%
spotPython tuning: 5.691103166702708 [##-----] 20.41%
spotPython tuning: 5.691103166702708 [##-----] 22.86%
spotPython tuning: 5.691103166702708 [###-----] 25.44%
spotPython tuning: 5.691103166702708 [###-----] 33.91%
spotPython tuning: 5.691103166702708 [#####-----] 45.79%
spotPython tuning: 5.691103166702708 [#####-----] 58.45%
spotPython tuning: 5.691103166702708 [#####---] 71.14%
spotPython tuning: 5.691103166702708 [#####--] 83.66%
spotPython tuning: 5.691103166702708 [#####] 95.99%
spotPython tuning: 5.691103166702708 [#####] 100.00% Done...
```

```
<spotPython.spot.spot.Spot at 0x28d897040>
```

10.9 Step 9: Results

```

SAVE = False
LOAD = False

if SAVE:
    result_file_name = "res_" + experiment_name + ".pkl"
    with open(result_file_name, 'wb') as f:
        pickle.dump(spot_tuner, f)

if LOAD:
    result_file_name = "res_ch10-friedman-hpt-0_maans03_60min_20init_1K_2023-04-14_10-11-1
    with open(result_file_name, 'rb') as f:
        spot_tuner = pickle.load(f)

```

After the hyperparameter tuning run is finished, the progress of the hyperparameter tuning can be visualized. The following code generates the progress plot from `spot_tuner.plot_progress`.

```

spot_tuner.plot_progress(log_y=False,
    filename="./figures/" + experiment_name+"_progress.png")

```

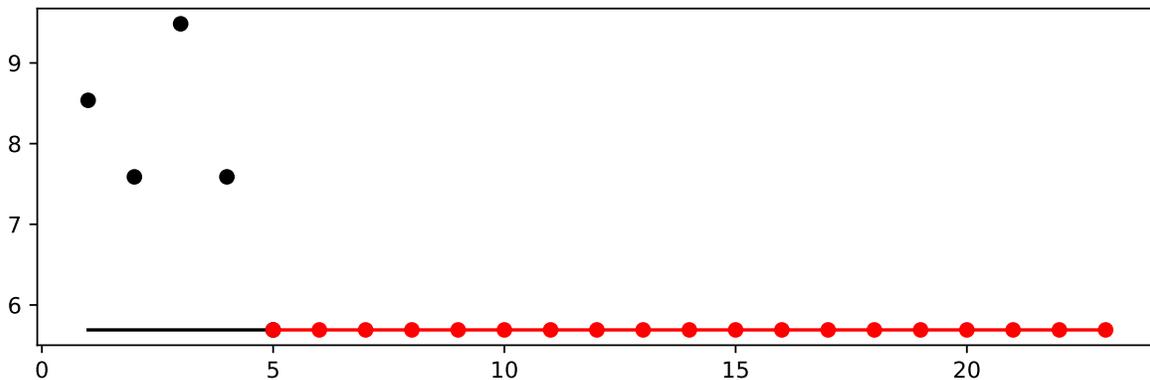


Figure 10.1: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

- Print the results

```

print(gen_design_table(fun_control=fun_control,
    spot=spot_tuner))

```

name	type	default	lower	upper	tuned	transform
-----	-----	-----	-----	-----	-----	-----

C	float	1.0		0.1		10.0		3.6280771109650245		None
kernel	factor	rbf		0.0		2.0				1.0
degree	int	3		3.0		3.0				3.0
gamma	factor	scale		0.0		1.0				0.0
coef0	float	0.0		0.0		0.0				0.0
shrinking	factor	0		0.0		1.0				1.0
probability	factor	0		0.0		1.0				0.0
tol	float	0.001		0.001		0.01		0.006642600916881275		None
cache_size	float	200.0		100.0		400.0		202.03372626175258		None
break_ties	factor	0		0.0		1.0				1.0

10.9.1 Show variable importance

```
spot_tuner.plot_importance(threshold=0.025, filename="./figures/" + experiment_name+"_impo
```

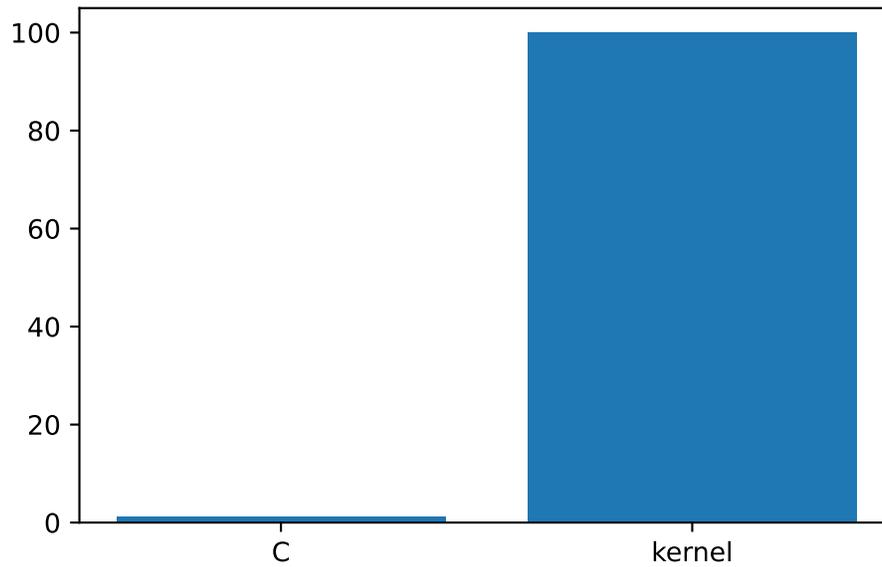


Figure 10.2: Variable importance plot, threshold 0.025.

10.9.2 Get Default Hyperparameters

```
from spotPython.hyperparameters.values import get_default_values, transform_hyper_parameter_values
values_default = get_default_values(fun_control)
values_default = transform_hyper_parameter_values(fun_control=fun_control, hyper_parameter_values=values_default)
```

```
{'C': 1.0,
 'kernel': 'rbf',
 'degree': 3,
 'gamma': 'scale',
 'coef0': 0.0,
 'shrinking': 0,
 'probability': 0,
 'tol': 0.001,
 'cache_size': 200.0,
 'break_ties': 0}
```

```
from sklearn.pipeline import make_pipeline
model_default = make_pipeline(fun_control["prep_model"], fun_control["core_model"](**values_default))
model_default
```

```
Pipeline(steps=[('standardscaler', StandardScaler()),
                 ('svc',
                  SVC(break_ties=0, cache_size=200.0, probability=0,
                      shrinking=0))])
```

10.9.3 Get SPOT Results

```
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))
print(X)
```

```
[[3.62807711e+00 1.00000000e+00 3.00000000e+00 0.00000000e+00
 0.00000000e+00 1.00000000e+00 0.00000000e+00 6.64260092e-03
 2.02033726e+02 1.00000000e+00]]
```

```
from spotPython.hyperparameters.values import assign_values, return_conf_list_from_var_dict
v_dict = assign_values(X, fun_control["var_name"])
return_conf_list_from_var_dict(var_dict=v_dict, fun_control=fun_control)
```

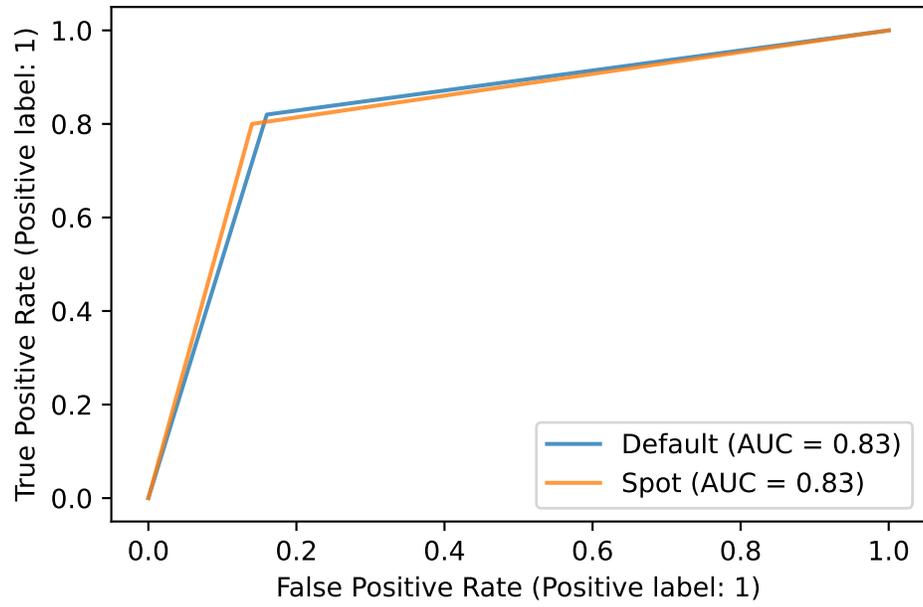
```
[{'C': 3.6280771109650245,
  'kernel': 'poly',
  'degree': 3,
  'gamma': 'scale',
  'coef0': 0.0,
  'shrinking': 1,
  'probability': 0,
  'tol': 0.006642600916881275,
  'cache_size': 202.03372626175258,
  'break_ties': 1}]
```

```
from spotPython.hyperparameters.values import get_one_sklearn_model_from_X
model_spot = get_one_sklearn_model_from_X(X, fun_control)
model_spot
```

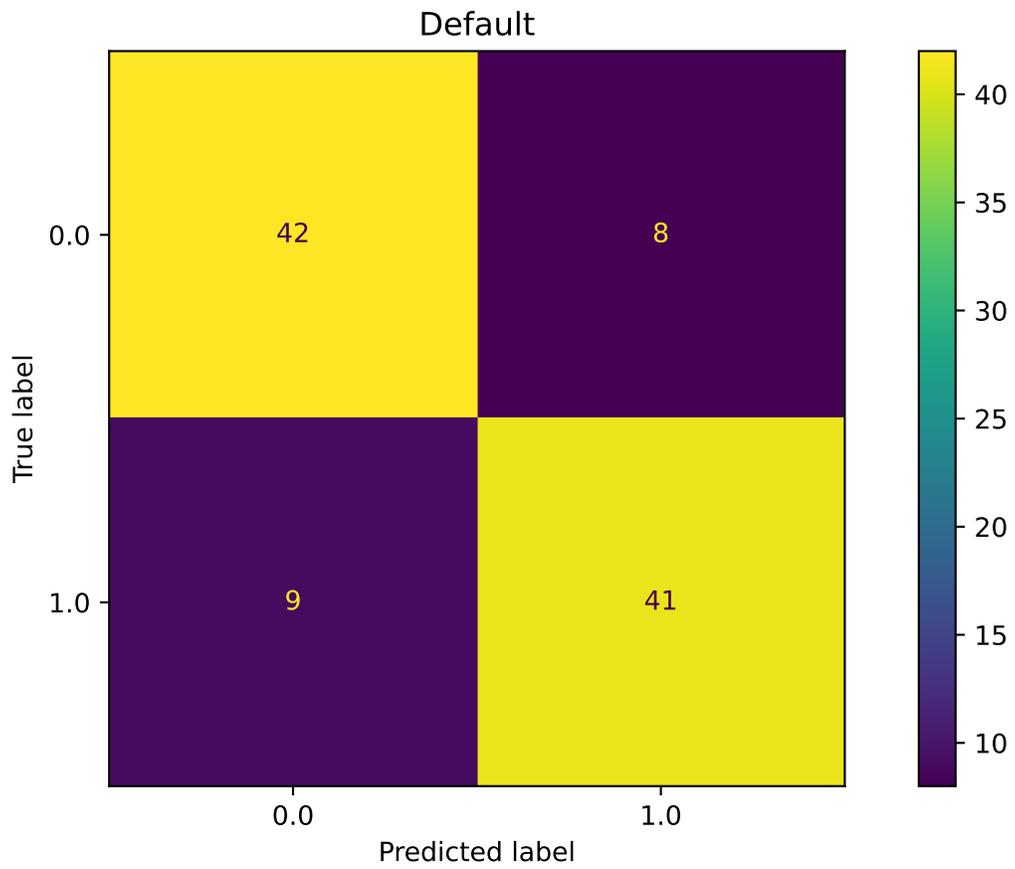
```
Pipeline(steps=[('standardscaler', StandardScaler()),
                 ('svc',
                  SVC(C=3.6280771109650245, break_ties=1,
                      cache_size=202.03372626175258, kernel='poly',
                      probability=0, shrinking=1, tol=0.006642600916881275))])
```

10.9.4 Plot: Compare Predictions

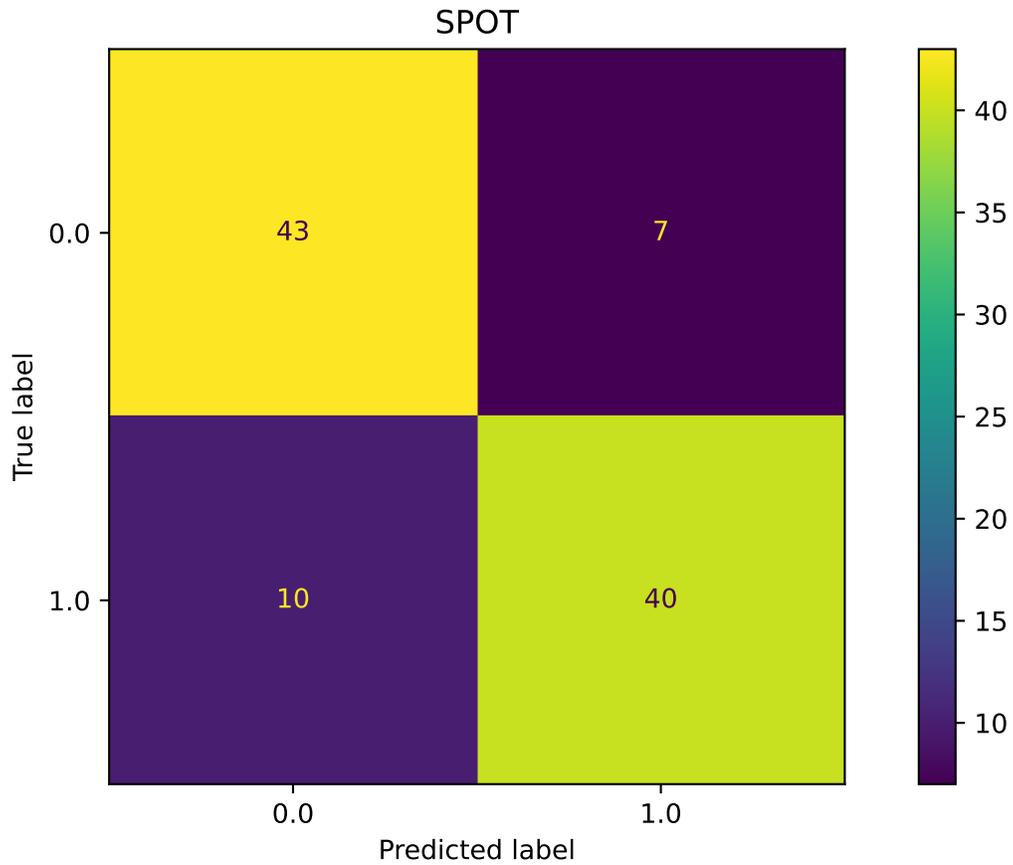
```
from spotPython.plot.validation import plot_roc
plot_roc([model_default, model_spot], fun_control, model_names=["Default", "Spot"])
```



```
from spotPython.plot.validation import plot_confusion_matrix
plot_confusion_matrix(model_default, fun_control, title = "Default")
```



```
plot_confusion_matrix(model_spot, fun_control, title="SPOT")
```



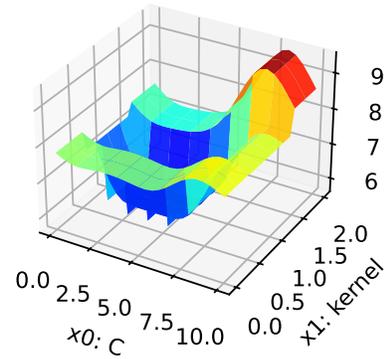
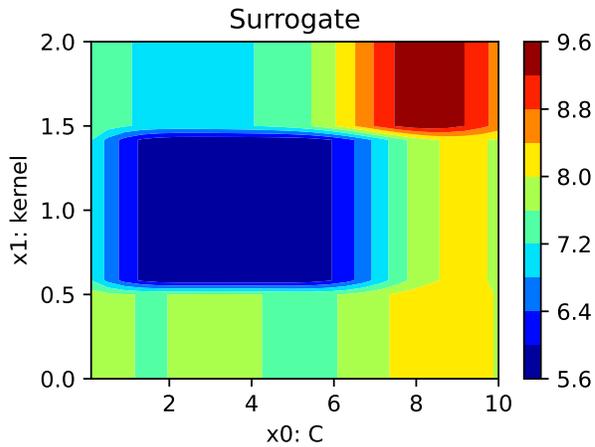
```
min(spot_tuner.y), max(spot_tuner.y)
```

```
(5.691103166702708, 9.485171944504513)
```

10.9.5 Detailed Hyperparameter Plots

```
filename = "./figures/" + experiment_name
spot_tuner.plot_important_hyperparameter_contour(filename=filename)
```

```
C: 1.1399176173997725
kernel: 100.0
```



10.9.6 Parallel Coordinates Plot

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Unable to display output for mime type(s): text/html

10.9.7 Plot all Combinations of Hyperparameters

- Warning: this may take a while.

```
PLOT_ALL = False
if PLOT_ALL:
    n = spot_tuner.k
    for i in range(n-1):
        for j in range(i+1, n):
            spot_tuner.plot_contour(i=i, j=j, min_z=min_z, max_z = max_z)
```

11 HPT: PyTorch With fashionMNIST

In this tutorial, we will show how `spotPython` can be integrated into the PyTorch training workflow.

This document refers to the following software versions:

- python: 3.10.10
- torch: 2.0.1
- torchvision: 0.15.0

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

`spotPython` can be installed via `pip`. Alternatively, the source code can be downloaded from gitHub: <https://github.com/sequential-parameter-optimization/spotPython>.

```
!pip install spotPython
```

- Uncomment the following lines if you want to for (re-)installation the latest version of `spotPython` from gitHub.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

11.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time, initial design size and the device that is used.

 Caution: Run time and initial design size should be increased for real experiments

- MAX_TIME is set to one minute for demonstration purposes. For real experiments, this should be increased to at least 1 hour.
- INIT_SIZE is set to 5 for demonstration purposes. For real experiments, this should be increased to at least 10.

 Note: Device selection

- The device can be selected by setting the variable DEVICE.
- Since we are using a simple neural net, the setting "cpu" is preferred (on Mac).
- If you have a GPU, you can use "cuda:0" instead.
- If DEVICE is set to None, spotPython will automatically select the device.
 - This might result in "mps" on Macs, which is not the best choice for simple neural nets.

```
MAX_TIME = 1
INIT_SIZE = 5
DEVICE = "cpu" # "cuda:0"
```

```
from spotPython.utils.device import getDevice
DEVICE = getDevice(DEVICE)
print(DEVICE)
```

cpu

```
import os
import copy
import socket
from datetime import datetime
from dateutil.tz import tzlocal
start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '11-torch' + "_" + HOSTNAME + "_" + str(MAX_TIME) + "min_" + str(INIT_SIZE)
experiment_name = experiment_name.replace(':', '-')
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')
```

11-torch_bartz09_1min_5init_2023-06-19_02-34-00

11.2 Step 2: Initialization of the Empty fun_control Dictionary

spotPython uses a Python dictionary for storing the information required for the hyperparameter tuning process, which was described in Section 14.2.

 Caution: Tensorboard does not work under Windows

- Since tensorboard does not work under Windows, we recommend setting the parameter `tensorboard_path` to `None` if you are working under Windows.

```
from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="classification",
    tensorboard_path="runs/11_spot_hpt_torch_fashion_mnist",
    device=DEVICE)
```

11.3 Step 3: PyTorch Data Loading

11.3.1 Load fashionMNIST Data

```
from torchvision import datasets, transforms
from torchvision.transforms import ToTensor
def load_data(data_dir="./data"):
    # Download training data from open datasets.
    training_data = datasets.FashionMNIST(
        root=data_dir,
        train=True,
        download=True,
        transform=ToTensor(),
    )
    # Download test data from open datasets.
    test_data = datasets.FashionMNIST(
        root=data_dir,
        train=False,
        download=True,
        transform=ToTensor(),
    )
    return training_data, test_data
```

```

train, test = load_data()
train.data.shape, test.data.shape

```

```
(torch.Size([60000, 28, 28]), torch.Size([10000, 28, 28]))
```

```

n_samples = len(train)
# add the dataset to the fun_control
fun_control.update({"data": None,
                   "train": train,
                   "test": test,
                   "n_samples": n_samples,
                   "target_column": None})

```

11.4 Step 4: Specification of the Preprocessing Model

After the training and test data are specified and added to the `fun_control` dictionary, `spotPython` allows the specification of a data preprocessing pipeline, e.g., for the scaling of the data or for the one-hot encoding of categorical variables, see Section 14.4. This feature is not used here, so we do not change the default value (which is `None`).

11.5 Step 5: Select Model (algorithm) and `core_model_hyper_dict`

`spotPython` implements a class which is similar to the class described in the PyTorch tutorial. The class is called `Net_fashionMNIST` and is implemented in the file `netfashionMNIST.py`. The class is imported here.

```

from torch import nn
import spotPython.torch.netcore as netcore

class Net_fashionMNIST(netcore.Net_Core):
    def __init__(self, l1, l2, lr_mult, batch_size, epochs, k_folds, patience, optimizer,
                 super(Net_fashionMNIST, self).__init__(
                     lr_mult=lr_mult,
                     batch_size=batch_size,
                     epochs=epochs,

```

```

        k_folds=k_folds,
        patience=patience,
        optimizer=optimizer,
        sgd_momentum=sgd_momentum,
    )
    self.flatten = nn.Flatten()
    self.linear_relu_stack = nn.Sequential(
        nn.Linear(28 * 28, 11),
        nn.ReLU(),
        nn.Linear(11, 12),
        nn.ReLU(),
        nn.Linear(12, 10)
    )

    def forward(self, x):
        x = self.flatten(x)
        logits = self.linear_relu_stack(x)
        return logits

```

This class inherits from the class `Net_Core` which is implemented in the file `netcore.py`, see Section 14.5.1.

```

from spotPython.data.torch_hyper_dict import TorchHyperDict
from spotPython.torch.netfashionMNIST import Net_fashionMNIST
from spotPython.hyperparameters.values import add_core_model_to_fun_control
fun_control = add_core_model_to_fun_control(core_model=Net_fashionMNIST,
                                           fun_control=fun_control,
                                           hyper_dict=TorchHyperDict,
                                           filename=None)

```

11.5.1 The Search Space

11.5.2 Configuring the Search Space With `spotPython`

11.5.2.1 The `hyper_dict` Hyperparameters for the Selected Algorithm

`spotPython` uses JSON files for the specification of the hyperparameters, which were described in Section 14.5.5.

The corresponding entries for the `core_model` class are shown below.

```
fun_control['core_model_hyper_dict']
```

```
{'l1': {'type': 'int',  
       'default': 5,  
       'transform': 'transform_power_2_int',  
       'lower': 2,  
       'upper': 9},  
'l2': {'type': 'int',  
       'default': 5,  
       'transform': 'transform_power_2_int',  
       'lower': 2,  
       'upper': 9},  
'lr_mult': {'type': 'float',  
            'default': 1.0,  
            'transform': 'None',  
            'lower': 0.1,  
            'upper': 10.0},  
'batch_size': {'type': 'int',  
               'default': 4,  
               'transform': 'transform_power_2_int',  
               'lower': 1,  
               'upper': 4},  
'epochs': {'type': 'int',  
            'default': 3,  
            'transform': 'transform_power_2_int',  
            'lower': 3,  
            'upper': 4},  
'k_folds': {'type': 'int',  
            'default': 1,  
            'transform': 'None',  
            'lower': 1,  
            'upper': 1},  
'patience': {'type': 'int',  
              'default': 5,  
              'transform': 'None',  
              'lower': 2,  
              'upper': 10},  
'optimizer': {'levels': ['Adadelata',  
                          'Adagrad',  
                          'Adam',  
                          'AdamW',  
                          'SparseAdam',
```

```

'Adamax',
'ASGD',
'NAdam',
'RAdam',
'RMSprop',
'Rprop',
'SGD'],
'type': 'factor',
'default': 'SGD',
'transform': 'None',
'core_model_parameter_type': 'str',
'lower': 0,
'upper': 12},
'sgd_momentum': {'type': 'float',
'default': 0.0,
'transform': 'None',
'lower': 0.0,
'upper': 1.0}}

```

11.6 Step 6: Modify hyper_dict Hyperparameters for the Selected Algorithm aka core_model

spotPython provides functions for modifying the hyperparameters, their bounds and factors as well as for activating and de-activating hyperparameters without re-compilation of the Python source code. These functions were described in Section 14.6.

11.6.1 Modify hyperparameter of type numeric and integer (boolean)

The hyperparameter `k_folds` is not used, it is de-activated here by setting the lower and upper bound to the same value.

 Caution: Small net size, number of epochs, and patience for demonstration purposes

- Net sizes 11 and 12 as well as `epochs` and `patience` are set to small values for demonstration purposes. These values are too small for a real application.
- More reasonable values are, e.g.:
 - `fun_control = modify_hyper_parameter_bounds(fun_control, "11", bounds=[2, 7])`
 - `fun_control = modify_hyper_parameter_bounds(fun_control,`

```
"epochs", bounds=[7, 9]) and
- fun_control = modify_hyper_parameter_bounds(fun_control,
"patience", bounds=[2, 7])
```

```
from spotPython.hyperparameters.values import modify_hyper_parameter_bounds
fun_control = modify_hyper_parameter_bounds(fun_control, "k_folds", bounds=[0, 0])
fun_control = modify_hyper_parameter_bounds(fun_control, "patience", bounds=[2, 2])
fun_control = modify_hyper_parameter_bounds(fun_control, "epochs", bounds=[2, 3])
fun_control = modify_hyper_parameter_bounds(fun_control, "l1", bounds=[2, 5])
fun_control = modify_hyper_parameter_bounds(fun_control, "l2", bounds=[2, 5])
```

11.6.2 Modify hyperparameter of type factor

```
from spotPython.hyperparameters.values import modify_hyper_parameter_levels
fun_control = modify_hyper_parameter_levels(fun_control, "optimizer", ["Adam", "AdamW", "Ad
```

11.6.3 Optimizers

Optimizers are described in Section [14.6.1](#).

```
fun_control = modify_hyper_parameter_bounds(fun_control,
"lr_mult", bounds=[1e-3, 1e-3])
fun_control = modify_hyper_parameter_bounds(fun_control,
"sgd_momentum", bounds=[0.9, 0.9])
```

11.7 Step 7: Selection of the Objective (Loss) Function

11.7.1 Evaluation

The evaluation procedure requires the specification of two elements:

1. the way how the data is split into a train and a test set and
2. the loss function (and a metric).

These are described in Section [19.7.1](#).

The key "loss_function" specifies the loss function which is used during the optimization, see Section [14.7.5](#).

We will use CrossEntropy loss for the multiclass-classification task.

```
from torch.nn import CrossEntropyLoss
loss_function = CrossEntropyLoss()
fun_control.update({
    "loss_function": loss_function,
    "shuffle": True,
    "eval": "train_hold_out"
})
```

11.7.2 Metric

```
from torchmetrics import Accuracy
metric_torch = Accuracy(task="multiclass", num_classes=10).to(fun_control["device"])
fun_control.update({"metric_torch": metric_torch})
```

11.8 Step 8: Calling the SPOT Function

11.8.1 Preparing the SPOT Call

The following code passes the information about the parameter ranges and bounds to `spot`.

```
# extract the variable types, names, and bounds
from spotPython.hyperparameters.values import (get_bound_values,
    get_var_name,
    get_var_type,)
var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
fun_control.update({"var_type": var_type,
    "var_name": var_name})
lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")

from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))
```

name	type	default	lower	upper	transform
-----	-----	-----	-----	-----	-----

l1	int	5	2	5	transform_power_2_int	
l2	int	5	2	5	transform_power_2_int	
lr_mult	float	1.0	0.001	0.001	None	
batch_size	int	4	1	4	transform_power_2_int	
epochs	int	3	2	3	transform_power_2_int	
k_folds	int	1	0	0	None	
patience	int	5	2	2	None	
optimizer	factor	SGD	0	3	None	
sgd_momentum	float	0.0	0.9	0.9	None	

11.8.2 The Objective Function `fun_torch`

The objective function `fun_torch` is selected next. It implements an interface from PyTorch's training, validation, and testing methods to `spotPython`.

```
from spotPython.fun.hypertorch import HyperTorch
fun = HyperTorch().fun_torch
```

11.8.3 Starting the Hyperparameter Tuning

```
import numpy as np
from spotPython.spot import spot
from math import inf
spot_tuner = spot.Spot(fun=fun,
                      lower = lower,
                      upper = upper,
                      fun_evals = inf,
                      fun_repeats = 1,
                      max_time = MAX_TIME,
                      noise = False,
                      tolerance_x = np.sqrt(np.spacing(1)),
                      var_type = var_type,
                      var_name = var_name,
                      infill_criterion = "y",
                      n_points = 1,
                      seed=123,
                      log_level = 50,
                      show_models= False,
                      show_progress= True,
                      fun_control = fun_control,
```


MulticlassAccuracy: 0.2035000026226044 | Loss: 2.0989609602292378 | Acc: 0.2035000000000000.
Returned to Spot: Validation loss: 2.098960960229238

config: {'l1': 8, 'l2': 8, 'lr_mult': 0.001, 'batch_size': 8, 'epochs': 4, 'k_folds': 0, 'pa'
Epoch: 1 |

MulticlassAccuracy: 0.0995416641235352 | Loss: 2.3227417463461557 | Acc: 0.0995416666666667.
Epoch: 2 |

MulticlassAccuracy: 0.0995416641235352 | Loss: 2.3156459696292879 | Acc: 0.0995416666666667.
Epoch: 3 |

MulticlassAccuracy: 0.0995416641235352 | Loss: 2.3093050432999931 | Acc: 0.0995416666666667.
Epoch: 4 |

MulticlassAccuracy: 0.0999583303928375 | Loss: 2.3032854072252911 | Acc: 0.0999583333333333.
Returned to Spot: Validation loss: 2.303285407225291

config: {'l1': 32, 'l2': 16, 'lr_mult': 0.001, 'batch_size': 2, 'epochs': 8, 'k_folds': 0, 'p'
Epoch: 1 |

MulticlassAccuracy: 0.2719583213329315 | Loss: 2.1422173360188803 | Acc: 0.2719583333333334.
Epoch: 2 |

MulticlassAccuracy: 0.3155833184719086 | Loss: 1.9287186470131079 | Acc: 0.3155833333333333.
Epoch: 3 |

MulticlassAccuracy: 0.3675000071525574 | Loss: 1.6976838712344566 | Acc: 0.3675000000000000.
Epoch: 4 |

MulticlassAccuracy: 0.4343749880790710 | Loss: 1.4833368577162425 | Acc: 0.4343750000000000.
Epoch: 5 |

MulticlassAccuracy: 0.5630833506584167 | Loss: 1.3038433362829189 | Acc: 0.5630833333333334.
Epoch: 6 |

MulticlassAccuracy: 0.5956666469573975 | Loss: 1.1723824600316584 | Acc: 0.5956666666666667.
Epoch: 7 |

MulticlassAccuracy: 0.6169166564941406 | Loss: 1.0750376099298398 | Acc: 0.6169166666666667.
Epoch: 8 |

MulticlassAccuracy: 0.6428750157356262 | Loss: 1.0028487622166673 | Acc: 0.6428750000000000.
Returned to Spot: Validation loss: 1.0028487622166673

config: {'l1': 4, 'l2': 8, 'lr_mult': 0.001, 'batch_size': 4, 'epochs': 4, 'k_folds': 0, 'pa
Epoch: 1 |

MulticlassAccuracy: 0.1016250029206276 | Loss: 2.2951913241545361 | Acc: 0.1016250000000000.
Epoch: 2 |

MulticlassAccuracy: 0.0960833355784416 | Loss: 2.2737737801869708 | Acc: 0.09608333333333333.
Epoch: 3 |

MulticlassAccuracy: 0.1021666675806046 | Loss: 2.2533710262775419 | Acc: 0.10216666666666667.
Epoch: 4 |

MulticlassAccuracy: 0.1210833340883255 | Loss: 2.2334476508696874 | Acc: 0.12108333333333333.
Returned to Spot: Validation loss: 2.2334476508696874

config: {'l1': 16, 'l2': 32, 'lr_mult': 0.001, 'batch_size': 8, 'epochs': 8, 'k_folds': 0, 'p
Epoch: 1 |

MulticlassAccuracy: 0.1032499969005585 | Loss: 2.2762544226646422 | Acc: 0.1032500000000000.
Epoch: 2 |

MulticlassAccuracy: 0.2049583345651627 | Loss: 2.2346506589253745 | Acc: 0.20495833333333333.
Epoch: 3 |

MulticlassAccuracy: 0.1946666687726974 | Loss: 2.1924666829903923 | Acc: 0.19466666666666667.
Epoch: 4 |

MulticlassAccuracy: 0.1942083388566971 | Loss: 2.1497441027164461 | Acc: 0.19420833333333333.
Epoch: 5 |

MulticlassAccuracy: 0.1945416629314423 | Loss: 2.1075240010420480 | Acc: 0.19454166666666667.
Epoch: 6 |

MulticlassAccuracy: 0.1961666643619537 | Loss: 2.0650026981433234 | Acc: 0.1961666666666667.
Epoch: 7 |

MulticlassAccuracy: 0.2014999985694885 | Loss: 2.0221298006772996 | Acc: 0.2015000000000000.
Epoch: 8 |

MulticlassAccuracy: 0.2169999927282333 | Loss: 1.9782093891302746 | Acc: 0.2170000000000000.
Returned to Spot: Validation loss: 1.9782093891302746

config: {'l1': 8, 'l2': 16, 'lr_mult': 0.001, 'batch_size': 8, 'epochs': 8, 'k_folds': 0, 'p'
Epoch: 1 |

MulticlassAccuracy: 0.0982916653156281 | Loss: 2.2982495419979094 | Acc: 0.0982916666666667.
Epoch: 2 |

MulticlassAccuracy: 0.0982916653156281 | Loss: 2.2705754169623056 | Acc: 0.0982916666666667.
Epoch: 3 |

MulticlassAccuracy: 0.0982916653156281 | Loss: 2.2382629959980647 | Acc: 0.0982916666666667.
Epoch: 4 |

MulticlassAccuracy: 0.1042499989271164 | Loss: 2.2041005131006242 | Acc: 0.1042500000000000.
Epoch: 5 |

MulticlassAccuracy: 0.1448333263397217 | Loss: 2.1684542604287467 | Acc: 0.1448333333333333.
Epoch: 6 |

MulticlassAccuracy: 0.1692083328962326 | Loss: 2.1305531595547995 | Acc: 0.1692083333333333.
Epoch: 7 |

MulticlassAccuracy: 0.1992083340883255 | Loss: 2.0902105062007905 | Acc: 0.1992083333333333.
Epoch: 8 |

MulticlassAccuracy: 0.2788749933242798 | Loss: 2.0482383878628414 | Acc: 0.2788750000000000.
Returned to Spot: Validation loss: 2.0482383878628414
spotPython tuning: 1.0028487622166673 [####-----] 44.27%

```

config: {'l1': 16, 'l2': 16, 'lr_mult': 0.001, 'batch_size': 2, 'epochs': 4, 'k_folds': 0, 'l': 16}
Epoch: 1 |

MulticlassAccuracy: 0.1944999992847443 | Loss: 2.2520305745303633 | Acc: 0.1945000000000000.
Epoch: 2 |

MulticlassAccuracy: 0.2562916576862335 | Loss: 2.1541367422242961 | Acc: 0.2562916666666666.
Epoch: 3 |

MulticlassAccuracy: 0.3050416707992554 | Loss: 2.0353351381421088 | Acc: 0.3050416666666667.
Epoch: 4 |

MulticlassAccuracy: 0.3342083394527435 | Loss: 1.9166675697068374 | Acc: 0.3342083333333333.
Returned to Spot: Validation loss: 1.9166675697068374
spotPython tuning: 1.0028487622166673 [#####] 100.00% Done...

<spotPython.spot.spot.Spot at 0x17fc53af0>

```

11.9 Step 9: Tensorboard

The textual output shown in the console (or code cell) can be visualized with Tensorboard as described in Section 14.9, see also the description in the documentation: [Tensorboard](#).

11.10 Step 10: Results

After the hyperparameter tuning run is finished, the results can be analyzed as described in Section 14.10.

```

SAVE = False
LOAD = False

if SAVE:
    result_file_name = "res_" + experiment_name + ".pkl"
    with open(result_file_name, 'wb') as f:
        pickle.dump(spot_tuner, f)

if LOAD:

```

```

result_file_name = "ADD THE NAME here, e.g.: res_ch10-friedman-hpt-0_maans03_60min_20i
with open(result_file_name, 'rb') as f:
    spot_tuner = pickle.load(f)

```

After the hyperparameter tuning run is finished, the progress of the hyperparameter tuning can be visualized. The following code generates the progress plot from `spot_tuner.plot_progress`.

```

spot_tuner.plot_progress(log_y=False,
    filename="./figures/" + experiment_name+"_progress.png")

```




Figure 11.1: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

- Print the results

```

print(gen_design_table(fun_control=fun_control,
    spot=spot_tuner))

```

name	type	default	lower	upper	tuned	transform
l1	int	5	2.0	5.0	5.0	transform_power_2_int
l2	int	5	2.0	5.0	4.0	transform_power_2_int
lr_mult	float	1.0	0.001	0.001	0.001	None
batch_size	int	4	1.0	4.0	1.0	transform_power_2_int
epochs	int	3	2.0	3.0	3.0	transform_power_2_int
k_folds	int	1	0.0	0.0	0.0	None
patience	int	5	2.0	2.0	2.0	None
optimizer	factor	SGD	0.0	3.0	3.0	None
sgd_momentum	float	0.0	0.9	0.9	0.9	None

11.10.1 Show variable importance

```
spot_tuner.plot_importance(threshold=0.025, filename="./figures/" + experiment_name+"_impo
```




Figure 11.2: Variable importance plot, threshold 0.025.

11.10.2 Get the Tuned Architecture (SPOT Results)

The architecture of the spotPython model can be obtained by the following code:

```
from spotPython.hyperparameters.values import get_one_core_model_from_X
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))
model_spot = get_one_core_model_from_X(X, fun_control)
model_spot
```

```
Net_fashionMNIST(
  (flatten): Flatten(start_dim=1, end_dim=-1)
  (linear_relu_stack): Sequential(
    (0): Linear(in_features=784, out_features=32, bias=True)
    (1): ReLU()
    (2): Linear(in_features=32, out_features=16, bias=True)
    (3): ReLU()
    (4): Linear(in_features=16, out_features=10, bias=True)
```

```
)  
)
```

11.10.3 Get Default Hyperparameters

```
fc = fun_control  
fc.update({"core_model_hyper_dict":  
          hyper_dict[fun_control["core_model"].__name__]})  
model_default = get_one_core_model_from_X(X_start, fun_control=fc)  
model_default
```

```
Net_fashionMNIST(  
  (flatten): Flatten(start_dim=1, end_dim=-1)  
  (linear_relu_stack): Sequential(  
    (0): Linear(in_features=784, out_features=32, bias=True)  
    (1): ReLU()  
    (2): Linear(in_features=32, out_features=32, bias=True)  
    (3): ReLU()  
    (4): Linear(in_features=32, out_features=10, bias=True)  
  )  
)
```

11.10.4 Evaluation of the Default and the Tuned Architectures

The method `train_tuned` takes a model architecture without trained weights and trains this model with the train data. The train data is split into train and validation data. The validation data is used for early stopping. The trained model weights are saved as a dictionary.

```
from spotPython.torch.traintest import train_tuned  
train_tuned(net=model_default, train_dataset=train, shuffle=True,  
            loss_function=fun_control["loss_function"],  
            metric=fun_control["metric_torch"],  
            device = fun_control["device"],  
            show_batch_interval=1_000_000,  
            path=None,  
            task=fun_control["task"])
```

Epoch: 1 |

MulticlassAccuracy: 0.3391666710376740 | Loss: 2.0737699147065483 | Acc: 0.3391666666666667.
Epoch: 2 |

MulticlassAccuracy: 0.4791249930858612 | Loss: 1.5249479449590047 | Acc: 0.4791250000000000.
Epoch: 3 |

MulticlassAccuracy: 0.6147916913032532 | Loss: 1.1953332616488139 | Acc: 0.6147916666666666.
Epoch: 4 |

MulticlassAccuracy: 0.6680416464805603 | Loss: 1.0364151503245036 | Acc: 0.6680416666666666.
Epoch: 5 |

MulticlassAccuracy: 0.6905000209808350 | Loss: 0.9377396509448688 | Acc: 0.6905000000000000.
Epoch: 6 |

MulticlassAccuracy: 0.7003750205039978 | Loss: 0.8730384647846222 | Acc: 0.7003750000000000.
Epoch: 7 |

MulticlassAccuracy: 0.7117083072662354 | Loss: 0.8279227907458941 | Acc: 0.7117083333333334.
Epoch: 8 |

MulticlassAccuracy: 0.7172083258628845 | Loss: 0.7946865118145943 | Acc: 0.7172083333333333.
Returned to Spot: Validation loss: 0.7946865118145943

```
from spotPython.torch.traintest import test_tuned
test_tuned(net=model_default, test_dataset=test,
            loss_function=fun_control["loss_function"],
            metric=fun_control["metric_torch"],
            shuffle=False,
            device = fun_control["device"],
            task=fun_control["task"])
```

MulticlassAccuracy: 0.7125999927520752 | Loss: 0.8081323936939240 | Acc: 0.7126000000000000.

Final evaluation: Validation loss: 0.808132393693924

Final evaluation: Validation metric: 0.7125999927520752

(0.808132393693924, nan, tensor(0.7126))

The following code trains the model `model_spot`. If `path` is set to a filename, e.g., `path = "model_spot_trained.pt"`, the weights of the trained model will be saved to this file.

```
train_tuned(net=model_spot, train_dataset=train,
            loss_function=fun_control["loss_function"],
            metric=fun_control["metric_torch"],
            shuffle=True,
            device = fun_control["device"],
            path=None,
            task=fun_control["task"])
```

Epoch: 1 |

Batch: 10000. Batch Size: 2. Training Loss (running): 2.232

MulticlassAccuracy: 0.3914583325386047 | Loss: 1.9770964104533195 | Acc: 0.3914583333333334.
Epoch: 2 |

Batch: 10000. Batch Size: 2. Training Loss (running): 1.882

MulticlassAccuracy: 0.4615833461284637 | Loss: 1.6711871655707558 | Acc: 0.4615833333333333.
Epoch: 3 |

Batch: 10000. Batch Size: 2. Training Loss (running): 1.598

MulticlassAccuracy: 0.5159583091735840 | Loss: 1.4528864771264294 | Acc: 0.5159583333333333.
Epoch: 4 |

Batch: 10000. Batch Size: 2. Training Loss (running): 1.398

MulticlassAccuracy: 0.5601249933242798 | Loss: 1.2964802738521248 | Acc: 0.5601250000000000.
Epoch: 5 |

Batch: 10000. Batch Size: 2. Training Loss (running): 1.260

MulticlassAccuracy: 0.5947083234786987 | Loss: 1.1789221649660417 | Acc: 0.5947083333333333.
Epoch: 6 |

Batch: 10000. Batch Size: 2. Training Loss (running): 1.143

MulticlassAccuracy: 0.6173333525657654 | Loss: 1.0856274762209506 | Acc: 0.6173333333333333.
Epoch: 7 |

Batch: 10000. Batch Size: 2. Training Loss (running): 1.055

MulticlassAccuracy: 0.6342083215713501 | Loss: 1.0114365832787007 | Acc: 0.6342083333333334.
Epoch: 8 |

Batch: 10000. Batch Size: 2. Training Loss (running): 0.979

MulticlassAccuracy: 0.6452500224113464 | Loss: 0.9519951666158935 | Acc: 0.6452500000000000.
Returned to Spot: Validation loss: 0.9519951666158935

```
test_tuned(net=model_spot, test_dataset=test,
            shuffle=False,
            loss_function=fun_control["loss_function"],
            metric=fun_control["metric_torch"],
            device = fun_control["device"],
            task=fun_control["task"])
```

MulticlassAccuracy: 0.6423000097274780 | Loss: 0.9578003666810692 | Acc: 0.6423000000000000.
Final evaluation: Validation loss: 0.9578003666810692
Final evaluation: Validation metric: 0.642300009727478

(0.9578003666810692, nan, tensor(0.6423))

11.10.5 Detailed Hyperparameter Plots

```
filename = "./figures/" + experiment_name
spot_tuner.plot_important_hyperparameter_contour(filename=filename)
```

l1: 100.0
batch_size: 2.3422106092726107
optimizer: 26.08413824178054









Figure 11.3: Contour plots.





11.10.6 Parallel Coordinates Plot

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Parallel coordinates plots

Unable to display output for mime type(s): text/html

11.10.7 Plot all Combinations of Hyperparameters

- Warning: this may take a while.

```
PLOT_ALL = False
if PLOT_ALL:
    n = spot_tuner.k
    for i in range(n-1):
        for j in range(i+1, n):
            spot_tuner.plot_contour(i=i, j=j, min_z=min_z, max_z = max_z)
```

12 HPT: PyTorch With cifar10 Data

In this tutorial, we will show how `spotPython` can be integrated into the PyTorch training workflow.

This document refers to the following software versions:

- python: 3.10.10
- torch: 2.0.1
- torchvision: 0.15.0

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

`spotPython` can be installed via `pip`. Alternatively, the source code can be downloaded from gitHub: <https://github.com/sequential-parameter-optimization/spotPython>.

```
!pip install spotPython
```

- Uncomment the following lines if you want to for (re-)installation the latest version of `spotPython` from gitHub.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

12.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time, initial design size and the device that is used.

 Caution: Run time and initial design size should be increased for real experiments

- MAX_TIME is set to one minute for demonstration purposes. For real experiments, this should be increased to at least 1 hour.
- INIT_SIZE is set to 5 for demonstration purposes. For real experiments, this should be increased to at least 10.

 Note: Device selection

- The device can be selected by setting the variable DEVICE.
- Since we are using a simple neural net, the setting "cpu" is preferred (on Mac).
- If you have a GPU, you can use "cuda:0" instead.
- If DEVICE is set to None, spotPython will automatically select the device.
 - This might result in "mps" on Macs, which is not the best choice for simple neural nets.

```
MAX_TIME = 1
INIT_SIZE = 5
DEVICE = "cpu" # "cuda:0" None
```

```
from spotPython.utils.device import getDevice
DEVICE = getDevice(DEVICE)
print(DEVICE)
```

cpu

```
import os
import copy
import socket
from datetime import datetime
from dateutil.tz import tzlocal
start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '12-torch' + "_" + HOSTNAME + "_" + str(MAX_TIME) + "min_" + str(INIT_SIZE)
experiment_name = experiment_name.replace(':', '-')
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')
```

12-torch_bartz09_1min_5init_2023-06-19_02-50-41

12.2 Step 2: Initialization of the Empty fun_control Dictionary

spotPython uses a Python dictionary for storing the information required for the hyperparameter tuning process, which was described in Section 14.2.

 Caution: Tensorboard does not work under Windows

- Since tensorboard does not work under Windows, we recommend setting the parameter `tensorboard_path` to `None` if you are working under Windows.

```
from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="classification",
    tensorboard_path="runs/12_spot_hpt_torch_cifar10",
    device=DEVICE)
```

12.3 Step 3: PyTorch Data Loading

12.3.1 Load Data Cifar10 Data

```
from torchvision import datasets, transforms
import torchvision
def load_data(data_dir="./data"):
    transform = transforms.Compose([
        transforms.ToTensor(),
        transforms.Normalize((0.5, 0.5, 0.5), (0.5, 0.5, 0.5))
    ])

    trainset = torchvision.datasets.CIFAR10(
        root=data_dir, train=True, download=True, transform=transform)

    testset = torchvision.datasets.CIFAR10(
        root=data_dir, train=False, download=True, transform=transform)

    return trainset, testset
train, test = load_data()
```

Files already downloaded and verified

Files already downloaded and verified

- Since this works fine, we can add the data loading to the `fun_control` dictionary:

```
n_samples = len(train)
# add the dataset to the fun_control
fun_control.update({"data": None, # dataset,
                  "train": train,
                  "test": test,
                  "n_samples": n_samples,
                  "target_column": None})
```

12.4 Step 4: Specification of the Preprocessing Model

After the training and test data are specified and added to the `fun_control` dictionary, `spotPython` allows the specification of a data preprocessing pipeline, e.g., for the scaling of the data or for the one-hot encoding of categorical variables, see Section 14.4. This feature is not used here, so we do not change the default value (which is `None`).

12.5 Step 5: Select Model (algorithm) and `core_model_hyper_dict`

12.5.1 Implementing a Configurable Neural Network With `spotPython`

`spotPython` includes the `Net_CIFAR10` class which is implemented in the file `netcifar10.py`. The class is imported here.

This class inherits from the class `Net_Core` which is implemented in the file `netcore.py`, see Section 14.5.1.

```
from spotPython.torch.netcifar10 import Net_CIFAR10
from spotPython.data.torch_hyper_dict import TorchHyperDict
from spotPython.hyperparameters.values import add_core_model_to_fun_control
fun_control = add_core_model_to_fun_control(core_model=Net_CIFAR10,
                                          fun_control=fun_control,
                                          hyper_dict=TorchHyperDict,
                                          filename=None)
```

12.5.2 The Search Space

12.5.3 Configuring the Search Space With spotPython

12.5.3.1 The hyper_dict Hyperparameters for the Selected Algorithm

spotPython uses JSON files for the specification of the hyperparameters, which were described in Section 14.5.5.

The corresponding entries for the `core_model` class are shown below.

```
fun_control['core_model_hyper_dict']

{'l1': {'type': 'int',
        'default': 5,
        'transform': 'transform_power_2_int',
        'lower': 2,
        'upper': 9},
 'l2': {'type': 'int',
        'default': 5,
        'transform': 'transform_power_2_int',
        'lower': 2,
        'upper': 9},
 'lr_mult': {'type': 'float',
             'default': 1.0,
             'transform': 'None',
             'lower': 0.1,
             'upper': 10.0},
 'batch_size': {'type': 'int',
                'default': 4,
                'transform': 'transform_power_2_int',
                'lower': 1,
                'upper': 4},
 'epochs': {'type': 'int',
            'default': 3,
            'transform': 'transform_power_2_int',
            'lower': 3,
            'upper': 4},
 'k_folds': {'type': 'int',
             'default': 1,
             'transform': 'None',
             'lower': 1,
```

```

    'upper': 1},
'patience': {'type': 'int',
             'default': 5,
             'transform': 'None',
             'lower': 2,
             'upper': 10},
'optimizer': {'levels': ['Adadelata',
                        'Adagrad',
                        'Adam',
                        'AdamW',
                        'SparseAdam',
                        'Adamax',
                        'ASGD',
                        'NAdam',
                        'RAdam',
                        'RMSprop',
                        'Rprop',
                        'SGD'],
             'type': 'factor',
             'default': 'SGD',
             'transform': 'None',
             'class_name': 'torch.optim',
             'core_model_parameter_type': 'str',
             'lower': 0,
             'upper': 12},
'sgd_momentum': {'type': 'float',
                 'default': 0.0,
                 'transform': 'None',
                 'lower': 0.0,
                 'upper': 1.0}}

```

12.6 Step 6: Modify `hyper_dict` Hyperparameters for the Selected Algorithm aka `core_model`

spotPython provides functions for modifying the hyperparameters, their bounds and factors as well as for activating and de-activating hyperparameters without re-compilation of the Python source code. These functions were described in [Section 14.6](#).

12.6.1 Step 5: Modify hyper_dict Hyperparameters for the Selected Algorithm aka core_model

12.6.1.1 Modify Hyperparameters of Type numeric and integer (boolean)

The hyperparameter `k_folds` is not used, it is de-activated here by setting the lower and upper bound to the same value.

 Caution: Small net size, number of epochs, and patience for demonstration purposes

- Net sizes 11 and 12 as well as `epochs` and `patience` are set to small values for demonstration purposes. These values are too small for a real application.
- More reasonable values are, e.g.:

```
- fun_control = modify_hyper_parameter_bounds(fun_control, "11",  
  bounds=[2, 7])  
- fun_control = modify_hyper_parameter_bounds(fun_control,  
  "epochs", bounds=[7, 9]) and  
- fun_control = modify_hyper_parameter_bounds(fun_control,  
  "patience", bounds=[2, 7])
```

```
from spotPython.hyperparameters.values import modify_hyper_parameter_bounds  
fun_control = modify_hyper_parameter_bounds(fun_control, "k_folds", bounds=[0, 0])  
fun_control = modify_hyper_parameter_bounds(fun_control, "patience", bounds=[2, 2])  
fun_control = modify_hyper_parameter_bounds(fun_control, "epochs", bounds=[2, 3])  
fun_control = modify_hyper_parameter_bounds(fun_control, "11", bounds=[2, 5])  
fun_control = modify_hyper_parameter_bounds(fun_control, "12", bounds=[2, 5])
```

12.6.2 Modify hyperparameter of type factor

```
from spotPython.hyperparameters.values import modify_hyper_parameter_levels  
fun_control = modify_hyper_parameter_levels(fun_control, "optimizer", ["Adam", "AdamW", "Ad
```

12.6.3 Optimizers

Optimizers can be selected as described in Section [19.6.2](#).

Optimizers are described in Section [14.6.1](#).

```
fun_control = modify_hyper_parameter_bounds(fun_control,
    "lr_mult", bounds=[1e-3, 1e-3])
fun_control = modify_hyper_parameter_bounds(fun_control,
    "sgd_momentum", bounds=[0.9, 0.9])
```

12.7 Step 7: Selection of the Objective (Loss) Function

12.7.1 Evaluation

The evaluation procedure requires the specification of two elements:

1. the way how the data is split into a train and a test set and
2. the loss function (and a metric).

These are described in Section [19.7.1](#).

The key "loss_function" specifies the loss function which is used during the optimization, see Section [14.7.5](#).

We will use CrossEntropy loss for the multiclass-classification task.

```
from torch.nn import CrossEntropyLoss
loss_function = CrossEntropyLoss()
fun_control.update({
    "loss_function": loss_function,
    "shuffle": True,
    "eval": "train_hold_out"
})
```

12.7.2 Metric

```
import torchmetrics
metric_torch = torchmetrics.Accuracy(task="multiclass",
    num_classes=10).to(fun_control["device"])
fun_control.update({"metric_torch": metric_torch})
```

12.8 Step 8: Calling the SPOT Function

12.8.1 Preparing the SPOT Call

The following code passes the information about the parameter ranges and bounds to `spot`.

```
# extract the variable types, names, and bounds
from spotPython.hyperparameters.values import (get_bound_values,
        get_var_name,
        get_var_type,)
var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
fun_control.update({"var_type": var_type,
                   "var_name": var_name})
lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")

from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))
```

name	type	default	lower	upper	transform
l1	int	5	2	5	transform_power_2_int
l2	int	5	2	5	transform_power_2_int
lr_mult	float	1.0	0.001	0.001	None
batch_size	int	4	1	4	transform_power_2_int
epochs	int	3	2	3	transform_power_2_int
k_folds	int	1	0	0	None
patience	int	5	2	2	None
optimizer	factor	SGD	0	3	None
sgd_momentum	float	0.0	0.9	0.9	None

12.8.2 The Objective Function `fun_torch`

The objective function `fun_torch` is selected next. It implements an interface from PyTorch's training, validation, and testing methods to `spotPython`.

```
from spotPython.fun.hypertorch import HyperTorch
fun = HyperTorch().fun_torch
```


MulticlassAccuracy: 0.0979000031948090 | Loss: 2.3118295257568358 | Acc: 0.0979000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.0984499976038933 | Loss: 2.3107973091125489 | Acc: 0.0984500000000000.
Epoch: 4 |

MulticlassAccuracy: 0.0979499965906143 | Loss: 2.3097029823303221 | Acc: 0.0979500000000000.
Epoch: 5 |

MulticlassAccuracy: 0.0974999964237213 | Loss: 2.3084314373016359 | Acc: 0.0975000000000000.
Epoch: 6 |

MulticlassAccuracy: 0.0978000015020370 | Loss: 2.3069485197067259 | Acc: 0.0978000000000000.
Epoch: 7 |

MulticlassAccuracy: 0.0994499996304512 | Loss: 2.3052204282760620 | Acc: 0.0994500000000000.
Epoch: 8 |

MulticlassAccuracy: 0.1077999994158745 | Loss: 2.3031790950775148 | Acc: 0.1078000000000000.
Returned to Spot: Validation loss: 2.303179095077515

config: {'l1': 8, 'l2': 8, 'lr_mult': 0.001, 'batch_size': 8, 'epochs': 4, 'k_folds': 0, 'pa
Epoch: 1 |

MulticlassAccuracy: 0.1019499972462654 | Loss: 2.3290661005973816 | Acc: 0.1019500000000000.
Epoch: 2 |

MulticlassAccuracy: 0.0920000001788139 | Loss: 2.3277609358787537 | Acc: 0.0920000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.0856499969959259 | Loss: 2.3266119544029236 | Acc: 0.0856500000000000.
Epoch: 4 |

MulticlassAccuracy: 0.0936999991536140 | Loss: 2.3255103428840638 | Acc: 0.0937000000000000.
Returned to Spot: Validation loss: 2.3255103428840638

config: {'l1': 32, 'l2': 16, 'lr_mult': 0.001, 'batch_size': 2, 'epochs': 8, 'k_folds': 0, 'p
Epoch: 1 |

MulticlassAccuracy: 0.1010999977588654 | Loss: 2.3073783339738845 | Acc: 0.1011000000000000.
Epoch: 2 |

MulticlassAccuracy: 0.1086999997496605 | Loss: 2.2941057253479959 | Acc: 0.1087000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.1623499989509583 | Loss: 2.2612908222079278 | Acc: 0.1623500000000000.
Epoch: 4 |

MulticlassAccuracy: 0.1683499962091446 | Loss: 2.2120475108027460 | Acc: 0.1683500000000000.
Epoch: 5 |

MulticlassAccuracy: 0.1793999969959259 | Loss: 2.1673231732010843 | Acc: 0.1794000000000000.
Epoch: 6 |

MulticlassAccuracy: 0.2073999941349030 | Loss: 2.1304403829693794 | Acc: 0.2074000000000000.
Epoch: 7 |

MulticlassAccuracy: 0.2124000042676926 | Loss: 2.1028314747035504 | Acc: 0.2124000000000000.
Epoch: 8 |

MulticlassAccuracy: 0.2163500040769577 | Loss: 2.0804890911757945 | Acc: 0.2163500000000000.
Returned to Spot: Validation loss: 2.0804890911757945

config: {'l1': 4, 'l2': 8, 'lr_mult': 0.001, 'batch_size': 4, 'epochs': 4, 'k_folds': 0, 'pa
Epoch: 1 |

MulticlassAccuracy: 0.0993499979376793 | Loss: 2.3089312206268309 | Acc: 0.0993500000000000.
Epoch: 2 |

MulticlassAccuracy: 0.0978500023484230 | Loss: 2.3069013849258422 | Acc: 0.0978500000000000.
Epoch: 3 |

MulticlassAccuracy: 0.0964500010013580 | Loss: 2.3048164315223696 | Acc: 0.0964500000000000.
Epoch: 4 |

MulticlassAccuracy: 0.1124999970197678 | Loss: 2.3030379461765289 | Acc: 0.1125000000000000.
Returned to Spot: Validation loss: 2.303037946176529

config: {'l1': 16, 'l2': 32, 'lr_mult': 0.001, 'batch_size': 8, 'epochs': 8, 'k_folds': 0, 'p
Epoch: 1 |

MulticlassAccuracy: 0.1054000034928322 | Loss: 2.3139406742095949 | Acc: 0.1054000000000000.
Epoch: 2 |

MulticlassAccuracy: 0.1093000024557114 | Loss: 2.3129244773864746 | Acc: 0.1093000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.1137500032782555 | Loss: 2.3119277452468872 | Acc: 0.1137500000000000.
Epoch: 4 |

MulticlassAccuracy: 0.1186999976634979 | Loss: 2.3108698651313784 | Acc: 0.1187000000000000.
Epoch: 5 |

MulticlassAccuracy: 0.1224000006914139 | Loss: 2.3096745619773866 | Acc: 0.1224000000000000.
Epoch: 6 |

MulticlassAccuracy: 0.1256500035524368 | Loss: 2.3082214976310729 | Acc: 0.1256500000000000.
Epoch: 7 |

MulticlassAccuracy: 0.1280000060796738 | Loss: 2.3064146327018737 | Acc: 0.1280000000000000.
Epoch: 8 |

MulticlassAccuracy: 0.1294499933719635 | Loss: 2.3041261745452881 | Acc: 0.1294500000000000.
Returned to Spot: Validation loss: 2.304126174545288

config: {'l1': 8, 'l2': 16, 'lr_mult': 0.001, 'batch_size': 8, 'epochs': 8, 'k_folds': 0, 'p'
Epoch: 1 |

MulticlassAccuracy: 0.1017500013113022 | Loss: 2.3090413111686705 | Acc: 0.1017500000000000.
Epoch: 2 |

MulticlassAccuracy: 0.1014999970793724 | Loss: 2.3063754611015321 | Acc: 0.1015000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.1022500023245811 | Loss: 2.3024991001129149 | Acc: 0.1022500000000000.
Epoch: 4 |

MulticlassAccuracy: 0.1098999977111816 | Loss: 2.2978994737625120 | Acc: 0.1099000000000000.
Epoch: 5 |

```
MulticlassAccuracy: 0.1268000006675720 | Loss: 2.2934811944007873 | Acc: 0.1268000000000000.  
Epoch: 6 |
```

```
MulticlassAccuracy: 0.1454499959945679 | Loss: 2.2888825453758241 | Acc: 0.1454500000000000.  
Epoch: 7 |
```

```
MulticlassAccuracy: 0.1518999934196472 | Loss: 2.2838717328071594 | Acc: 0.1519000000000000.  
Epoch: 8 |
```

```
MulticlassAccuracy: 0.1623000055551529 | Loss: 2.2781795486450194 | Acc: 0.1623000000000000.  
Returned to Spot: Validation loss: 2.2781795486450194  
spotPython tuning: 2.0804890911757945 [#####] 100.00% Done...
```

```
<spotPython.spot.spot.Spot at 0x2bbce6d10>
```

12.9 Step 9: Tensorboard

The textual output shown in the console (or code cell) can be visualized with Tensorboard as described in Section 14.9, see also the description in the documentation: [Tensorboard](#).

12.10 Step 10: Results

After the hyperparameter tuning run is finished, the results can be analyzed as described in Section 14.10.

```
SAVE = False  
LOAD = False  
  
if SAVE:  
    result_file_name = "res_" + experiment_name + ".pkl"  
    with open(result_file_name, 'wb') as f:  
        pickle.dump(spot_tuner, f)  
  
if LOAD:  
    result_file_name = "ADD THE NAME here, e.g.: res_ch10-friedman-hpt-0_maans03_60min_20i  
    with open(result_file_name, 'rb') as f:  
        spot_tuner = pickle.load(f)
```

After the hyperparameter tuning run is finished, the progress of the hyperparameter tuning can be visualized. The following code generates the progress plot from `?@fig-progress`.

```
spot_tuner.plot_progress(log_y=False,
                        filename="./figures/" + experiment_name+"_progress.png")
```




Figure 12.1: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

- Print the results

```
print(gen_design_table(fun_control=fun_control,
                      spot=spot_tuner))
```

name	type	default	lower	upper	tuned	transform
l1	int	5	2.0	5.0	5.0	transform_power_2_int
l2	int	5	2.0	5.0	4.0	transform_power_2_int
lr_mult	float	1.0	0.001	0.001	0.001	None
batch_size	int	4	1.0	4.0	1.0	transform_power_2_int
epochs	int	3	2.0	3.0	3.0	transform_power_2_int
k_folds	int	1	0.0	0.0	0.0	None
patience	int	5	2.0	2.0	2.0	None
optimizer	factor	SGD	0.0	3.0	3.0	None
sgd_momentum	float	0.0	0.9	0.9	0.9	None

12.10.1 Show variable importance

```
spot_tuner.plot_importance(threshold=0.025, filename="./figures/" + experiment_name+"_impo
```




Figure 12.2: Variable importance plot, threshold 0.025.

12.10.2 Get the Tuned Architecture (SPOT Results)

The architecture of the spotPython model can be obtained by the following code:

```
from spotPython.hyperparameters.values import get_one_core_model_from_X
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))
model_spot = get_one_core_model_from_X(X, fun_control)
model_spot
```

```
Net_CIFAR10(
  (conv1): Conv2d(3, 6, kernel_size=(5, 5), stride=(1, 1))
  (pool): MaxPool2d(kernel_size=2, stride=2, padding=0, dilation=1, ceil_mode=False)
  (conv2): Conv2d(6, 16, kernel_size=(5, 5), stride=(1, 1))
  (fc1): Linear(in_features=400, out_features=32, bias=True)
  (fc2): Linear(in_features=32, out_features=16, bias=True)
  (fc3): Linear(in_features=16, out_features=10, bias=True)
)
```

12.10.3 Evaluation of the Tuned Architecture

```
from spotPython.torch.traintest import (  
    train_tuned,  
    test_tuned,  
    )  
  
train_tuned(net=model_spot, train_dataset=train,  
            loss_function=fun_control["loss_function"],  
            metric=fun_control["metric_torch"],  
            shuffle=True,  
            device = fun_control["device"],  
            path=None,  
            task=fun_control["task"],)
```

Epoch: 1 |

Batch: 10000. Batch Size: 2. Training Loss (running): 2.308

MulticlassAccuracy: 0.1005500033497810 | Loss: 2.3031486242294310 | Acc: 0.1005500000000000.

Epoch: 2 |

Batch: 10000. Batch Size: 2. Training Loss (running): 2.300

MulticlassAccuracy: 0.1138999983668327 | Loss: 2.2927414273738860 | Acc: 0.1139000000000000.

Epoch: 3 |

Batch: 10000. Batch Size: 2. Training Loss (running): 2.287

MulticlassAccuracy: 0.1597000062465668 | Loss: 2.2735267727375033 | Acc: 0.1597000000000000.

Epoch: 4 |

Batch: 10000. Batch Size: 2. Training Loss (running): 2.266

MulticlassAccuracy: 0.2150499969720840 | Loss: 2.2502061260223387 | Acc: 0.2150500000000000.

Epoch: 5 |

Batch: 10000. Batch Size: 2. Training Loss (running): 2.243

MulticlassAccuracy: 0.2282000035047531 | Loss: 2.223145659089090 | Acc: 0.2282000000000000.
Epoch: 6 |

Batch: 10000. Batch Size: 2. Training Loss (running): 2.213

MulticlassAccuracy: 0.2391999959945679 | Loss: 2.1825061973452566 | Acc: 0.2392000000000000.
Epoch: 7 |

Batch: 10000. Batch Size: 2. Training Loss (running): 2.167

MulticlassAccuracy: 0.2432000041007996 | Loss: 2.1375639670968054 | Acc: 0.2432000000000000.
Epoch: 8 |

Batch: 10000. Batch Size: 2. Training Loss (running): 2.129

MulticlassAccuracy: 0.2525500059127808 | Loss: 2.0974310400962830 | Acc: 0.2525500000000000.
Returned to Spot: Validation loss: 2.097431040096283

If path is set to a filename, e.g., path = "model_spot_trained.pt", the weights of the trained model will be loaded from this file.

```
test_tuned(net=model_spot, test_dataset=test,
           shuffle=False,
           loss_function=fun_control["loss_function"],
           metric=fun_control["metric_torch"],
           device = fun_control["device"],
           task=fun_control["task"],)
```

MulticlassAccuracy: 0.2560999989509583 | Loss: 2.0896654647111892 | Acc: 0.2561000000000000.
Final evaluation: Validation loss: 2.089665464711189
Final evaluation: Validation metric: 0.25609999895095825

(2.089665464711189, nan, tensor(0.2561))

12.10.4 Cross-validated Evaluations

🔥 Caution: Cross-validated Evaluations

- The number of folds is set to 1 by default.
- Here it was changed to 3 for demonstration purposes.
- Set the number of folds to a reasonable value, e.g., 10.
- This can be done by setting the `k_folds` attribute of the model as follows:
- `setattr(model_spot, "k_folds", 10)`

```
from spotPython.torch.traintest import evaluate_cv
# modify k-folds:
setattr(model_spot, "k_folds", 3)
df_eval, df_preds, df_metrics = evaluate_cv(net=model_spot,
      dataset=fun_control["data"],
      loss_function=fun_control["loss_function"],
      metric=fun_control["metric_torch"],
      task=fun_control["task"],
      writer=fun_control["writer"],
      writerId="model_spot_cv",
      device = fun_control["device"])
```

Error in Net_Core. Call to evaluate_cv() failed. err=TypeError("Expected sequence or array-1

```
metric_name = type(fun_control["metric_torch"]).__name__
print(f"loss: {df_eval}, Cross-validated {metric_name}: {df_metrics}")
```

loss: nan, Cross-validated MulticlassAccuracy: nan

12.10.5 Detailed Hyperparameter Plots

```
filename = "./figures/" + experiment_name
spot_tuner.plot_important_hyperparameter_contour(filename=filename)
```

```
l1: 41.759595356993984
batch_size: 100.0
optimizer: 0.035981536916571195
```









Figure 12.3: Contour plots.





12.10.6 Parallel Coordinates Plot

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Parallel coordinates plots

Unable to display output for mime type(s): text/html

12.10.7 Plot all Combinations of Hyperparameters

- Warning: this may take a while.

```
PLOT_ALL = False
if PLOT_ALL:
    n = spot_tuner.k
    for i in range(n-1):
        for j in range(i+1, n):
            spot_tuner.plot_contour(i=i, j=j, min_z=min_z, max_z = max_z)
```

13 HPT: River

River is a Python library for online machine learning (Montiel et al. 2021). It aims to be the most user-friendly library for doing machine learning on streaming data. River is the result of a merger between creme and scikit-multiflow.

13.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time, initial design size and the device that is used.

 Caution: Run time and initial design size should be increased for real experiments

- MAX_TIME is set to one minute for demonstration purposes. For real experiments, this should be increased to at least 1 hour.
- INIT_SIZE is set to 5 for demonstration purposes. For real experiments, this should be increased to at least 10.
- K is set to 0.1 for demonstration purposes. For real experiments, this should be increased to at least 1.

```
MAX_TIME = 1
INIT_SIZE = 5
K = .1
```

10-river_bartz09_1min_5init_2023-06-19_03-04-54

13.1.1 river Hyperparameter Tuning: HATR with Friedman Drift Data

- This notebook exemplifies hyperparameter tuning with SPOT (spotPython and spotRiver).
- The hyperparameter software SPOT was developed in R (statistical programming language), see Open Access book “Hyperparameter Tuning for Machine and Deep Learning with R - A Practical Guide”, available here: <https://link.springer.com/book/10.1007/978-981-19-5170-1>.

- This notebook demonstrates hyperparameter tuning for `river`. It is based on the notebook “Incremental decision trees in river: the Hoeffding Tree case”, see: <https://riverml.xyz/0.15.0/recipes/on-hoeffding-trees/#42-regression-tree-splitters>.
- Here we will use the river HTR and HATR functions as in “Incremental decision trees in river: the Hoeffding Tree case”, see: <https://riverml.xyz/0.15.0/recipes/on-hoeffding-trees/#42-regression-tree-splitters>.

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

13.2 Step 2: Initialization of the `fun_control` Dictionary

```
from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="regression",
                               tensorboard_path=None)
```

13.3 Step 3: Load the Friedman Drift Data

```
horizon = 7*24
k = K
n_total = int(k*100_000)
n_samples = n_total
p_1 = int(k*25_000)
p_2 = int(k*50_000)
position=(p_1, p_2)
n_train = 1_000
a = n_train + p_1 - 12
b = a + 12
```

- Since we also need a `river` version of the data below for plotting the model, the corresponding data set is generated here. Note: `spotRiver` uses the `train` and `test` data sets, while `river` uses the `X` and `y` data sets

```

from river.datasets import synth
import pandas as pd
dataset = synth.FriedmanDrift(
    drift_type='gra',
    position=position,
    seed=123
)
data_dict = {key: [] for key in list(dataset.take(1))[0][0].keys()}
data_dict["y"] = []
for x, y in dataset.take(n_total):
    for key, value in x.items():
        data_dict[key].append(value)
    data_dict["y"].append(y)
df = pd.DataFrame(data_dict)
# Add column names x1 until x10 to the first 10 columns of the dataframe and the column name y
df.columns = [f"x{i}" for i in range(1, 11)] + ["y"]

train = df[:n_train]
test = df[n_train:]
target_column = "y"
#
fun_control.update({"data": None, # dataset,
                  "train": train,
                  "test": test,
                  "n_samples": n_samples,
                  "target_column": target_column})

```

13.4 Step 4: Specification of the Preprocessing Model

```

from river import preprocessing
prep_model = preprocessing.StandardScaler()
fun_control.update({"prep_model": prep_model})

```

13.5 Step 5: Select algorithm and core_model_hyper_dict

- The river model (HATR) is selected.
- Furthermore, the corresponding hyperparameters, see: <https://riverml.xyz/0.15.0/api/tree/HoeffdingTreeRegressor/> are selected (incl. type information, names, and bounds).
- The corresponding hyperparameter dictionary is added to the fun_control dictionary.
- Alternatively, you can load a local hyper_dict. Simply set river_hyper_dict.json as the filename. If filename is set to None, the hyper_dict is loaded from the spotRiver package.

```
from river.tree import HoeffdingAdaptiveTreeRegressor
from spotRiver.data.river_hyper_dict import RiverHyperDict
from spotPython.hyperparameters.values import add_core_model_to_fun_control
core_model = HoeffdingAdaptiveTreeRegressor
fun_control = add_core_model_to_fun_control(core_model=core_model,
                                          fun_control=fun_control,
                                          hyper_dict=RiverHyperDict,
                                          filename=None)
```

The corresponding entries for the core_model class are shown below.

```
fun_control['core_model_hyper_dict']
```

```
{'grace_period': {'type': 'int',
                  'default': 200,
                  'transform': 'None',
                  'lower': 10,
                  'upper': 1000},
 'max_depth': {'type': 'int',
                'default': 20,
                'transform': 'transform_power_2_int',
                'lower': 2,
                'upper': 20},
 'delta': {'type': 'float',
            'default': 1e-07,
            'transform': 'None',
            'lower': 1e-08,
            'upper': 1e-06},
 'tau': {'type': 'float',
          'default': 0.05,
          'transform': 'None',
          'lower': 0.01,
```

```

'upper': 0.1},
'leaf_prediction': {'levels': ['mean', 'model', 'adaptive'],
'   type': 'factor',
'   default': 'mean',
'   transform': 'None',
'   core_model_parameter_type': 'str',
'   lower': 0,
'   upper': 2},
'leaf_model': {'levels': ['LinearRegression', 'PAREgressor', 'Perceptron'],
'   type': 'factor',
'   default': 'LinearRegression',
'   transform': 'None',
'   class_name': 'river.linear_model',
'   core_model_parameter_type': 'instance()',
'   lower': 0,
'   upper': 2},
'model_selector_decay': {'type': 'float',
'   default': 0.95,
'   transform': 'None',
'   lower': 0.9,
'   upper': 0.99},
'splitter': {'levels': ['EBSTSplitter', 'TEBSTSplitter', 'QOSplitter'],
'   type': 'factor',
'   default': 'EBSTSplitter',
'   transform': 'None',
'   class_name': 'river.tree.splitter',
'   core_model_parameter_type': 'instance()',
'   lower': 0,
'   upper': 2},
'min_samples_split': {'type': 'int',
'   default': 5,
'   transform': 'None',
'   lower': 2,
'   upper': 10},
'bootstrap_sampling': {'levels': [0, 1],
'   type': 'factor',
'   default': 0,
'   transform': 'None',
'   core_model_parameter_type': 'bool',
'   lower': 0,
'   upper': 1},
'drift_window_threshold': {'type': 'int',
'   default': 300,

```

```

'transform': 'None',
'lower': 100,
'upper': 500},
'switch_significance': {'type': 'float',
'default': 0.05,
'transform': 'None',
'lower': 0.01,
'upper': 0.1},
'binary_split': {'levels': [0, 1],
'type': 'factor',
'default': 0,
'transform': 'None',
'core_model_parameter_type': 'bool',
'lower': 0,
'upper': 1},
'max_size': {'type': 'float',
'default': 500.0,
'transform': 'None',
'lower': 100.0,
'upper': 1000.0},
'memory_estimate_period': {'type': 'int',
'default': 1000000,
'transform': 'None',
'lower': 100000,
'upper': 1000000},
'stop_mem_management': {'levels': [0, 1],
'type': 'factor',
'default': 0,
'transform': 'None',
'core_model_parameter_type': 'bool',
'lower': 0,
'upper': 1},
'remove_poor_attrs': {'levels': [0, 1],
'type': 'factor',
'default': 0,
'transform': 'None',
'core_model_parameter_type': 'bool',
'lower': 0,
'upper': 1},
'merit_preprune': {'levels': [0, 1],
'type': 'factor',
'default': 0,
'transform': 'None',

```

```
'core_model_parameter_type': 'bool',  
'lower': 0,  
'upper': 1}}
```

13.6 Step 6: Modify hyper_dict Hyperparameters for the Selected Algorithm aka core_model

13.6.1 Modify hyperparameter of type factor

```
# fun_control = modify_hyper_parameter_levels(fun_control, "leaf_model", ["LinearRegression", "LogisticRegression"])  
# fun_control["core_model_hyper_dict"]
```

13.6.2 Modify hyperparameter of type numeric and integer (boolean)

```
from spotPython.hyperparameters.values import modify_hyper_parameter_bounds  
fun_control = modify_hyper_parameter_bounds(fun_control, "delta", bounds=[1e-10, 1e-6])  
# fun_control = modify_hyper_parameter_bounds(fun_control, "min_samples_split", bounds=[3, 10])  
fun_control = modify_hyper_parameter_bounds(fun_control, "merit_preprune", [0, 0])
```

13.7 Step 7: Selection of the Objective (Loss) Function

There are three metrics:

1. ``metric_river`` is used for the river based evaluation via ``eval_oml_iter_progressive``.
2. ``metric_sklearn`` is used for the sklearn based evaluation via ``eval_oml_horizon``.
3. ``metric_torch`` is used for the pytorch based evaluation.

```
import numpy as np  
from river import metrics  
from sklearn.metrics import mean_absolute_error  
  
from spotRiver.fun.hyperriver import HyperRiver  
fun = HyperRiver(seed=123, log_level=50).fun_oml_horizon  
weights = np.array([1, 1/1000, 1/1000])*10_000.0  
horizon = 7*24
```

```

oml_grace_period = 2
step = 100
weight_coeff = 1.0

fun_control.update({
    "horizon": horizon,
    "oml_grace_period": oml_grace_period,
    "weights": weights,
    "step": step,
    "log_level": 50,
    "weight_coeff": weight_coeff,
    "metric_river": metrics.MAE(),
    "metric_sklearn": mean_absolute_error
})

```

13.8 Step 8: Calling the SPOT Function

13.8.1 Prepare the SPOT Parameters

- Get types and variable names as well as lower and upper bounds for the hyperparameters.

```

from spotPython.hyperparameters.values import (
    get_var_type,
    get_var_name,
    get_bound_values
)

var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
fun_control.update({"var_type": var_type,
                   "var_name": var_name})

lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")

from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))

```

name	type	default	lower	upper	transform
-----	-----	-----	-----	-----	-----

grace_period	int	200		10		1000		None
max_depth	int	20		2		20		transform_pow
delta	float	1e-07		1e-10		1e-06		None
tau	float	0.05		0.01		0.1		None
leaf_prediction	factor	mean		0		2		None
leaf_model	factor	LinearRegression		0		2		None
model_selector_decay	float	0.95		0.9		0.99		None
splitter	factor	EBSTSplitter		0		2		None
min_samples_split	int	5		2		10		None
bootstrap_sampling	factor	0		0		1		None
drift_window_threshold	int	300		100		500		None
switch_significance	float	0.05		0.01		0.1		None
binary_split	factor	0		0		1		None
max_size	float	500.0		100		1000		None
memory_estimate_period	int	1000000		100000		1e+06		None
stop_mem_management	factor	0		0		1		None
remove_poor_attrs	factor	0		0		1		None
merit_preprune	factor	0		0		0		None

13.8.2 Run the Spot Optimizer

- Run SPOT for approx. x mins (max_time).
- Note: the run takes longer, because the evaluation time of initial design (here: initi_size, 20 points) is not considered.

```
from spotPython.hyperparameters.values import get_default_hyperparameters_as_array
hyper_dict=RiverHyperDict().load()
X_start = get_default_hyperparameters_as_array(fun_control, hyper_dict)
```

```
from spotPython.spot import spot
from math import inf
import numpy as np
spot_tuner = spot.Spot(fun=fun,
                      lower = lower,
                      upper = upper,
                      fun_evals = inf,
                      fun_repeats = 1,
                      max_time = MAX_TIME,
                      noise = False,
                      tolerance_x = np.sqrt(np.spacing(1)),
                      var_type = var_type,
```

```

        var_name = var_name,
        infill_criterion = "y",
        n_points = 1,
        seed=123,
        log_level = 50,
        show_models= False,
        show_progress= True,
        fun_control = fun_control,
        design_control={"init_size": INIT_SIZE,
                        "repeats": 1},
        surrogate_control={"noise": True,
                           "cod_type": "norm",
                           "min_theta": -4,
                           "max_theta": 3,
                           "n_theta": len(var_name),
                           "model_fun_evals": 10_000,
                           "log_level": 50
                          })

spot_tuner.run(X_start=X_start)

```

spotPython tuning: 2.145751611487961 [##-----] 15.43%

spotPython tuning: 2.145751611487961 [###-----] 34.06%

spotPython tuning: 2.145751611487961 [#####-----] 51.79%

spotPython tuning: 2.145751611487961 [#####] 100.00% Done...

<spotPython.spot.spot.Spot at 0x2a6e36b00>

13.9 Step 9: Results

```

import pickle
SAVE = False
LOAD = False

if SAVE:
    result_file_name = "res_" + experiment_name + ".pkl"

```

```


with open(result_file_name, 'wb') as f:
    pickle.dump(spot_tuner, f)

if LOAD:
    result_file_name = "res_ch10-friedman-hpt-0_maans03_60min_20init_1K_2023-04-14_10-11-1"
    with open(result_file_name, 'rb') as f:
        spot_tuner = pickle.load(f)

```

- Show the Progress of the hyperparameter tuning:

```
spot_tuner.plot_progress(log_y=True, filename="./figures/" + experiment_name+"_progress.pdf")
```



- Print the Results


```
print(gen_design_table(fun_control=fun_control, spot=spot_tuner))
```

name	type	default	lower	upper	
grace_period	int	200	10.0	1000.0	
max_depth	int	20	2.0	20.0	
delta	float	1e-07	1e-10	1e-06	4.068723023437
tau	float	0.05	0.01	0.1	0.0484260091
leaf_prediction	factor	mean	0.0	2.0	
leaf_model	factor	LinearRegression	0.0	2.0	
model_selector_decay	float	0.95	0.9	0.99	0.970713237
splitter	factor	EBSTSplitter	0.0	2.0	
min_samples_split	int	5	2.0	10.0	
bootstrap_sampling	factor	0	0.0	1.0	
drift_window_threshold	int	300	100.0	500.0	

switch_significance	float	0.05		0.01		0.1		0.040370639
binary_split	factor	0		0.0		1.0		
max_size	float	500.0		100.0		1000.0		454.140654
memory_estimate_period	int	1000000		100000.0		1000000.0		9
stop_mem_management	factor	0		0.0		1.0		
remove_poor_attrs	factor	0		0.0		1.0		
merit_preprune	factor	0		0.0		0.0		

13.9.1 Show variable importance

```
spot_tuner.plot_importance(threshold=0.0025, filename="./figures/" + experiment_name+"_imp
```



13.9.2 Build and Evaluate HTR Model with Tuned Hyperparameters

```
m = test.shape[0]
a = int(m/2)-50
b = int(m/2)
```

13.9.3 The Large Data Set (k=0.2)

Caution: Increased Friedman-Drift Data Set

- The Friedman-Drift Data Set is increased by a factor of two to show the transferability of the hyperparameter tuning results.
- Larger values of k lead to a longer run time.

```
horizon = 7*24
k = .2
n_total = int(k*100_000)
n_samples = n_total
p_1 = int(k*25_000)
p_2 = int(k*50_000)
position=(p_1, p_2)
n_train = 1_000
a = n_train + p_1 - 12
b = a + 12
dataset = synth.FriedmanDrift(
    drift_type='gra',
    position=position,
    seed=123
)
data_dict = {key: [] for key in list(dataset.take(1))[0][0].keys()}
data_dict["y"] = []
for x, y in dataset.take(n_total):
    for key, value in x.items():
        data_dict[key].append(value)
    data_dict["y"].append(y)
df = pd.DataFrame(data_dict)
# Add column names x1 until x10 to the first 10 columns of the dataframe and the column name
df.columns = [f"x{i}" for i in range(1, 11)] + ["y"]

train = df[:n_train]
test = df[n_train:]
target_column = "y"
#
fun_control.update({"data": None, # dataset,
                   "train": train,
                   "test": test,
                   "n_samples": n_samples,
```

```
"target_column": target_column})
```

13.9.4 Get Default Hyperparameters

```
# fun_control was modified, we generate a new one with the original
# default hyperparameters
from spotPython.hyperparameters.values import get_one_core_model_from_X
fc = fun_control
fc.update({"core_model_hyper_dict":
          hyper_dict[fun_control["core_model"].__name__]})
model_default = get_one_core_model_from_X(X_start, fun_control=fc)
model_default
```

```
HoeffdingAdaptiveTreeRegressor (
  grace_period=200
  max_depth=1048576
  delta=1e-07
  tau=0.05
  leaf_prediction="mean"
  leaf_model=LinearRegression (
    optimizer=SGD (
      lr=Constant (
        learning_rate=0.01
      )
    )
    loss=Squared ()
    l2=0.
    l1=0.
    intercept_init=0.
    intercept_lr=Constant (
      learning_rate=0.01
    )
    clip_gradient=1e+12
    initializer=Zeros ()
  )
  model_selector_decay=0.95
  nominal_attributes=None
  splitter=EBSTSplitter ()
  min_samples_split=5
  bootstrap_sampling=0
```

```

drift_window_threshold=300
drift_detector=ADWIN (
    delta=0.002
    clock=32
    max_buckets=5
    min_window_length=5
    grace_period=10
)
switch_significance=0.05
binary_split=0
max_size=500.
memory_estimate_period=1000000
stop_mem_management=0
remove_poor_attrs=0
merit_preprune=0
seed=None
)

```

```

from spotRiver.evaluation.eval_bml import eval_oml_horizon

```

```

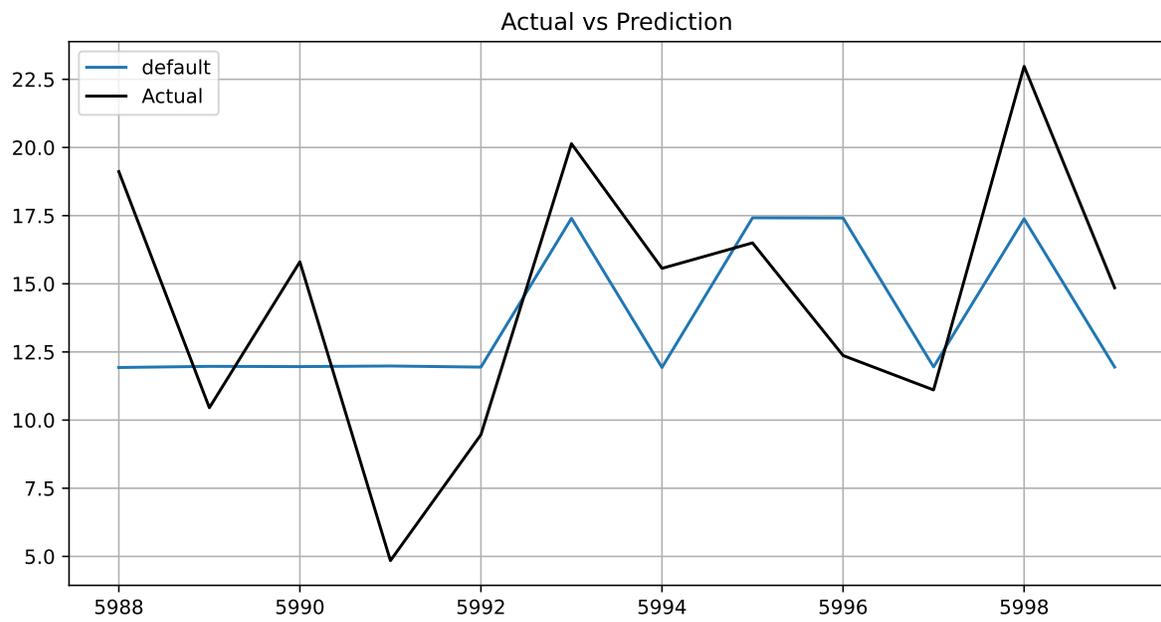

df_eval_default, df_true_default = eval_oml_horizon(
    model=model_default,
    train=fun_control["train"],
    test=fun_control["test"],
    target_column=fun_control["target_column"],
    horizon=fun_control["horizon"],
    oml_grace_period=fun_control["oml_grace_period"],
    metric=fun_control["metric_sklearn"],
)

```

```

from spotRiver.evaluation.eval_bml import plot_bml_oml_horizon_metrics, plot_bml_oml_horizon_predictions
df_labels=["default"]
plot_bml_oml_horizon_metrics(df_eval = [df_eval_default], log_y=False, df_labels=df_labels)
plot_bml_oml_horizon_predictions(df_true = [df_true_default[a:b]], target_column=target_co

```



13.9.5 Get SPOT Results

```
from spotPython.hyperparameters.values import get_one_core_model_from_X
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))
model_spot = get_one_core_model_from_X(X, fun_control)
model_spot
```

```
HoeffdingAdaptiveTreeRegressor (
  grace_period=657
  max_depth=256
  delta=4e-08
  tau=0.048426
  leaf_prediction="adaptive"
  leaf_model=LinearRegression (
    optimizer=SGD (
      lr=Constant (
        learning_rate=0.01
      )
    )
    loss=Squared ()
    l2=0.
    l1=0.
    intercept_init=0.
    intercept_lr=Constant (
      learning_rate=0.01
    )
    clip_gradient=1e+12
    initializer=Zeros ()
  )
  model_selector_decay=0.970713
  nominal_attributes=None
  splitter=QOSplitter (
    radius=0.25
    allow_multiway_splits=False
  )
  min_samples_split=5
  bootstrap_sampling=1
  drift_window_threshold=166
  drift_detector=ADWIN (
    delta=0.002
    clock=32
    max_buckets=5
```

```

    min_window_length=5
    grace_period=10
)
switch_significance=0.040371
binary_split=0
max_size=454.140654
memory_estimate_period=910594
stop_mem_management=1
remove_poor_attrs=1
merit_preprune=0
seed=None
)

```

```


df_eval_spot, df_true_spot = eval_oml_horizon(
    model=model_spot,
    train=fun_control["train"],
    test=fun_control["test"],
    target_column=fun_control["target_column"],
    horizon=fun_control["horizon"],
    oml_grace_period=fun_control["oml_grace_period"],
    metric=fun_control["metric_sklearn"],
)

```

```

df_labels=["default", "spot"]
plot_bml_oml_horizon_metrics(df_eval = [df_eval_default, df_eval_spot], log_y=False, df_la


```



```


a = int(m/2)+20
b = int(m/2)+50
plot_bml_oml_horizon_predictions(df_true = [df_true_default[a:b], df_true_spot[a:b]], targ

```




```
from spotPython.plot.validation import plot_actual_vs_predicted
plot_actual_vs_predicted(y_test=df_true_default["y"], y_pred=df_true_default["Prediction"])
plot_actual_vs_predicted(y_test=df_true_spot["y"], y_pred=df_true_spot["Prediction"], titl
```

Default



SPOT



13.9.6 Visualize Regression Trees

```
dataset_f = dataset.take(n_total)
for x, y in dataset_f:
    model_default.learn_one(x, y)
```

Caution: Large Trees

- Since the trees are large, the visualization is suppressed by default.
- To visualize the trees, uncomment the following line.

```
# model_default.draw()
```

```
model_default.summary
```

```
{'n_nodes': 35,
  'n_branches': 17,
  'n_leaves': 18,
  'n_active_leaves': 96,
  'n_inactive_leaves': 0,
  'height': 6,
  'total_observed_weight': 39002.0,
  'n_alternate_trees': 21,
  'n_pruned_alternate_trees': 6,
  'n_switch_alternate_trees': 2}
```

13.9.7 Spot Model

```
dataset_f = dataset.take(n_total)
for x, y in dataset_f:
    model_spot.learn_one(x, y)
```

Caution: Large Trees

- Since the trees are large, the visualization is suppressed by default.
- To visualize the trees, uncomment the following line.

```
# model_spot.draw()
```

```
model_spot.summary
```

```
{'n_nodes': 25,  
 'n_branches': 12,  
 'n_leaves': 13,  
 'n_active_leaves': 55,  
 'n_inactive_leaves': 0,  
 'height': 6,  
 'total_observed_weight': 39002.0,  
 'n_alternate_trees': 24,  
 'n_pruned_alternate_trees': 10,  
 'n_switch_alternate_trees': 0}
```

```
from spotPython.utils.eda import compare_two_tree_models  
print(compare_two_tree_models(model_default, model_spot))
```

Parameter	Default	Spot
n_nodes	35	25
n_branches	17	12
n_leaves	18	13
n_active_leaves	96	55
n_inactive_leaves	0	0
height	6	6
total_observed_weight	39002	39002
n_alternate_trees	21	24
n_pruned_alternate_trees	6	10
n_switch_alternate_trees	2	0


```
min(spot_tuner.y), max(spot_tuner.y)
```

```
(2.145751611487961, 13.362638642210635)
```

13.9.8 Detailed Hyperparameter Plots

```
filename = "./figures/" + experiment_name
spot_tuner.plot_important_hyperparameter_contour(filename=filename)
```

```
leaf_model: 100.0
memory_estimate_period: 0.1638816497501305
```



13.9.9 Parallel Coordinates Plots

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Unable to display output for mime type(s): text/html

13.9.10 Plot all Combinations of Hyperparameters

- Warning: this may take a while.

```
PLOT_ALL = False
if PLOT_ALL:
    n = spot_tuner.k
```

```
for i in range(n-1):
    for j in range(i+1, n):
        spot_tuner.plot_contour(i=i, j=j, min_z=min_z, max_z = max_z)
```

14 HPT: PyTorch With spotPython and Ray Tune on CIFAR10

In this tutorial, we will show how `spotPython` can be integrated into the PyTorch training workflow. It is based on the tutorial “Hyperparameter Tuning with Ray Tune” from the PyTorch documentation (PyTorch 2023a), which is an extension of the tutorial “Training a Classifier” (PyTorch 2023b) for training a CIFAR10 image classifier.

This document refers to the following software versions:

- python: 3.10.10
- torch: 2.0.1
- torchvision: 0.15.0

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

`spotPython` can be installed via `pip`¹.

```
!pip install spotPython
```

- Uncomment the following lines if you want to for (re-)installation the latest version of `spotPython` from gitHub.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

¹Alternatively, the source code can be downloaded from gitHub: <https://github.com/sequential-parameter-optimization/spotPython>.

Results that refer to the Ray Tune package are taken from https://PyTorch.org/tutorials/beginner/hyperparameter_tuning_tutorial.html².

14.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time, initial design size and the device that is used.

 Caution: Run time and initial design size should be increased for real experiments

- MAX_TIME is set to one minute for demonstration purposes. For real experiments, this should be increased to at least 1 hour.
- INIT_SIZE is set to 5 for demonstration purposes. For real experiments, this should be increased to at least 10.

 Note: Device selection

- The device can be selected by setting the variable DEVICE.
- Since we are using a simple neural net, the setting "cpu" is preferred (on Mac).
- If you have a GPU, you can use "cuda:0" instead.
- If DEVICE is set to None, spotPython will automatically select the device.
 - This might result in "mps" on Macs, which is not the best choice for simple neural nets.

```
MAX_TIME = 10
INIT_SIZE = 5
DEVICE = "cpu" # "cuda:0"
```

```
from spotPython.utils.device import getDevice
DEVICE = getDevice(DEVICE)
print(DEVICE)
```

cpu

```
import os
import copy
import socket
```

²We were not able to install Ray Tune on our system. Therefore, we used the results from the PyTorch tutorial.

```

import warnings
from datetime import datetime
from dateutil.tz import tzlocal
start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '14-torch' + "_" + HOSTNAME + "_" + str(MAX_TIME) + "min_" + str(INIT_SECONDS)
experiment_name = experiment_name.replace(':', '-')
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')
warnings.filterwarnings("ignore")

```

14-torch_bartz09_10min_5init_2023-06-19_03-33-31

14.2 Step 2: Initialization of the `fun_control` Dictionary

`spotPython` uses a Python dictionary for storing the information required for the hyperparameter tuning process. This dictionary is called `fun_control` and is initialized with the function `fun_control_init`. The function `fun_control_init` returns a skeleton dictionary. The dictionary is filled with the required information for the hyperparameter tuning process. It stores the hyperparameter tuning settings, e.g., the deep learning network architecture that should be tuned, the classification (or regression) problem, and the data that is used for the tuning. The dictionary is used as an input for the `SPOT` function.

 Caution: Tensorboard does not work under Windows

- Since tensorboard does not work under Windows, we recommend setting the parameter `tensorboard_path` to `None` if you are working under Windows.

```

from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="classification",
    tensorboard_path="runs/14_spot_ray_hpt_torch_cifar10",
    device=DEVICE,)

```

14.3 Step 3: PyTorch Data Loading

The data loading process is implemented in the same manner as described in the Section “Data loaders” in PyTorch (2023a). The data loaders are wrapped into the function

`load_data_cifar10` which is identical to the function `load_data` in PyTorch (2023a). A global data directory is used, which allows sharing the data directory between different trials. The method `load_data_cifar10` is part of the `spotPython` package and can be imported from `spotPython.data.torchdata`.

In the following step, the test and train data are added to the dictionary `fun_control`.

```
from spotPython.data.torchdata import load_data_cifar10
train, test = load_data_cifar10()
n_samples = len(train)
# add the dataset to the fun_control
fun_control.update({
    "train": train,
    "test": test,
    "n_samples": n_samples})
```

Files already downloaded and verified

Files already downloaded and verified

14.4 Step 4: Specification of the Preprocessing Model

After the training and test data are specified and added to the `fun_control` dictionary, `spotPython` allows the specification of a data preprocessing pipeline, e.g., for the scaling of the data or for the one-hot encoding of categorical variables. The preprocessing model is called `prep_model` (“preparation” or pre-processing) and includes steps that are not subject to the hyperparameter tuning process. The preprocessing model is specified in the `fun_control` dictionary. The preprocessing model can be implemented as a `sklearn` pipeline. The following code shows a typical preprocessing pipeline:

```
categorical_columns = ["cities", "colors"]
one_hot_encoder = OneHotEncoder(handle_unknown="ignore",
                                sparse_output=False)

prep_model = ColumnTransformer(
    transformers=[
        ("categorical", one_hot_encoder, categorical_columns),
    ],
    remainder=StandardScaler(),
)
```

Because the Ray Tune (`ray[tune]`) hyperparameter tuning as described in PyTorch (2023a) does not use a preprocessing model, the preprocessing model is set to `None` here.

```
prep_model = None
fun_control.update({"prep_model": prep_model})
```

14.5 Step 5: Select Model (algorithm) and `core_model_hyper_dict`

The same neural network model as implemented in the section “Configurable neural network” of the PyTorch tutorial (PyTorch 2023a) is used here. We will show the implementation from PyTorch (2023a) in Section 14.5.0.1 first, before the extended implementation with `spotPython` is shown in Section 14.5.0.2.

14.5.0.1 Implementing a Configurable Neural Network With Ray Tune

We used the same hyperparameters that are implemented as configurable in the PyTorch tutorial. We specify the layer sizes, namely 11 and 12, of the fully connected layers:

```
class Net(nn.Module):
    def __init__(self, l1=120, l2=84):
        super(Net, self).__init__()
        self.conv1 = nn.Conv2d(3, 6, 5)
        self.pool = nn.MaxPool2d(2, 2)
        self.conv2 = nn.Conv2d(6, 16, 5)
        self.fc1 = nn.Linear(16 * 5 * 5, l1)
        self.fc2 = nn.Linear(l1, l2)
        self.fc3 = nn.Linear(l2, 10)

    def forward(self, x):
        x = self.pool(F.relu(self.conv1(x)))
        x = self.pool(F.relu(self.conv2(x)))
        x = x.view(-1, 16 * 5 * 5)
        x = F.relu(self.fc1(x))
        x = F.relu(self.fc2(x))
        x = self.fc3(x)
        return x
```

The learning rate, i.e., `lr`, of the optimizer is made configurable, too:

```
optimizer = optim.SGD(net.parameters(), lr=config["lr"], momentum=0.9)
```

14.5.0.2 Implementing a Configurable Neural Network With spotPython

spotPython implements a class which is similar to the class described in the PyTorch tutorial. The class is called `Net_CIFAR10` and is implemented in the file `netcifar10.py`.

```
from torch import nn
import torch.nn.functional as F
import spotPython.torch.netcore as netcore

class Net_CIFAR10(netcore.Net_Core):
    def __init__(self, l1, l2, lr_mult, batch_size, epochs, k_folds, patience,
optimizer, sgd_momentum):
        super(Net_CIFAR10, self).__init__(
            lr_mult=lr_mult,
            batch_size=batch_size,
            epochs=epochs,
            k_folds=k_folds,
            patience=patience,
            optimizer=optimizer,
            sgd_momentum=sgd_momentum,
        )
        self.conv1 = nn.Conv2d(3, 6, 5)
        self.pool = nn.MaxPool2d(2, 2)
        self.conv2 = nn.Conv2d(6, 16, 5)
        self.fc1 = nn.Linear(16 * 5 * 5, 11)
        self.fc2 = nn.Linear(11, 12)
        self.fc3 = nn.Linear(12, 10)

    def forward(self, x):
        x = self.pool(F.relu(self.conv1(x)))
        x = self.pool(F.relu(self.conv2(x)))
        x = x.view(-1, 16 * 5 * 5)
        x = F.relu(self.fc1(x))
        x = F.relu(self.fc2(x))
        x = self.fc3(x)
        return x
```

14.5.1 The Net_Core class

`Net_CIFAR10` inherits from the class `Net_Core` which is implemented in the file `netcore.py`. It implements the additional attributes that are common to all neural network models. The `Net_Core` class is implemented in the file `netcore.py`. It implements hyperparameters as attributes, that are not used by the `core_model`, e.g.:

- optimizer (`optimizer`),
- learning rate (`lr`),
- batch size (`batch_size`),
- epochs (`epochs`),
- k_folds (`k_folds`), and
- early stopping criterion “patience” (`patience`).

Users can add further attributes to the class. The class `Net_Core` is shown below.

```
from torch import nn

class Net_Core(nn.Module):
    def __init__(self, lr_mult, batch_size, epochs, k_folds, patience,
                 optimizer, sgd_momentum):
        super(Net_Core, self).__init__()
        self.lr_mult = lr_mult
        self.batch_size = batch_size
        self.epochs = epochs
        self.k_folds = k_folds
        self.patience = patience
        self.optimizer = optimizer
        self.sgd_momentum = sgd_momentum
```

14.5.2 Comparison of the Approach Described in the PyTorch Tutorial With spotPython

Comparing the class `Net` from the PyTorch tutorial and the class `Net_CIFAR10` from `spotPython`, we see that the class `Net_CIFAR10` has additional attributes and does not inherit from `nn` directly. It adds an additional class, `Net_core`, that takes care of additional attributes that are common to all neural network models, e.g., the learning rate multiplier `lr_mult` or the batch size `batch_size`.

`spotPython`'s `core_model` implements an instance of the `Net_CIFAR10` class. In addition to the basic neural network model, the `core_model` can use these additional attributes. `spotPython`

provides methods for handling these additional attributes to guarantee 100% compatibility with the PyTorch classes. The method `add_core_model_to_fun_control` adds the hyperparameters and additional attributes to the `fun_control` dictionary. The method is shown below.

```
from spotPython.torch.netcifar10 import Net_CIFAR10
from spotPython.data.torch_hyper_dict import TorchHyperDict
from spotPython.hyperparameters.values import add_core_model_to_fun_control
core_model = Net_CIFAR10
fun_control = add_core_model_to_fun_control(core_model=core_model,
                                           fun_control=fun_control,
                                           hyper_dict=TorchHyperDict,
                                           filename=None)
```

14.5.3 The Search Space: Hyperparameters

In Section 14.5.4, we first describe how to configure the search space with `ray[tune]` (as shown in PyTorch (2023a)) and then how to configure the search space with `spotPython` in -14.

14.5.4 Configuring the Search Space With Ray Tune

Ray Tune's search space can be configured as follows (PyTorch 2023a):

```
config = {
    "l1": tune.sample_from(lambda _: 2**np.random.randint(2, 9)),
    "l2": tune.sample_from(lambda _: 2**np.random.randint(2, 9)),
    "lr": tune.loguniform(1e-4, 1e-1),
    "batch_size": tune.choice([2, 4, 8, 16])
}
```

The `tune.sample_from()` function enables the user to define sample methods to obtain hyperparameters. In this example, the `l1` and `l2` parameters should be powers of 2 between 4 and 256, so either 4, 8, 16, 32, 64, 128, or 256. The `lr` (learning rate) should be uniformly sampled between 0.0001 and 0.1. Lastly, the batch size is a choice between 2, 4, 8, and 16.

At each trial, `ray[tune]` will randomly sample a combination of parameters from these search spaces. It will then train a number of models in parallel and find the best performing one among these. `ray[tune]` uses the `ASHAScheduler` which will terminate bad performing trials early.

14.5.5 Configuring the Search Space With spotPython

14.5.5.1 The hyper_dict Hyperparameters for the Selected Algorithm

spotPython uses JSON files for the specification of the hyperparameters. Users can specify their individual JSON files, or they can use the JSON files provided by spotPython. The JSON file for the `core_model` is called `torch_hyper_dict.json`.

In contrast to `ray[tune]`, `spotPython` can handle numerical, boolean, and categorical hyperparameters. They can be specified in the JSON file in a similar way as the numerical hyperparameters as shown below. Each entry in the JSON file represents one hyperparameter with the following structure: `type`, `default`, `transform`, `lower`, and `upper`.

```
"factor_hyperparameter": {
  "levels": ["A", "B", "C"],
  "type": "factor",
  "default": "B",
  "transform": "None",
  "core_model_parameter_type": "str",
  "lower": 0,
  "upper": 2},
```

The corresponding entries for the `core_model` class are shown below.

```
fun_control['core_model_hyper_dict']

{'l1': {'type': 'int',
  'default': 5,
  'transform': 'transform_power_2_int',
  'lower': 2,
  'upper': 9},
'l2': {'type': 'int',
  'default': 5,
  'transform': 'transform_power_2_int',
  'lower': 2,
  'upper': 9},
'lr_mult': {'type': 'float',
  'default': 1.0,
  'transform': 'None',
  'lower': 0.1,
  'upper': 10.0},
'batch_size': {'type': 'int',
```

```

'default': 4,
'transform': 'transform_power_2_int',
'lower': 1,
'upper': 4},
'epochs': {'type': 'int',
'default': 3,
'transform': 'transform_power_2_int',
'lower': 3,
'upper': 4},
'k_folds': {'type': 'int',
'default': 1,
'transform': 'None',
'lower': 1,
'upper': 1},
'patience': {'type': 'int',
'default': 5,
'transform': 'None',
'lower': 2,
'upper': 10},
'optimizer': {'levels': ['Adadelata',
'Adagrad',
'Adam',
'AdamW',
'SparseAdam',
'Adamax',
'ASGD',
'NAdam',
'RAdam',
'RMSprop',
'Rprop',
'SGD'],
'type': 'factor',
'default': 'SGD',
'transform': 'None',
'class_name': 'torch.optim',
'core_model_parameter_type': 'str',
'lower': 0,
'upper': 12},
'sgd_momentum': {'type': 'float',
'default': 0.0,
'transform': 'None',
'lower': 0.0,
'upper': 1.0}}

```

14.6 Step 6: Modify `hyper_dict` Hyperparameters for the Selected Algorithm aka `core_model`

Ray tune (PyTorch 2023a) does not provide a way to change the specified hyperparameters without re-compilation. However, `spotPython` provides functions for modifying the hyperparameters, their bounds and factors as well as for activating and de-activating hyperparameters without re-compilation of the Python source code. These functions are described in the following.

14.6.0.1 Modify `hyper_dict` Hyperparameters for the Selected Algorithm aka `core_model`

After specifying the model, the corresponding hyperparameters, their types and bounds are loaded from the JSON file `torch_hyper_dict.json`. After loading, the user can modify the hyperparameters, e.g., the bounds. `spotPython` provides a simple rule for de-activating hyperparameters: If the lower and the upper bound are set to identical values, the hyperparameter is de-activated. This is useful for the hyperparameter tuning, because it allows to specify a hyperparameter in the JSON file, but to de-activate it in the `fun_control` dictionary. This is done in the next step.

14.6.0.2 Modify Hyperparameters of Type `numeric` and `integer` (`boolean`)

Since the hyperparameter `k_folds` is not used in the PyTorch tutorial, it is de-activated here by setting the lower and upper bound to the same value. Note, `k_folds` is of type “integer”.

```
from spotPython.hyperparameters.values import modify_hyper_parameter_bounds
fun_control = modify_hyper_parameter_bounds(fun_control,
    "batch_size", bounds=[1, 5])
fun_control = modify_hyper_parameter_bounds(fun_control,
    "k_folds", bounds=[0, 0])
fun_control = modify_hyper_parameter_bounds(fun_control,
    "patience", bounds=[3, 3])
```

14.6.0.3 Modify Hyperparameter of Type `factor`

In a similar manner as for the numerical hyperparameters, the categorical hyperparameters can be modified. New configurations can be chosen by adding or deleting levels. For example, the hyperparameter `optimizer` can be re-configured as follows:

In the following setting, two optimizers ("SGD" and "Adam") will be compared during the spotPython hyperparameter tuning. The hyperparameter optimizer is active.

```
from spotPython.hyperparameters.values import modify_hyper_parameter_levels
fun_control = modify_hyper_parameter_levels(fun_control,
                                           "optimizer", ["SGD", "Adam"])
```

The hyperparameter optimizer can be de-activated by choosing only one value (level), here: "SGD".

```
fun_control = modify_hyper_parameter_levels(fun_control, "optimizer", ["SGD"])
```

As discussed in Section 14.6.1, there are some issues with the LBFGS optimizer. Therefore, the usage of the LBFGS optimizer is not deactivated in spotPython by default. However, the LBFGS optimizer can be activated by adding it to the list of optimizers. Rprop was removed, because it does perform very poorly (as some pre-tests have shown). However, it can also be activated by adding it to the list of optimizers. Since SparseAdam does not support dense gradients, Adam was used instead. Therefore, there are 10 default optimizers:

```
fun_control = modify_hyper_parameter_levels(fun_control, "optimizer",
                                           ["Adadelta", "Adagrad", "Adam", "AdamW", "Adamax", "ASGD",
                                           "NAdam", "RAdam", "RMSprop", "SGD"])
```

14.6.1 Optimizers

Table 14.1 shows some of the optimizers available in PyTorch:

a denotes (0.9,0.999), b (0.5,1.2), and c (1e-6, 50), respectively. R denotes required, but unspecified. "m" denotes momentum, "w_d" weight_decay, "d" dampening, "n" nesterov, "r" rho, "l_s" learning rate for scaling delta, "l_d" lr_decay, "b" betas, "l" lambda, "a" alpha, "m_d" for momentum_decay, "e" etas, and "s_s" for step_sizes.

Table 14.1: Optimizers available in PyTorch (selection). The default values are shown in the table.

Optimizer	lr	m	w_d	d	n	r	l_s	l_d	b	l	a	m_d	e	s_s
Adadelta	-	-	0.	-	-	0.9	1.	-	-	-	-	-	-	-
Adagrad	1e-2	-	0.	-	-	-	-	0.	-	-	-	-	-	-
Adam	1e-3	-	0.	-	-	-	-	-	a	-	-	-	-	-
AdamW	1e-3	-	1e-2	-	-	-	-	-	a	-	-	-	-	-
SparseAdam	1e-3	-	-	-	-	-	-	-	a	-	-	-	-	-
Adamax	2e-3	-	0.	-	-	-	-	-	a	-	-	-	-	-

Optimizer	lr	m	w_d	d	n	r	l_s	l_d	b	l	a	m_d	e	s_s
ASGD	1e-2	.9	0.	-	F	-	-	-	-	1e-4	.75	-	-	-
LBFGS	1.	-	-	-	-	-	-	-	-	-	-	-	-	-
NAdam	2e-3	-	0.	-	-	-	-	-	<i>a</i>	-	-	0	-	-
RAdam	1e-3	-	0.	-	-	-	-	-	<i>a</i>	-	-	-	-	-
RMSprop	1e-2	0.	0.	-	-	-	-	-	<i>a</i>	-	-	-	-	-
Rprop	1e-2	-	-	-	-	-	-	-	-	-	<i>b</i>	<i>c</i>	-	-
SGD	<i>R</i>	0.	0.	0.	F	-	-	-	-	-	-	-	-	-

`spotPython` implements an `optimization` handler that maps the optimizer names to the corresponding PyTorch optimizers.

i A note on LBFGS

We recommend deactivating PyTorch’s LBFGS optimizer, because it does not perform very well. The PyTorch documentation, see <https://pytorch.org/docs/stable/generated/torch.optim.LBFGS.html#torch.optim.LBFGS>, states:

This is a very memory intensive optimizer (it requires additional `param_bytes * (history_size + 1)` bytes). If it doesn’t fit in memory try reducing the history size, or use a different algorithm.

Furthermore, the LBFGS optimizer is not compatible with the PyTorch tutorial. The reason is that the LBFGS optimizer requires the `closure` function, which is not implemented in the PyTorch tutorial. Therefore, the LBFGS optimizer is recommended here. Since there are ten optimizers in the portfolio, it is not recommended tuning the hyperparameters that effect one single optimizer only.

i A note on the learning rate

`spotPython` provides a multiplier for the default learning rates, `lr_mult`, because optimizers use different learning rates. Using a multiplier for the learning rates might enable a simultaneous tuning of the learning rates for all optimizers. However, this is not recommended, because the learning rates are not comparable across optimizers. Therefore, we recommend fixing the learning rate for all optimizers if multiple optimizers are used. This can be done by setting the lower and upper bounds of the learning rate multiplier to the same value as shown below.

Thus, the learning rate, which affects the SGD optimizer, will be set to a fixed value. We choose the default value of `1e-3` for the learning rate, because it is used in other PyTorch examples (it is also the default value used by `spotPython` as defined in the `optimizer_handler()` method). We recommend tuning the learning rate later, when a

reduced set of optimizers is fixed. Here, we will demonstrate how to select in a screening phase the optimizers that should be used for the hyperparameter tuning.

For the same reason, we will fix the `sgd_momentum` to 0.9.

```
fun_control = modify_hyper_parameter_bounds(fun_control,
                                           "lr_mult", bounds=[1.0, 1.0])
fun_control = modify_hyper_parameter_bounds(fun_control,
                                           "sgd_momentum", bounds=[0.9, 0.9])
```

14.7 Step 7: Selection of the Objective (Loss) Function

14.7.1 Evaluation: Data Splitting

The evaluation procedure requires the specification of the way how the data is split into a train and a test set and the loss function (and a metric). As a default, `spotPython` provides a standard hold-out data split and cross validation.

14.7.2 Hold-out Data Split

If a hold-out data split is used, the data will be partitioned into a training, a validation, and a test data set. The split depends on the setting of the `eval` parameter. If `eval` is set to `train_hold_out`, one data set, usually the original training data set, is split into a new training and a validation data set. The training data set is used for training the model. The validation data set is used for the evaluation of the hyperparameter configuration and early stopping to prevent overfitting. In this case, the original test data set is not used.

i Note

`spotPython` returns the hyperparameters of the machine learning and deep learning models, e.g., number of layers, learning rate, or optimizer, but not the model weights. Therefore, after the SPOT run is finished, the corresponding model with the optimized architecture has to be trained again with the best hyperparameter configuration. The training is performed on the training data set. The test data set is used for the final evaluation of the model.

Summarizing, the following splits are performed in the hold-out setting:

1. Run `spotPython` with `eval` set to `train_hold_out` to determine the best hyperparameter configuration.
2. Train the model with the best hyperparameter configuration (“architecture”) on

```
the training data set: train_tuned(model_spot, train, "model_spot.pt").
3. Test the model on the test data: test_tuned(model_spot, test,
"model_spot.pt")
```

These steps will be exemplified in the following sections.

In addition to this hold-out setting, `spotPython` provides another hold-out setting, where an explicit test data is specified by the user that will be used as the validation set. To choose this option, the `eval` parameter is set to `test_hold_out`. In this case, the training data set is used for the model training. Then, the explicitly defined test data set is used for the evaluation of the hyperparameter configuration (the validation).

14.7.3 Cross-Validation

The cross validation setting is used by setting the `eval` parameter to `train_cv` or `test_cv`. In both cases, the data set is split into k folds. The model is trained on $k - 1$ folds and evaluated on the remaining fold. This is repeated k times, so that each fold is used exactly once for evaluation. The final evaluation is performed on the test data set. The cross validation setting is useful for small data sets, because it allows to use all data for training and evaluation. However, it is computationally expensive, because the model has to be trained k times.

i Note

Combinations of the above settings are possible, e.g., cross validation can be used for training and hold-out for evaluation or *vice versa*. Also, cross validation can be used for training and testing. Because cross validation is not used in the `PyTorch` tutorial (`PyTorch 2023a`), it is not considered further here.

14.7.4 Overview of the Evaluation Settings

14.7.4.1 Settings for the Hyperparameter Tuning

An overview of the training evaluations is shown in Table 14.2. `"train_cv"` and `"test_cv"` use `sklearn.model_selection.KFold()` internally. More details on the data splitting are provided in Section 21.14 (in the Appendix).

Table 14.2: Overview of the evaluation settings.

eval	train	test	function	comment
"train_hold_out" ✓			train_one_epoch(), validate_one_epoch() for early stopping	splits the train data set internally
"test_hold_out" ✓	✓	✓	train_one_epoch(), validate_one_epoch() for early stopping	use the test data set for validate_one_epoch()
"train_cv" ✓	✓		evaluate_cv(net, train)	CV using the train data set
"test_cv"		✓	evaluate_cv(net, test)	CV using the test data set . Identical to "train_cv", uses only test data.

14.7.4.2 Settings for the Final Evaluation of the Tuned Architecture

14.7.4.2.1 Training of the Tuned Architecture

`train_tuned(model, train)`: train the model with the best hyperparameter configuration (or simply the default) on the training data set. It splits the `traindata` into new `train` and `validation` sets using `create_train_val_data_loaders()`, which calls `torch.utils.data.random_split()` internally. Currently, 60% of the data is used for training and 40% for validation. The `train` data is used for training the model with `train_hold_out()`. The `validation` data is used for early stopping using `validate_fold_or_hold_out()` on the validation data set.

14.7.4.2.2 Testing of the Tuned Architecture

`test_tuned(model, test)`: test the model on the test data set. No data splitting is performed. The (trained) model is evaluated using the `validate_fold_or_hold_out()` function. Note: During training, `"shuffle"` is set to `True`, whereas during testing, `"shuffle"` is set to `False`.

Section [21.14.1.4](#) describes the final evaluation of the tuned architecture.

```
fun_control.update({
    "eval": "train_hold_out",
    "path": "torch_model.pt",
    "shuffle": True})
```

14.7.5 Evaluation: Loss Functions and Metrics

The key "loss_function" specifies the loss function which is used during the optimization. There are several different loss functions under PyTorch's `nn` package. For example, a simple loss is `MSELoss`, which computes the mean-squared error between the output and the target. In this tutorial we will use `CrossEntropyLoss`, because it is also used in the PyTorch tutorial.

```
from torch.nn import CrossEntropyLoss
loss_function = CrossEntropyLoss()
fun_control.update({"loss_function": loss_function})
```

In addition to the loss functions, `spotPython` provides access to a large number of metrics.

- The key "metric_sklearn" is used for metrics that follow the `scikit-learn` conventions.
- The key "river_metric" is used for the river based evaluation (Montiel et al. 2021) via `eval_oml_iter_progressive`, and
- the key "metric_torch" is used for the metrics from `TorchMetrics`.

`TorchMetrics` is a collection of more than 90 PyTorch metrics, see <https://torchmetrics.readthedocs.io/en/latest/>. Because the PyTorch tutorial uses the accuracy as metric, we use the same metric here. Currently, accuracy is computed in the tutorial's example code. We will use `TorchMetrics` instead, because it offers more flexibility, e.g., it can be used for regression and classification. Furthermore, `TorchMetrics` offers the following advantages:

- * A standardized interface to increase reproducibility
- * Reduces Boilerplate
- * Distributed-training compatible
- * Rigorously tested
- * Automatic accumulation over batches
- * Automatic synchronization between multiple devices

Therefore, we set

```
import torchmetrics
metric_torch = torchmetrics.Accuracy(task="multiclass", num_classes=10).to(fun_control["device"])
fun_control.update({"metric_torch": metric_torch})
```

14.8 Step 8: Calling the SPOT Function

14.8.1 Preparing the SPOT Call

The following code passes the information about the parameter ranges and bounds to `spot`.

```
from spotPython.hyperparameters.values import (
    get_var_type,
    get_var_name,
    get_bound_values
)

var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
fun_control.update({"var_type": var_type,
                   "var_name": var_name})

lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")
```

Now, the dictionary `fun_control` contains all information needed for the hyperparameter tuning. Before the hyperparameter tuning is started, it is recommended to take a look at the experimental design. The method `gen_design_table` generates a design table as follows:

```
from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))
```

name	type	default	lower	upper	transform
l1	int	5	2	9	transform_power_2_int
l2	int	5	2	9	transform_power_2_int
lr_mult	float	1.0	1	1	None
batch_size	int	4	1	5	transform_power_2_int
epochs	int	3	3	4	transform_power_2_int
k_folds	int	1	0	0	None
patience	int	5	3	3	None
optimizer	factor	SGD	0	9	None
sgd_momentum	float	0.0	0.9	0.9	None

This allows to check if all information is available and if the information is correct. `?@tbl-design` shows the experimental design for the hyperparameter tuning. The table shows the

hyperparameters, their types, default values, lower and upper bounds, and the transformation function. The transformation function is used to transform the hyperparameter values from the unit hypercube to the original domain. The transformation function is applied to the hyperparameter values before the evaluation of the objective function. Hyperparameter transformations are shown in the column “transform”, e.g., the 11 default is 5, which results in the value $2^5 = 32$ for the network, because the transformation `transform_power_2_int` was selected in the JSON file. The default value of the `batch_size` is set to 4, which results in a batch size of $2^4 = 16$.

14.8.2 The Objective Function `fun_torch`

The objective function `fun_torch` is selected next. It implements an interface from PyTorch’s training, validation, and testing methods to `spotPython`.

```
from spotPython.fun.hypertorch import HyperTorch
fun = HyperTorch().fun_torch
```

14.8.3 Using Default Hyperparameters or Results from Previous Runs

We add the default setting to the initial design:

```
from spotPython.hyperparameters.values import get_default_hyperparameters_as_array
hyper_dict=TorchHyperDict().load()
X_start = get_default_hyperparameters_as_array(fun_control, hyper_dict)
```

14.8.4 Starting the Hyperparameter Tuning

The `spotPython` hyperparameter tuning is started by calling the `Spot` function. Here, we will run the tuner for approximately 30 minutes (`max_time`). Note: the initial design is always evaluated in the `spotPython` run. As a consequence, the run may take longer than specified by `max_time`, because the evaluation time of initial design (here: `init_size`, 10 points) is performed independently of `max_time`. During the run, results from the training is shown. These results can be visualized with Tensorboard as will be shown in Section 14.9.

```
from spotPython.spot import spot
from math import inf
import numpy as np
spot_tuner = spot.Spot(fun=fun,
                      lower = lower,
```


MulticlassAccuracy: 0.4885500073432922 | Loss: 1.3688855980873107 | Acc: 0.4885500000000000.
Epoch: 5 |

MulticlassAccuracy: 0.5167999863624573 | Loss: 1.3167984383583069 | Acc: 0.5168000000000000.
Epoch: 6 |

MulticlassAccuracy: 0.5378999710083008 | Loss: 1.2682793026924133 | Acc: 0.5379000000000000.
Epoch: 7 |

MulticlassAccuracy: 0.5465000271797180 | Loss: 1.2487945146560668 | Acc: 0.5465000000000000.
Epoch: 8 |

MulticlassAccuracy: 0.5631999969482422 | Loss: 1.2205491739273071 | Acc: 0.5632000000000000.
Epoch: 9 |

MulticlassAccuracy: 0.5613999962806702 | Loss: 1.2304849659919739 | Acc: 0.5614000000000000.
Epoch: 10 |

MulticlassAccuracy: 0.5752500295639038 | Loss: 1.1978971423149110 | Acc: 0.5752500000000000.
Epoch: 11 |

MulticlassAccuracy: 0.5883499979972839 | Loss: 1.1682895747184754 | Acc: 0.5883500000000000.
Epoch: 12 |

MulticlassAccuracy: 0.5803499817848206 | Loss: 1.2077535195350646 | Acc: 0.5803500000000000.
Epoch: 13 |

MulticlassAccuracy: 0.5759999752044678 | Loss: 1.2306316658020020 | Acc: 0.5760000000000000.
Epoch: 14 |

MulticlassAccuracy: 0.5906000137329102 | Loss: 1.1822338700294495 | Acc: 0.5906000000000000.
Early stopping at epoch 13
Returned to Spot: Validation loss: 1.1822338700294495

config: {'l1': 16, 'l2': 16, 'lr_mult': 1.0, 'batch_size': 8, 'epochs': 8, 'k_folds': 0, 'pa
Epoch: 1 |

MulticlassAccuracy: 0.4542500078678131 | Loss: 1.5091321352601050 | Acc: 0.4542500000000000.
Epoch: 2 |

MulticlassAccuracy: 0.4963499903678894 | Loss: 1.4083931375503540 | Acc: 0.4963500000000000.
Epoch: 3 |

MulticlassAccuracy: 0.5075500011444092 | Loss: 1.3952710754871369 | Acc: 0.5075499999999999.
Epoch: 4 |

MulticlassAccuracy: 0.5388500094413757 | Loss: 1.2818053160846234 | Acc: 0.5388500000000001.
Epoch: 5 |

MulticlassAccuracy: 0.5512499809265137 | Loss: 1.2799427218079567 | Acc: 0.5512500000000000.
Epoch: 6 |

MulticlassAccuracy: 0.5591999888420105 | Loss: 1.2349592106938363 | Acc: 0.5592000000000000.
Epoch: 7 |

MulticlassAccuracy: 0.5665000081062317 | Loss: 1.2287030029177666 | Acc: 0.5665000000000000.
Epoch: 8 |

MulticlassAccuracy: 0.5521000027656555 | Loss: 1.2753997170805931 | Acc: 0.5521000000000000.
Returned to Spot: Validation loss: 1.275399717080593

config: {'l1': 256, 'l2': 128, 'lr_mult': 1.0, 'batch_size': 2, 'epochs': 16, 'k_folds': 0,
Epoch: 1 |

MulticlassAccuracy: 0.1001000031828880 | Loss: 2.3082210160255432 | Acc: 0.1001000000000000.
Epoch: 2 |

MulticlassAccuracy: 0.1024499982595444 | Loss: 2.3105591154098510 | Acc: 0.1024500000000000.
Epoch: 3 |

MulticlassAccuracy: 0.0992500036954880 | Loss: 2.3095418009757998 | Acc: 0.0992500000000000.
Epoch: 4 |

MulticlassAccuracy: 0.1027500033378601 | Loss: 2.3071149304151537 | Acc: 0.1027500000000000.
Epoch: 5 |

MulticlassAccuracy: 0.0991000011563301 | Loss: 2.3067904354810715 | Acc: 0.0991000000000000.
Epoch: 6 |

MulticlassAccuracy: 0.0999500006437302 | Loss: 2.3113588804483416 | Acc: 0.0999500000000000.
Epoch: 7 |

MulticlassAccuracy: 0.0991000011563301 | Loss: 2.3086068964004518 | Acc: 0.0991000000000000.
Epoch: 8 |

MulticlassAccuracy: 0.0988999977707863 | Loss: 2.3113879865884779 | Acc: 0.0989000000000000.
Early stopping at epoch 7
Returned to Spot: Validation loss: 2.311387986588478

config: {'l1': 8, 'l2': 32, 'lr_mult': 1.0, 'batch_size': 4, 'epochs': 8, 'k_folds': 0, 'pat.
Epoch: 1 |

MulticlassAccuracy: 0.4166499972343445 | Loss: 1.5726691699266433 | Acc: 0.4166500000000000.
Epoch: 2 |

MulticlassAccuracy: 0.4670000076293945 | Loss: 1.4920029521405698 | Acc: 0.4670000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.5202500224113464 | Loss: 1.3516685632601380 | Acc: 0.5202500000000000.
Epoch: 4 |

MulticlassAccuracy: 0.5358499884605408 | Loss: 1.3217709792330861 | Acc: 0.5358500000000000.
Epoch: 5 |

MulticlassAccuracy: 0.5334500074386597 | Loss: 1.3406620404675602 | Acc: 0.5334500000000000.
Epoch: 6 |

MulticlassAccuracy: 0.5503000020980835 | Loss: 1.2966141935270279 | Acc: 0.5503000000000000.
Epoch: 7 |

MulticlassAccuracy: 0.5745499730110168 | Loss: 1.2292844279039652 | Acc: 0.5745500000000000.
Epoch: 8 |

MulticlassAccuracy: 0.5806000232696533 | Loss: 1.2071203859100119 | Acc: 0.5806000000000000.
Returned to Spot: Validation loss: 1.207120385910012

config: {'l1': 64, 'l2': 512, 'lr_mult': 1.0, 'batch_size': 16, 'epochs': 16, 'k_folds': 0,
Epoch: 1 |

MulticlassAccuracy: 0.4485999941825867 | Loss: 1.4967560271263123 | Acc: 0.4486000000000000.
Epoch: 2 |

MulticlassAccuracy: 0.4823000133037567 | Loss: 1.4261294353008269 | Acc: 0.4823000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.5031999945640564 | Loss: 1.3617523387908936 | Acc: 0.5032000000000000.
Epoch: 4 |

MulticlassAccuracy: 0.5171499848365784 | Loss: 1.3235158298015595 | Acc: 0.5171500000000000.
Epoch: 5 |

MulticlassAccuracy: 0.5337499976158142 | Loss: 1.2952653534412384 | Acc: 0.5337499999999999.
Epoch: 6 |

MulticlassAccuracy: 0.5358999967575073 | Loss: 1.2782614757061004 | Acc: 0.5359000000000000.
Epoch: 7 |

MulticlassAccuracy: 0.5461500287055969 | Loss: 1.2658197217941285 | Acc: 0.5461500000000000.
Epoch: 8 |

MulticlassAccuracy: 0.5529000163078308 | Loss: 1.2501332284450530 | Acc: 0.5528999999999999.
Epoch: 9 |

MulticlassAccuracy: 0.5569999814033508 | Loss: 1.2343813209533692 | Acc: 0.5570000000000001.
Epoch: 10 |

MulticlassAccuracy: 0.5604500174522400 | Loss: 1.2296385282993316 | Acc: 0.5604500000000000.
Epoch: 11 |

MulticlassAccuracy: 0.5650500059127808 | Loss: 1.2153698269367219 | Acc: 0.5650500000000001.
Epoch: 12 |

MulticlassAccuracy: 0.5697500109672546 | Loss: 1.2049174281597137 | Acc: 0.5697500000000000.
Epoch: 13 |

MulticlassAccuracy: 0.5668500065803528 | Loss: 1.2107061998367310 | Acc: 0.5668500000000000.
Epoch: 14 |

MulticlassAccuracy: 0.5727999806404114 | Loss: 1.1986185997486114 | Acc: 0.5728000000000000.
Epoch: 15 |

MulticlassAccuracy: 0.5730000138282776 | Loss: 1.1946961738824844 | Acc: 0.5730000000000000.
Epoch: 16 |

MulticlassAccuracy: 0.5774499773979187 | Loss: 1.1870610105276107 | Acc: 0.5774500000000000.
Returned to Spot: Validation loss: 1.1870610105276107

config: {'l1': 128, 'l2': 16, 'lr_mult': 1.0, 'batch_size': 32, 'epochs': 16, 'k_folds': 0,
Epoch: 1 |

MulticlassAccuracy: 0.4137000143527985 | Loss: 1.5898399522781372 | Acc: 0.4137000000000000.
Epoch: 2 |

MulticlassAccuracy: 0.4796000123023987 | Loss: 1.4382932362556458 | Acc: 0.4796000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.5036000013351440 | Loss: 1.3805855596542358 | Acc: 0.5036000000000000.
Epoch: 4 |

MulticlassAccuracy: 0.5412499904632568 | Loss: 1.2722473503112792 | Acc: 0.5412500000000000.
Epoch: 5 |

MulticlassAccuracy: 0.5504000186920166 | Loss: 1.2480867130279540 | Acc: 0.5504000000000000.
Epoch: 6 |

MulticlassAccuracy: 0.5694000124931335 | Loss: 1.2153308080673217 | Acc: 0.5694000000000000.
Epoch: 7 |

MulticlassAccuracy: 0.5759999752044678 | Loss: 1.1881502468109131 | Acc: 0.5760000000000000.
Epoch: 8 |

MulticlassAccuracy: 0.5806999802589417 | Loss: 1.1821760972976685 | Acc: 0.5807000000000000.
Epoch: 9 |

MulticlassAccuracy: 0.5892000198364258 | Loss: 1.1729242265701294 | Acc: 0.5891999999999999.
Epoch: 10 |

MulticlassAccuracy: 0.5850499868392944 | Loss: 1.1787366871833802 | Acc: 0.5850500000000000.
Epoch: 11 |

MulticlassAccuracy: 0.5993000268936157 | Loss: 1.1413288003444673 | Acc: 0.5993000000000001.
Epoch: 12 |

MulticlassAccuracy: 0.6014500260353088 | Loss: 1.1668504621505738 | Acc: 0.6014500000000000.
Epoch: 13 |

MulticlassAccuracy: 0.6021999716758728 | Loss: 1.1557737874984741 | Acc: 0.6022000000000000.
Epoch: 14 |

MulticlassAccuracy: 0.5997999906539917 | Loss: 1.1702202931880952 | Acc: 0.5998000000000000.
Early stopping at epoch 13
Returned to Spot: Validation loss: 1.1702202931880952

spotPython tuning: 1.1702202931880952 [###-----] 25.93%

config: {'l1': 128, 'l2': 512, 'lr_mult': 1.0, 'batch_size': 32, 'epochs': 16, 'k_folds': 0,
Epoch: 1 |

MulticlassAccuracy: 0.4648999869823456 | Loss: 1.4701293358802796 | Acc: 0.4649000000000000.
Epoch: 2 |

MulticlassAccuracy: 0.5274500250816345 | Loss: 1.3237315242767334 | Acc: 0.5274500000000000.
Epoch: 3 |

MulticlassAccuracy: 0.5690500140190125 | Loss: 1.2135559646606446 | Acc: 0.5690499999999999.
Epoch: 4 |

MulticlassAccuracy: 0.5766500234603882 | Loss: 1.1901183177947998 | Acc: 0.5766500000000000.
Epoch: 5 |

MulticlassAccuracy: 0.5976499915122986 | Loss: 1.1567914697647095 | Acc: 0.5976500000000000.
Epoch: 6 |

MulticlassAccuracy: 0.5945000052452087 | Loss: 1.1695243903160095 | Acc: 0.5945000000000000.
Epoch: 7 |

MulticlassAccuracy: 0.4708999991416931 | Loss: 1.4370097077369690 | Acc: 0.4709000000000000.
Epoch: 6 |

MulticlassAccuracy: 0.4718999862670898 | Loss: 1.4181426862716675 | Acc: 0.4719000000000000.
Epoch: 7 |

MulticlassAccuracy: 0.4939500093460083 | Loss: 1.3904836480140685 | Acc: 0.4939500000000000.
Epoch: 8 |

MulticlassAccuracy: 0.4995500147342682 | Loss: 1.3769942201614380 | Acc: 0.4995500000000000.
Epoch: 9 |

MulticlassAccuracy: 0.5084499716758728 | Loss: 1.3722874144554138 | Acc: 0.5084500000000000.
Epoch: 10 |

MulticlassAccuracy: 0.5225999951362610 | Loss: 1.3297592091560364 | Acc: 0.5226000000000000.
Epoch: 11 |

MulticlassAccuracy: 0.5249000191688538 | Loss: 1.3007804115295409 | Acc: 0.5249000000000000.
Epoch: 12 |

MulticlassAccuracy: 0.5105000138282776 | Loss: 1.3568648647308350 | Acc: 0.5105000000000000.
Epoch: 13 |

MulticlassAccuracy: 0.5199000239372253 | Loss: 1.3266702429771424 | Acc: 0.5199000000000000.
Epoch: 14 |

MulticlassAccuracy: 0.5394499897956848 | Loss: 1.2927523935317993 | Acc: 0.5394500000000000.
Epoch: 15 |

MulticlassAccuracy: 0.5360500216484070 | Loss: 1.3086393711090087 | Acc: 0.5360500000000000.
Epoch: 16 |

MulticlassAccuracy: 0.5379999876022339 | Loss: 1.3059302357673646 | Acc: 0.5380000000000000.
Returned to Spot: Validation loss: 1.3059302357673646
spotPython tuning: 1.1702202931880952 [#####--] 80.14%

config: {'l1': 16, 'l2': 32, 'lr_mult': 1.0, 'batch_size': 8, 'epochs': 16, 'k_folds': 0, 'p
Epoch: 1 |

MulticlassAccuracy: 0.0975499972701073 | Loss: 2.3062089468955995 | Acc: 0.0975500000000000.
Epoch: 2 |

MulticlassAccuracy: 0.0998999997973442 | Loss: 2.3083316393852233 | Acc: 0.0999000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.1017000004649162 | Loss: 2.3034695069313051 | Acc: 0.1017000000000000.
Epoch: 4 |

MulticlassAccuracy: 0.1020999997854233 | Loss: 2.3041320311546327 | Acc: 0.1021000000000000.
Epoch: 5 |

MulticlassAccuracy: 0.1000500023365021 | Loss: 2.3084358304977415 | Acc: 0.1000500000000000.
Epoch: 6 |

MulticlassAccuracy: 0.1020999997854233 | Loss: 2.3038669599533081 | Acc: 0.1021000000000000.
Early stopping at epoch 5
Returned to Spot: Validation loss: 2.303866959953308

spotPython tuning: 1.1702202931880952 [#####-] 89.35%

config: {'l1': 16, 'l2': 512, 'lr_mult': 1.0, 'batch_size': 32, 'epochs': 16, 'k_folds': 0,
Epoch: 1 |

MulticlassAccuracy: 0.4500499963760376 | Loss: 1.5282039699554444 | Acc: 0.4500500000000000.
Epoch: 2 |

MulticlassAccuracy: 0.4803999960422516 | Loss: 1.4396946838378906 | Acc: 0.4804000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.4852499961853027 | Loss: 1.4204688949584960 | Acc: 0.4852500000000000.
Epoch: 4 |

MulticlassAccuracy: 0.4954499900341034 | Loss: 1.3889629602432252 | Acc: 0.4954500000000000.
Epoch: 5 |

MulticlassAccuracy: 0.5120000243186951 | Loss: 1.3607494637489319 | Acc: 0.5120000000000000.
Epoch: 6 |

MulticlassAccuracy: 0.5115000009536743 | Loss: 1.3528809277534486 | Acc: 0.5115000000000000.
Epoch: 7 |

MulticlassAccuracy: 0.5188000202178955 | Loss: 1.3332977566719055 | Acc: 0.5188000000000000.
Epoch: 8 |

MulticlassAccuracy: 0.5234000086784363 | Loss: 1.3171866333007813 | Acc: 0.5234000000000000.
Epoch: 9 |

MulticlassAccuracy: 0.5331500172615051 | Loss: 1.3003805235862731 | Acc: 0.5331500000000000.
Epoch: 10 |

MulticlassAccuracy: 0.5362499952316284 | Loss: 1.2932402163505554 | Acc: 0.5362500000000000.
Epoch: 11 |

MulticlassAccuracy: 0.5390999913215637 | Loss: 1.2862701608657836 | Acc: 0.5391000000000000.
Epoch: 12 |

MulticlassAccuracy: 0.5392000079154968 | Loss: 1.2832893109321595 | Acc: 0.5392000000000000.
Epoch: 13 |

MulticlassAccuracy: 0.5435000061988831 | Loss: 1.2712353020668030 | Acc: 0.5435000000000000.
Epoch: 14 |

MulticlassAccuracy: 0.5435500144958496 | Loss: 1.2644357828140258 | Acc: 0.5435500000000000.
Epoch: 15 |

MulticlassAccuracy: 0.5423499941825867 | Loss: 1.2824453946113585 | Acc: 0.5423500000000000.
Epoch: 16 |

MulticlassAccuracy: 0.5516499876976013 | Loss: 1.2494336812019349 | Acc: 0.5516500000000000.
Returned to Spot: Validation loss: 1.2494336812019349
spotPython tuning: 1.1702202931880952 [#####] 100.00% Done...

<spotPython.spot.spot.Spot at 0x28d6ab820>

14.9 Step 9: Tensorboard

The textual output shown in the console (or code cell) can be visualized with Tensorboard.

14.9.1 Tensorboard: Start Tensorboard

Start TensorBoard through the command line to visualize data you logged. Specify the root log directory as used in `fun_control = fun_control_init(task="regression", tensorboard_path="runs/24_spot_torch_regression")` as the `tensorboard_path`. The argument `logdir` points to directory where TensorBoard will look to find event files that it can display. TensorBoard will recursively walk the directory structure rooted at `logdir`, looking for `.tfevents` files.

```
tensorboard --logdir=runs
```

Go to the URL it provides or to <http://localhost:6006/>. The following figures show some screenshots of Tensorboard.




Figure 14.1: Tensorboard

Trial ID	Show Metrics	l1	l2	batch_size	epochs	patience	optimizer	fun_torch: loss
1686135261.24...	<input type="checkbox"/>	64.000	512.00	16.000	16.000	3.0000	Adagrad	1.1765
1686135486.0...	<input type="checkbox"/>	64.000	256.00	16.000	16.000	3.0000	Adagrad	1.1963
1686134673.15...	<input type="checkbox"/>	128.00	8.0000	32.000	16.000	3.0000	AdamW	1.2062
1686134773.50...	<input type="checkbox"/>	16.000	16.000	8.0000	8.0000	3.0000	NAdam	1.2880
1686135837.96...	<input type="checkbox"/>	64.000	256.00	16.000	16.000	3.0000	Adam	1.3155
1686135032.11...	<input type="checkbox"/>	8.0000	32.000	4.0000	8.0000	3.0000	Adamax	1.3435
1686135637.40...	<input type="checkbox"/>	64.000	128.00	2.0000	8.0000	3.0000	Adam	1.5804
1686135892.6...	<input type="checkbox"/>	16.000	32.000	8.0000	16.000	3.0000	RMSprop	2.1542
1686134917.07...	<input type="checkbox"/>	256.00	128.00	2.0000	16.000	3.0000	RMSprop	2.3099

Figure 14.2: Tensorboard

14.9.2 Saving the State of the Notebook

The state of the notebook can be saved and reloaded as follows:

```
import pickle
SAVE = False
LOAD = False

if SAVE:
    result_file_name = "res_" + experiment_name + ".pkl"
    with open(result_file_name, 'wb') as f:
        pickle.dump(spot_tuner, f)

if LOAD:
    result_file_name = "add_the_name_of_the_result_file_here.pkl"
    with open(result_file_name, 'rb') as f:
        spot_tuner = pickle.load(f)
```

14.10 Step 10: Results

After the hyperparameter tuning run is finished, the progress of the hyperparameter tuning can be visualized. The following code generates the progress plot from `?@fig-progress`.

```
spot_tuner.plot_progress(log_y=False,
    filename="./figures/" + experiment_name+"_progress.png")
```




Figure 14.3: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

`?@fig-progress` shows a typical behaviour that can be observed in many hyperparameter studies (Bartz et al. 2022): the largest improvement is obtained during the evaluation of the initial design. The surrogate model based optimization refines the results. `?@fig-progress` also illustrates one major difference between `ray[tune]` as used in PyTorch (2023a) and `spotPython`: the `ray[tune]` uses a random search and will generate results similar to the *black* dots, whereas `spotPython` uses a surrogate model based optimization and presents results represented by *red* dots in `?@fig-progress`. The surrogate model based optimization is considered to be more efficient than a random search, because the surrogate model guides the search towards promising regions in the hyperparameter space.

In addition to the improved (“optimized”) hyperparameter values, `spotPython` allows a statistical analysis, e.g., a sensitivity analysis, of the results. We can print the results of the hyperparameter tuning, see `?@tbl-results`. The table shows the hyperparameters, their types, default values, lower and upper bounds, and the transformation function. The column “tuned” shows the tuned values. The column “importance” shows the importance of the hyperparameters. The column “stars” shows the importance of the hyperparameters in stars. The importance is computed by the SPOT software.

```
from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control=fun_control, spot=spot_tuner))
```

name	type	default	lower	upper	tuned	transform
l1	int	5	2.0	9.0	7.0	transform_power_2_int
l2	int	5	2.0	9.0	4.0	transform_power_2_int
lr_mult	float	1.0	1.0	1.0	1.0	None
batch_size	int	4	1.0	5.0	5.0	transform_power_2_int

epochs	int	3		3.0		4.0		4.0		transform_power_2_int	
k_folds	int	1		0.0		0.0		0.0		None	
patience	int	5		3.0		3.0		3.0		None	
optimizer	factor	SGD		0.0		9.0		3.0		None	
sgd_momentum	float	0.0		0.9		0.9		0.9		None	

To visualize the most important hyperparameters, `spotPython` provides the function `plot_importance`. The following code generates the importance plot from `?@fig-importance`.

```
spot_tuner.plot_importance(threshold=0.025,
                           filename="./figures/" + experiment_name+"_importance.png")
```




Figure 14.4: Variable importance plot, threshold 0.025.

14.10.1 Get the Tuned Architecture (SPOT Results)

The architecture of the `spotPython` model can be obtained as follows. First, the numerical representation of the hyperparameters are obtained, i.e., the numpy array `X` is generated. This array is then used to generate the model `model_spot` by the function `get_one_core_model_from_X`. The model `model_spot` has the following architecture:

```

from spotPython.hyperparameters.values import get_one_core_model_from_X
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))
model_spot = get_one_core_model_from_X(X, fun_control)
model_spot

```

```

Net_CIFAR10(
  (conv1): Conv2d(3, 6, kernel_size=(5, 5), stride=(1, 1))
  (pool): MaxPool2d(kernel_size=2, stride=2, padding=0, dilation=1, ceil_mode=False)
  (conv2): Conv2d(6, 16, kernel_size=(5, 5), stride=(1, 1))
  (fc1): Linear(in_features=400, out_features=128, bias=True)
  (fc2): Linear(in_features=128, out_features=16, bias=True)
  (fc3): Linear(in_features=16, out_features=10, bias=True)
)

```

14.10.2 Get Default Hyperparameters

In a similar manner as in Section 14.10.1, the default hyperparameters can be obtained.

```

# fun_control was modified, we generate a new one with the original
# default hyperparameters
from spotPython.hyperparameters.values import get_one_core_model_from_X
fc = fun_control
fc.update({"core_model_hyper_dict":
  hyper_dict[fun_control["core_model"].__name__]})
model_default = get_one_core_model_from_X(X_start, fun_control=fc)
model_default

```

```

Net_CIFAR10(
  (conv1): Conv2d(3, 6, kernel_size=(5, 5), stride=(1, 1))
  (pool): MaxPool2d(kernel_size=2, stride=2, padding=0, dilation=1, ceil_mode=False)
  (conv2): Conv2d(6, 16, kernel_size=(5, 5), stride=(1, 1))
  (fc1): Linear(in_features=400, out_features=32, bias=True)
  (fc2): Linear(in_features=32, out_features=32, bias=True)
  (fc3): Linear(in_features=32, out_features=10, bias=True)
)

```

14.10.3 Evaluation of the Default Architecture

The method `train_tuned` takes a model architecture without trained weights and trains this model with the train data. The train data is split into train and validation data. The validation

data is used for early stopping. The trained model weights are saved as a dictionary. This evaluation is similar to the final evaluation in PyTorch (2023a).

```
from spotPython.torch.traintest import (
    train_tuned,
    test_tuned,
)
train_tuned(net=model_default, train_dataset=train, shuffle=True,
            loss_function=fun_control["loss_function"],
            metric=fun_control["metric_torch"],
            device = fun_control["device"], show_batch_interval=1_000_000,
            path=None,
            task=fun_control["task"],)

test_tuned(net=model_default, test_dataset=test,
            loss_function=fun_control["loss_function"],
            metric=fun_control["metric_torch"],
            shuffle=False,
            device = fun_control["device"],
            task=fun_control["task"],)
```

Epoch: 1 |

MulticlassAccuracy: 0.1036999970674515 | Loss: 2.3042341236114501 | Acc: 0.1037000000000000.
Epoch: 2 |

MulticlassAccuracy: 0.1106000021100044 | Loss: 2.3031450422286985 | Acc: 0.1106000000000000.
Epoch: 3 |

MulticlassAccuracy: 0.1171500012278557 | Loss: 2.3022329263687134 | Acc: 0.1171500000000000.
Epoch: 4 |

MulticlassAccuracy: 0.1272000074386597 | Loss: 2.3012789381027221 | Acc: 0.1272000000000000.
Epoch: 5 |

MulticlassAccuracy: 0.1410499960184097 | Loss: 2.2999784152984617 | Acc: 0.1410500000000000.
Epoch: 6 |

MulticlassAccuracy: 0.1469500064849854 | Loss: 2.2978194450378417 | Acc: 0.1469500000000000.
Epoch: 7 |

MulticlassAccuracy: 0.1459999978542328 | Loss: 2.2935358682632447 | Acc: 0.1460000000000000.
Epoch: 8 |

MulticlassAccuracy: 0.1375499963760376 | Loss: 2.2826718236923216 | Acc: 0.1375500000000000.
Returned to Spot: Validation loss: 2.2826718236923216

MulticlassAccuracy: 0.1374000012874603 | Loss: 2.2815687587738038 | Acc: 0.1374000000000000.
Final evaluation: Validation loss: 2.281568758773804
Final evaluation: Validation metric: 0.13740000128746033

(2.281568758773804, nan, tensor(0.1374))

14.10.4 Evaluation of the Tuned Architecture

The following code trains the model `model_spot`.

If `path` is set to a filename, e.g., `path = "model_spot_trained.pt"`, the weights of the trained model will be saved to this file.

If `path` is set to a filename, e.g., `path = "model_spot_trained.pt"`, the weights of the trained model will be loaded from this file.

```
train_tuned(net=model_spot, train_dataset=train,
            loss_function=fun_control["loss_function"],
            metric=fun_control["metric_torch"],
            shuffle=True,
            device = fun_control["device"],
            path=None,
            task=fun_control["task"],)
test_tuned(net=model_spot, test_dataset=test,
           shuffle=False,
           loss_function=fun_control["loss_function"],
           metric=fun_control["metric_torch"],
           device = fun_control["device"],
           task=fun_control["task"],)
```

Epoch: 1 |

MulticlassAccuracy: 0.4389500021934509 | Loss: 1.5250833265304566 | Acc: 0.4389500000000000.
Epoch: 2 |

MulticlassAccuracy: 0.4782499969005585 | Loss: 1.4415405203819276 | Acc: 0.4782500000000000.
Epoch: 3 |

MulticlassAccuracy: 0.5252000093460083 | Loss: 1.3205223499298095 | Acc: 0.5252000000000000.
Epoch: 4 |

MulticlassAccuracy: 0.5540999770164490 | Loss: 1.2503223680496216 | Acc: 0.5541000000000000.
Epoch: 5 |

MulticlassAccuracy: 0.5724499821662903 | Loss: 1.2086611101150513 | Acc: 0.5724500000000000.
Epoch: 6 |

MulticlassAccuracy: 0.5636500120162964 | Loss: 1.2337835314750671 | Acc: 0.5636500000000000.
Epoch: 7 |

MulticlassAccuracy: 0.5816000103950500 | Loss: 1.1978520113945008 | Acc: 0.5816000000000000.
Epoch: 8 |

MulticlassAccuracy: 0.5947999954223633 | Loss: 1.1535573720932006 | Acc: 0.5948000000000000.
Epoch: 9 |

MulticlassAccuracy: 0.5975499749183655 | Loss: 1.1575893239021302 | Acc: 0.5975500000000000.
Epoch: 10 |

MulticlassAccuracy: 0.5904999971389771 | Loss: 1.1804296854019165 | Acc: 0.5905000000000000.
Epoch: 11 |

MulticlassAccuracy: 0.5957499742507935 | Loss: 1.1608618884086608 | Acc: 0.5957500000000000.
Early stopping at epoch 10
Returned to Spot: Validation loss: 1.1608618884086608

MulticlassAccuracy: 0.6010000109672546 | Loss: 1.1575195035233665 | Acc: 0.6010000000000000.
Final evaluation: Validation loss: 1.1575195035233665
Final evaluation: Validation metric: 0.6010000109672546

(1.1575195035233665, nan, tensor(0.6010))

14.10.5 Detailed Hyperparameter Plots

The contour plots in this section visualize the interactions of the three most important hyperparameters. Since some of these hyperparameters take factorial or integer values, sometimes step-like fitness landscapes (or response surfaces) are generated. SPOT draws the interactions of the main hyperparameters by default. It is also possible to visualize all interactions.

```
filename = "./figures/" + experiment_name
spot_tuner.plot_important_hyperparameter_contour(filename=filename)
```

```
epochs: 100.0
optimizer: 0.803179786530439
```




Figure 14.5: Contour plots.

The figures (`?@fig-contour`) show the contour plots of the loss as a function of the hyperparameters. These plots are very helpful for benchmark studies and for understanding neural networks. `spotPython` provides additional tools for a visual inspection of the results and give valuable insights into the hyperparameter tuning process. This is especially useful for model explainability, transparency, and trustworthiness. In addition to the contour plots, `?@fig-parallel` shows the parallel plot of the hyperparameters.

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Parallel coordinates plots

Unable to display output for mime type(s): text/html

14.11 Summary and Outlook

This tutorial presents the hyperparameter tuning open source software `spotPython` for `PyTorch`. To show its basic features, a comparison with the “official” `PyTorch` hyperparameter tuning tutorial (PyTorch 2023a) is presented. Some of the advantages of `spotPython` are:

- Numerical and categorical hyperparameters.
- Powerful surrogate models.
- Flexible approach and easy to use.
- Simple JSON files for the specification of the hyperparameters.
- Extension of default and user specified network classes.
- Noise handling techniques.
- Interaction with `tensorboard`.

Currently, only rudimentary parallel and distributed neural network training is possible, but these capabilities will be extended in the future. The next version of `spotPython` will also include a more detailed documentation and more examples.

! Important

Important: This tutorial does not present a complete benchmarking study (Bartz-Beielstein et al. 2020). The results are only preliminary and highly dependent on the local configuration (hard- and software). Our goal is to provide a first impression of the performance of the hyperparameter tuning package `spotPython`. To demonstrate its capabilities, a quick comparison with `ray[tune]` was performed. `ray[tune]` was chosen, because it is presented as “an industry standard tool for distributed hyperparameter tuning.” The results should be interpreted with care.

14.12 Appendix

14.12.1 Sample Output From Ray Tune’s Run

The output from `ray[tune]` could look like this (PyTorch 2023b):

Number of trials: 10 (10 TERMINATED)

l1	l2	lr	batch_size	loss	accuracy	training_iteration
64	4	0.00011629	2	1.87273	0.244	2
32	64	0.000339763	8	1.23603	0.567	8
8	16	0.00276249	16	1.1815	0.5836	10
4	64	0.000648721	4	1.31131	0.5224	8
32	16	0.000340753	8	1.26454	0.5444	8
8	4	0.000699775	8	1.99594	0.1983	2
256	8	0.0839654	16	2.3119	0.0993	1
16	128	0.0758154	16	2.33575	0.1327	1
16	8	0.0763312	16	2.31129	0.1042	4
128	16	0.000124903	4	2.26917	0.1945	1

Best trial config: {'l1': 8, 'l2': 16, 'lr': 0.00276249, 'batch_size': 16, 'data_dir': '..'}

Best trial final validation loss: 1.181501

Best trial final validation accuracy: 0.5836

Best trial test set accuracy: 0.5806

15 HPT: sklearn RandomForestClassifier VBDP Data

This document refers to the following software versions:

- python: 3.10.10

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

spotPython can be installed via pip. Alternatively, the source code can be downloaded from gitHub: <https://github.com/sequential-parameter-optimization/spotPython>.

```
!pip install spotPython
```

- Uncomment the following lines if you want to for (re-)installation the latest version of spotPython from gitHub.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

15.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time and the initial design size.

```

MAX_TIME = 1
INIT_SIZE = 5
ORIGINAL = False

import os
import copy
import socket
from datetime import datetime
from dateutil.tz import tzlocal
start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '16-rf-sklearn' + "_" + HOSTNAME + "_" + str(MAX_TIME) + "min_" + str(INIT_SIZE)
experiment_name = experiment_name.replace(':', '-')
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')

```

16-rf-sklearn_bartz09_1min_5init_2023-06-19_04-02-41

```

import warnings
warnings.filterwarnings("ignore")

```

15.2 Step 2: Initialization of the Empty fun_control Dictionary

 Caution: Tensorboard does not work under Windows

- Since tensorboard does not work under Windows, we recommend setting the parameter `tensorboard_path` to `None` if you are working under Windows.

```

from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="classification",
    tensorboard_path="runs/16_spot_hpt_sklearn_classification")

```

15.3 Step 3: PyTorch Data Loading

15.3.1 Load Data: Classification VBDP

```
import pandas as pd
if ORIGINAL == True:
    train_df = pd.read_csv('./data/VBDP/trainnn.csv')
    test_df = pd.read_csv('./data/VBDP/testtt.csv')
else:
    train_df = pd.read_csv('./data/VBDP/train.csv')
    # remove the id column
    train_df = train_df.drop(columns=['id'])

from sklearn.preprocessing import OrdinalEncoder
n_samples = train_df.shape[0]
n_features = train_df.shape[1] - 1
target_column = "prognosis"
# Encoder our prognosis labels as integers for easier decoding later
enc = OrdinalEncoder()
train_df[target_column] = enc.fit_transform(train_df[[target_column]])
train_df.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
print(train_df.shape)
train_df.head()
```

(707, 65)

	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	...	x56	x57	x58	x59	x60	x61	x62	x63
0	1.0	1.0	0.0	1.0	1.0	1.0	1.0	0.0	1.0	1.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	1.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0	1.0	1.0	1.0	0.0	1.0	1.0	1.0	1.0	1.0	...	1.0	1.0	1.0	1.0	1.0	0.0	1.0	1.0
3	0.0	0.0	1.0	1.0	1.0	1.0	0.0	1.0	0.0	1.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	...	0.0	1.0	0.0	0.0	1.0	1.0	1.0	0.0

The full data set `train_df` 64 features. The target column is labeled as `prognosis`.

15.3.2 Holdout Train and Test Data

We split out a hold-out test set (25% of the data) so we can calculate an example MAP@K

```

import numpy as np
from sklearn.model_selection import train_test_split
X_train, X_test, y_train, y_test = train_test_split(train_df.drop(target_column, axis=1),
                                                    random_state=42,
                                                    test_size=0.25,
                                                    stratify=train_df[target_column])
train = pd.DataFrame(np.hstack((X_train, np.array(y_train).reshape(-1, 1))))
test = pd.DataFrame(np.hstack((X_test, np.array(y_test).reshape(-1, 1))))
train.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
test.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
print(train.shape)
print(test.shape)
train.head()

```

(530, 65)

(177, 65)

	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	...	x56	x57	x58	x59	x60	x61	x62	x63
0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0	0.0	0.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	1.0	1.0	1.0	0.0
3	1.0	1.0	0.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	0.0	0.0	0.0	1.0	0.0	0.0	1.0	1.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

```

# add the dataset to the fun_control
fun_control.update({"data": train_df, # full dataset,
                  "train": train,
                  "test": test,
                  "n_samples": n_samples,
                  "target_column": target_column})

```

15.4 Step 4: Specification of the Preprocessing Model

Data preprocessing can be very simple, e.g., you can ignore it. Then you would choose the `prep_model` "None":

```

prep_model = None
fun_control.update({"prep_model": prep_model})

```

A default approach for numerical data is the `StandardScaler` (mean 0, variance 1). This can be selected as follows:

```
# prep_model = StandardScaler()
# fun_control.update({"prep_model": prep_model})
```

Even more complicated pre-processing steps are possible, e.g., the following pipeline:

```
# categorical_columns = []
# one_hot_encoder = OneHotEncoder(handle_unknown="ignore", sparse_output=False)
# prep_model = ColumnTransformer(
#     transformers=[
#         ("categorical", one_hot_encoder, categorical_columns),
#     ],
#     remainder=StandardScaler(),
# )
```

15.5 Step 5: Select Model (algorithm) and `core_model_hyper_dict`

The selection of the algorithm (ML model) that should be tuned is done by specifying the its name from the `sklearn` implementation. For example, the SVC support vector machine classifier is selected as follows:

```
fun_control = add_core_model_to_fun_control(SVC, fun_control, SklearnHyperDict)
```

Other `core_models` are, e.g.,:

- `RidgeCV`
- `GradientBoostingRegressor`
- `ElasticNet`
- `RandomForestClassifier`
- `LogisticRegression`
- `KNeighborsClassifier`
- `RandomForestClassifier`
- `GradientBoostingClassifier`
- `HistGradientBoostingClassifier`

We will use the `RandomForestClassifier` classifier in this example.

```

from sklearn.linear_model import RidgeCV
from sklearn.ensemble import RandomForestClassifier
from sklearn.svm import SVC
from sklearn.linear_model import LogisticRegression
from sklearn.neighbors import KNeighborsClassifier
from sklearn.ensemble import GradientBoostingClassifier
from sklearn.ensemble import GradientBoostingRegressor
from sklearn.linear_model import ElasticNet
from spotPython.hyperparameters.values import add_core_model_to_fun_control
from spotPython.data.sklearn_hyper_dict import SklearnHyperDict
from spotPython.fun.hypersklearn import HyperSklearn

# core_model = RidgeCV
# core_model = GradientBoostingRegressor
# core_model = ElasticNet
core_model = RandomForestClassifier
# core_model = SVC
# core_model = LogisticRegression
# core_model = KNeighborsClassifier
# core_model = GradientBoostingClassifier
fun_control = add_core_model_to_fun_control(core_model=core_model,
                                          fun_control=fun_control,
                                          hyper_dict=SklearnHyperDict,
                                          filename=None)

```

Now `fun_control` has the information from the JSON file. The available hyperparameters are:

```
print(*fun_control["core_model_hyper_dict"].keys(), sep="\n")
```

```

n_estimators
criterion
max_depth
min_samples_split
min_samples_leaf
min_weight_fraction_leaf
max_features
max_leaf_nodes
min_impurity_decrease
bootstrap
oob_score

```

15.6 Step 6: Modify `hyper_dict` Hyperparameters for the Selected Algorithm aka `core_model`

15.6.1 Modify hyperparameter of type numeric and integer (boolean)

Numeric and boolean values can be modified using the `modify_hyper_parameter_bounds` method. For example, to change the `tol` hyperparameter of the `SVC` model to the interval `[1e-3, 1e-2]`, the following code can be used:

```
fun_control = modify_hyper_parameter_bounds(fun_control, "tol", bounds=[1e-3, 1e-2])
```

```
from spotPython.hyperparameters.values import modify_hyper_parameter_bounds
# fun_control = modify_hyper_parameter_bounds(fun_control, "tol", bounds=[1e-3, 1e-2])
```

15.6.2 Modify hyperparameter of type factor

`spotPython` provides functions for modifying the hyperparameters, their bounds and factors as well as for activating and de-activating hyperparameters without re-compilation of the Python source code. These functions were described in Section 14.6.

Factors can be modified with the `modify_hyper_parameter_levels` function. For example, to exclude the `sigmoid` kernel from the tuning, the `kernel` hyperparameter of the `SVC` model can be modified as follows:

```
fun_control = modify_hyper_parameter_levels(fun_control, "kernel", ["linear", "rbf"])
```

The new setting can be controlled via:

```
fun_control["core_model_hyper_dict"]["kernel"]
```

```
from spotPython.hyperparameters.values import modify_hyper_parameter_levels
# XGBoost:
# fun_control = modify_hyper_parameter_levels(fun_control, "loss", ["log_loss"])
```

i Note: `RandomForestClassifier` and Out-of-bag Estimation

Since `oob_score` requires the `bootstrap` hyperparameter to `True`, we set the `oob_score` parameter to `False`. The `oob_score` is later discussed in Section 15.7.3.

```
fun_control = modify_hyper_parameter_bounds(fun_control, "bootstrap", bounds=[0, 1])
fun_control = modify_hyper_parameter_bounds(fun_control, "oob_score", bounds=[0, 0])
```

15.6.3 Optimizers

Optimizers are described in Section [14.6.1](#).

15.6.4 Selection of the Objective: Metric and Loss Functions

- Machine learning models are optimized with respect to a metric, for example, the accuracy function.
- Deep learning, e.g., neural networks are optimized with respect to a loss function, for example, the `cross_entropy` function and evaluated with respect to a metric, for example, the accuracy function.

15.7 Step 7: Selection of the Objective (Loss) Function

The loss function, that is usually used in deep learning for optimizing the weights of the net, is stored in the `fun_control` dictionary as `"loss_function"`.

15.7.1 Metric Function

There are two different types of metrics in `spotPython`:

1. `"metric_river"` is used for the river based evaluation via `eval_oml_iter_progressive`.
2. `"metric_sklearn"` is used for the sklearn based evaluation.

We will consider multi-class classification metrics, e.g., `mapk_score` and `top_k_accuracy_score`.

i Predict Probabilities

In this multi-class classification example the machine learning algorithm should return the probabilities of the specific classes (`"predict_proba"`) instead of the predicted values.

We set `"predict_proba"` to `True` in the `fun_control` dictionary.

15.7.1.1 The MAPK Metric

To select the MAPK metric, the following two entries can be added to the `fun_control` dictionary:

```
"metric_sklearn": mapk_score"  
"metric_params": {"k": 3}.
```

15.7.1.2 Other Metrics

Alternatively, other metrics for multi-class classification can be used, e.g., `* top_k_accuracy_score` or `* roc_auc_score`

The metric `roc_auc_score` requires the parameter `"multi_class"`, e.g.,

```
"multi_class": "ovr".
```

This is set in the `fun_control` dictionary.

i Weights

`spotPython` performs a minimization, therefore, metrics that should be maximized have to be multiplied by `-1`. This is done by setting `"weights"` to `-1`.

- The complete setup for the metric in our example is:

```
from spotPython.utils.metrics import mapk_score  
fun_control.update({  
    "weights": -1,  
    "metric_sklearn": mapk_score,  
    "predict_proba": True,  
    "metric_params": {"k": 3},  
})
```

15.7.2 Evaluation on Hold-out Data

- The default method for computing the performance is `"eval_holdout"`.
- Alternatively, cross-validation can be used for every machine learning model.
- Specifically for `RandomForests`, the OOB-score can be used.

```
fun_control.update({
    "eval": "train_hold_out",
})
```

15.7.3 OOB Score

Using the OOB-Score is a very efficient way to estimate the performance of a random forest classifier. The OOB-Score is calculated on the training data and does not require a hold-out test set. If the OOB-Score is used, the key “eval” in the `fun_control` dictionary should be set to `"oob_score"` as shown below.

i OOB-Score

In addition to setting the key `"eval"` in the `fun_control` dictionary to `"oob_score"`, the keys `"oob_score"` and `"bootstrap"` have to be set to `True`, because the OOB-Score requires the bootstrap method.

- Uncomment the following lines to use the OOB-Score:

```
fun_control.update({
    "eval": "eval_oob_score",
})
fun_control = modify_hyper_parameter_bounds(fun_control, "bootstrap", bounds=[1, 1])
fun_control = modify_hyper_parameter_bounds(fun_control, "oob_score", bounds=[1, 1])
```

15.7.3.1 Cross Validation

Instead of using the OOB-score, the classical cross validation can be used. The number of folds is set by the key `"k_folds"`. For example, to use 5-fold cross validation, the key `"k_folds"` is set to 5. Uncomment the following line to use cross validation:

```
# fun_control.update({
#     "eval": "train_cv",
#     "k_folds": 10,
# })
```

15.8 Step 8: Calling the SPOT Function

15.8.1 Preparing the SPOT Call

- Get types and variable names as well as lower and upper bounds for the hyperparameters.

```
# extract the variable types, names, and bounds
from spotPython.hyperparameters.values import (get_bound_values,
        get_var_name,
        get_var_type,)
var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
fun_control.update({"var_type": var_type,
                   "var_name": var_name})
lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")

from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))
```

name	type	default	lower	upper	transform
n_estimators	int	7	5	10	transform_power_2_int
criterion	factor	gini	0	2	None
max_depth	int	10	1	20	transform_power_2_int
min_samples_split	int	2	2	100	None
min_samples_leaf	int	1	1	25	None
min_weight_fraction_leaf	float	0.0	0	0.01	None
max_features	factor	sqrt	0	1	transform_none_to_None
max_leaf_nodes	int	10	7	12	transform_power_2_int
min_impurity_decrease	float	0.0	0	0.01	None
bootstrap	factor	1	1	1	None
oob_score	factor	0	1	1	None

15.8.2 The Objective Function

The objective function is selected next. It implements an interface from `sklearn`'s training, validation, and testing methods to `spotPython`.

```
from spotPython.fun.hyper sklearn import HyperSklearn
fun = HyperSklearn().fun_sklearn
```

15.8.3 Run the Spot Optimizer

- Run SPOT for approx. x mins (`max_time`).
- Note: the run takes longer, because the evaluation time of initial design (here: `init_size`, 20 points) is not considered.

```
from spotPython.hyperparameters.values import get_default_hyperparameters_as_array
hyper_dict=SklearnHyperDict().load()
X_start = get_default_hyperparameters_as_array(fun_control, hyper_dict)
X_start
```

```
array([[ 7.,  0., 10.,  2.,  1.,  0.,  0., 10.,  0.,  1.,  0.]])
```

```
import numpy as np
from spotPython.spot import spot
from math import inf
spot_tuner = spot.Spot(fun=fun,
                      lower = lower,
                      upper = upper,
                      fun_evals = inf,
                      fun_repeats = 1,
                      max_time = MAX_TIME,
                      noise = False,
                      tolerance_x = np.sqrt(np.spacing(1)),
                      var_type = var_type,
                      var_name = var_name,
                      infill_criterion = "y",
                      n_points = 1,
                      seed=123,
                      log_level = 50,
                      show_models= False,
                      show_progress= True,
                      fun_control = fun_control,
                      design_control={"init_size": INIT_SIZE,
                                     "repeats": 1},
                      surrogate_control={"noise": True,
                                       "cod_type": "norm",
```

```
        "min_theta": -4,  
        "max_theta": 3,  
        "n_theta": len(var_name),  
        "model_fun_evals": 10_000,  
        "log_level": 50  
    })  
spot_tuner.run(X_start=X_start)
```

```
spotPython tuning: -0.3361635220125786 [-----] 1.39%  
spotPython tuning: -0.33962264150943394 [-----] 2.82%  
spotPython tuning: -0.33962264150943394 [-----] 4.50%  
spotPython tuning: -0.33962264150943394 [#-----] 5.84%  
spotPython tuning: -0.33962264150943394 [#-----] 7.50%  
spotPython tuning: -0.33962264150943394 [#-----] 9.59%  
spotPython tuning: -0.33962264150943394 [#-----] 11.08%  
spotPython tuning: -0.35094339622641507 [#-----] 13.59%  
spotPython tuning: -0.35094339622641507 [#-----] 14.87%  
spotPython tuning: -0.35094339622641507 [##-----] 15.84%  
spotPython tuning: -0.3641509433962264 [##-----] 18.40%  
spotPython tuning: -0.3641509433962264 [##-----] 20.81%  
spotPython tuning: -0.3641509433962264 [###-----] 25.00%  
spotPython tuning: -0.3641509433962264 [###-----] 29.38%  
spotPython tuning: -0.3641509433962264 [###-----] 34.01%
```

```
spotPython tuning: -0.3641509433962264 [####-----] 39.17%
spotPython tuning: -0.3641509433962264 [####-----] 43.64%
spotPython tuning: -0.3694968553459119 [#####-----] 47.48%
spotPython tuning: -0.3694968553459119 [#####-----] 51.84%
spotPython tuning: -0.3694968553459119 [#####-----] 57.05%
spotPython tuning: -0.3694968553459119 [#####-----] 61.60%
spotPython tuning: -0.3694968553459119 [#####-----] 65.70%
spotPython tuning: -0.3694968553459119 [#####-----] 69.71%
spotPython tuning: -0.3694968553459119 [#####-----] 73.86%
spotPython tuning: -0.3694968553459119 [#####-----] 77.96%
spotPython tuning: -0.3694968553459119 [#####-----] 82.20%
spotPython tuning: -0.3694968553459119 [#####-----] 86.53%
spotPython tuning: -0.3694968553459119 [#####-----] 92.21%
spotPython tuning: -0.3694968553459119 [#####-----] 98.87%
spotPython tuning: -0.3694968553459119 [#####-----] 100.00% Done...

<spotPython.spot.spot.Spot at 0x2b9d75f60>
```

15.9 Step 9: Tensorboard

The textual output shown in the console (or code cell) can be visualized with Tensorboard as described in Section 14.9, see also the description in the documentation: [Tensorboard](#).

15.10 Step 10: Results

After the hyperparameter tuning run is finished, the progress of the hyperparameter tuning can be visualized. The following code generates the progress plot from `?@fig-progress`.

```
spot_tuner.plot_progress(log_y=False,  
                        filename="./figures/" + experiment_name+"_progress.png")
```




Figure 15.1: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

- Print the results

```
print(gen_design_table(fun_control=fun_control,  
                      spot=spot_tuner))
```

name	type	default	lower	upper	tuned
n_estimators	int	7	5.0	10.0	10.0
criterion	factor	gini	0.0	2.0	1.0
max_depth	int	10	1.0	20.0	8.0
min_samples_split	int	2	2.0	100.0	15.0
min_samples_leaf	int	1	1.0	25.0	1.0
min_weight_fraction_leaf	float	0.0	0.0	0.01	0.009298535461845475
max_features	factor	sqrt	0.0	1.0	0.0
max_leaf_nodes	int	10	7.0	12.0	10.0
min_impurity_decrease	float	0.0	0.0	0.01	0.0
bootstrap	factor	1	1.0	1.0	1.0
oob_score	factor	0	1.0	1.0	1.0

15.10.1 Show variable importance

```
spot_tuner.plot_importance(threshold=0.025, filename="./figures/" + experiment_name+"_imp
```




Figure 15.2: Variable importance plot, threshold 0.025.

15.10.2 Get Default Hyperparameters

```
from spotPython.hyperparameters.values import get_default_values, transform_hyper_parameter_values_default = get_default_values(fun_control) values_default = transform_hyper_parameter_values(fun_control=fun_control, hyper_parameter_values_default
```

```
{'n_estimators': 128, 'criterion': 'gini', 'max_depth': 1024, 'min_samples_split': 2, 'min_samples_leaf': 1, 'min_weight_fraction_leaf': 0.0, 'max_features': 'sqrt', 'max_leaf_nodes': 1024, 'min_impurity_decrease': 0.0,
```

```
'bootstrap': 1,  
'oob_score': 0}
```

```
from sklearn.pipeline import make_pipeline  
model_default = make_pipeline(fun_control["prep_model"], fun_control["core_model"](**value  
model_default
```

```
Pipeline(steps=[('nonetype', None),  
                ('randomforestclassifier',  
                 RandomForestClassifier(bootstrap=1, max_depth=1024,  
                                       max_leaf_nodes=1024, n_estimators=128,  
                                       oob_score=0))])
```

15.10.3 Get SPOT Results

```
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))  
print(X)
```

```
[[1.00000000e+01 1.00000000e+00 8.00000000e+00 1.50000000e+01  
 1.00000000e+00 9.29853546e-03 0.00000000e+00 1.00000000e+01  
 0.00000000e+00 1.00000000e+00 1.00000000e+00]]
```

```
from spotPython.hyperparameters.values import assign_values, return_conf_list_from_var_dict  
v_dict = assign_values(X, fun_control["var_name"])  
return_conf_list_from_var_dict(var_dict=v_dict, fun_control=fun_control)
```

```
[{'n_estimators': 1024,  
  'criterion': 'entropy',  
  'max_depth': 256,  
  'min_samples_split': 15,  
  'min_samples_leaf': 1,  
  'min_weight_fraction_leaf': 0.009298535461845475,  
  'max_features': 'sqrt',  
  'max_leaf_nodes': 1024,  
  'min_impurity_decrease': 0.0,  
  'bootstrap': 1,  
  'oob_score': 1}]
```

```

from spotPython.hyperparameters.values import get_one_sklearn_model_from_X
model_spot = get_one_sklearn_model_from_X(X, fun_control)
model_spot

```

```

RandomForestClassifier(bootstrap=1, criterion='entropy', max_depth=256,
                        max_leaf_nodes=1024, min_samples_split=15,
                        min_weight_fraction_leaf=0.009298535461845475,
                        n_estimators=1024, oob_score=1)

```

15.10.4 Evaluate SPOT Results

- Fetch the data.

```

from spotPython.utils.convert import get_Xy_from_df
X_train, y_train = get_Xy_from_df(fun_control["train"], fun_control["target_column"])
X_test, y_test = get_Xy_from_df(fun_control["test"], fun_control["target_column"])
X_test.shape, y_test.shape

```

```
((177, 64), (177,))
```

- Fit the model with the tuned hyperparameters. This gives one result:

```

model_spot.fit(X_train, y_train)
y_pred = model_spot.predict_proba(X_test)
res = mapk_score(y_true=y_test, y_pred=y_pred, k=3)
res

```

```
0.35404896421845566
```

```

def repeated_eval(n, model):
    res_values = []
    for i in range(n):
        model.fit(X_train, y_train)
        y_pred = model.predict_proba(X_test)
        res = mapk_score(y_true=y_test, y_pred=y_pred, k=3)
        res_values.append(res)
    mean_res = np.mean(res_values)
    print(f"mean_res: {mean_res}")
    std_res = np.std(res_values)

```

```

print(f"std_res: {std_res}")
min_res = np.min(res_values)
print(f"min_res: {min_res}")
max_res = np.max(res_values)
print(f"max_res: {max_res}")
median_res = np.median(res_values)
print(f"median_res: {median_res}")
return mean_res, std_res, min_res, max_res, median_res

```

15.10.5 Handling Non-deterministic Results

- Because the model is non-deterministic, we perform $n = 30$ runs and calculate the mean and standard deviation of the performance metric.

```
_ = repeated_eval(30, model_spot)
```

```

mean_res: 0.3603578154425611
std_res: 0.005587383690827419
min_res: 0.3483992467043315
max_res: 0.37099811676082867
median_res: 0.3615819209039548

```

15.10.6 Evaluation of the Default Hyperparameters

```
model_default.fit(X_train, y_train)["randomforestclassifier"]
```

```

RandomForestClassifier(bootstrap=1, max_depth=1024, max_leaf_nodes=1024,
                        n_estimators=128, oob_score=0)

```

- One evaluation of the default hyperparameters is performed on the hold-out test set.

```

y_pred = model_default.predict_proba(X_test)
mapk_score(y_true=y_test, y_pred=y_pred, k=3)

```

```
0.3295668549905838
```


Since one single evaluation is not meaningful, we perform, similar to the evaluation of the SPOT results, $n = 30$ runs of the default setting and calculate the mean and standard deviation of the performance metric.

```
_ = repeated_eval(30, model_default)
```


```
mean_res: 0.338480853735091
std_res: 0.013688038473151918
min_res: 0.31544256120527303
max_res: 0.3681732580037665
median_res: 0.3366290018832392
```

15.10.7 Plot: Compare Predictions

```
from spotPython.plot.validation import plot_confusion_matrix
plot_confusion_matrix(model_default, fun_control, title = "Default")
```



```
plot_confusion_matrix(model_spot, fun_control, title="SPOT")
```



```
min(spot_tuner.y), max(spot_tuner.y)
```

```
(-0.3694968553459119, -0.2858490566037736)
```

15.10.8 Cross-validated Evaluations

```
from spotPython.sklearn.traintest import evaluate_cv
fun_control.update({
    "eval": "train_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)
```

```
(0.3660377358490566, None)
```

```
fun_control.update({
    "eval": "test_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)
```

(0.29793028322440085, None)

- This is the evaluation that will be used in the comparison:







```
fun_control.update({
    "eval": "data_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)
```







(0.3598859825620389, None)







15.10.9 Detailed Hyperparameter Plots



```
filename = "./figures/" + experiment_name
spot_tuner.plot_important_hyperparameter_contour(filename=filename)
```

```
n_estimators: 0.06242752772054072
min_samples_split: 0.36314083309831824
min_samples_leaf: 0.04240927390389511
max_features: 100.0
max_leaf_nodes: 100.0
```









15.10.10 Parallel Coordinates Plot

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Unable to display output for mime type(s): text/html

15.10.11 Plot all Combinations of Hyperparameters

- Warning: this may take a while.

```
PLOT_ALL = False
if PLOT_ALL:
    n = spot_tuner.k
    for i in range(n-1):
        for j in range(i+1, n):
            spot_tuner.plot_contour(i=i, j=j, min_z=min_z, max_z = max_z)
```

16 HPT: sklearn XGB Classifier VBDP Data

This document refers to the following software versions:

- python: 3.10.10

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

spotPython can be installed via pip. Alternatively, the source code can be downloaded from gitHub: <https://github.com/sequential-parameter-optimization/spotPython>.

```
!pip install spotPython
```

- Uncomment the following lines if you want to for (re-)installation the latest version of spotPython from gitHub.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

16.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time and the initial design size.

```
MAX_TIME = 1
INIT_SIZE = 5
ORIGINAL = False
```

```

import os
import copy
import socket
from datetime import datetime
from dateutil.tz import tzlocal
start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '17-xgb-sklearn' + "_" + HOSTNAME + "_" + str(MAX_TIME) + "min_" + str(I
experiment_name = experiment_name.replace(':', '-')
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')

```

17-xgb-sklearn_bartz09_1min_5init_2023-06-19_04-09-05

```

import warnings
warnings.filterwarnings("ignore")

```

16.2 Step 2: Initialization of the Empty fun_control Dictionary

 Caution: Tensorboard does not work under Windows

- Since tensorboard does not work under Windows, we recommend setting the parameter `tensorboard_path` to `None` if you are working under Windows.

```

from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="classification",
    tensorboard_path="runs/16_spot_hpt_sklearn_classification")

```

16.3 Step 3: PyTorch Data Loading

16.3.1 1. Load Data: Classification VBDP

```
import pandas as pd
if ORIGINAL == True:
    train_df = pd.read_csv('./data/VBDP/trainnn.csv')
    test_df = pd.read_csv('./data/VBDP/testtt.csv')
else:
    train_df = pd.read_csv('./data/VBDP/train.csv')
    # remove the id column
    train_df = train_df.drop(columns=['id'])

from sklearn.preprocessing import OrdinalEncoder
n_samples = train_df.shape[0]
n_features = train_df.shape[1] - 1
target_column = "prognosis"
# Encoder our prognosis labels as integers for easier decoding later
enc = OrdinalEncoder()
train_df[target_column] = enc.fit_transform(train_df[[target_column]])
train_df.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
print(train_df.shape)
train_df.head()
```

(707, 65)

	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	...	x56	x57	x58	x59	x60	x61	x62	x63
0	1.0	1.0	0.0	1.0	1.0	1.0	1.0	0.0	1.0	1.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	1.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0	1.0	1.0	1.0	0.0	1.0	1.0	1.0	1.0	1.0	...	1.0	1.0	1.0	1.0	1.0	0.0	1.0	1.0
3	0.0	0.0	1.0	1.0	1.0	1.0	0.0	1.0	0.0	1.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	...	0.0	1.0	0.0	0.0	1.0	1.0	1.0	0.0

The full data set `train_df` 64 features. The target column is labeled as `prognosis`.

16.3.2 Holdout Train and Test Data

We split out a hold-out test set (25% of the data) so we can calculate an example MAP@K

```

import numpy as np
from sklearn.model_selection import train_test_split
X_train, X_test, y_train, y_test = train_test_split(train_df.drop(target_column, axis=1),
                                                    random_state=42,
                                                    test_size=0.25,
                                                    stratify=train_df[target_column])
train = pd.DataFrame(np.hstack((X_train, np.array(y_train).reshape(-1, 1))))
test = pd.DataFrame(np.hstack((X_test, np.array(y_test).reshape(-1, 1))))
train.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
test.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
print(train.shape)
print(test.shape)
train.head()

```

(530, 65)

(177, 65)

	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	...	x56	x57	x58	x59	x60	x61	x62	x63
0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0	0.0	0.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	1.0	1.0	1.0	0.0
3	1.0	1.0	0.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	0.0	0.0	0.0	1.0	0.0	0.0	1.0	1.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

```

# add the dataset to the fun_control
fun_control.update({"data": train_df, # full dataset,
                  "train": train,
                  "test": test,
                  "n_samples": n_samples,
                  "target_column": target_column})

```

16.4 Step 4: Specification of the Preprocessing Model

Data preprocessing can be very simple, e.g., you can ignore it. Then you would choose the `prep_model` "None":

```

prep_model = None
fun_control.update({"prep_model": prep_model})

```

A default approach for numerical data is the `StandardScaler` (mean 0, variance 1). This can be selected as follows:

```
# prep_model = StandardScaler()
# fun_control.update({"prep_model": prep_model})
```

Even more complicated pre-processing steps are possible, e.g., the following pipeline:

```
# categorical_columns = []
# one_hot_encoder = OneHotEncoder(handle_unknown="ignore", sparse_output=False)
# prep_model = ColumnTransformer(
#     transformers=[
#         ("categorical", one_hot_encoder, categorical_columns),
#     ],
#     remainder=StandardScaler(),
# )
```

16.5 Step 5: Select Model (algorithm) and `core_model_hyper_dict`

The selection of the algorithm (ML model) that should be tuned is done by specifying the its name from the `sklearn` implementation. For example, the SVC support vector machine classifier is selected as follows:

```
fun_control = add_core_model_to_fun_control(SVC, fun_control, SklearnHyperDict)
```

Other `core_models` are, e.g.,:

- `RidgeCV`
- `GradientBoostingRegressor`
- `ElasticNet`
- `RandomForestClassifier`
- `LogisticRegression`
- `KNeighborsClassifier`
- `RandomForestClassifier`
- `GradientBoostingClassifier`
- `HistGradientBoostingClassifier`

We will use the `RandomForestClassifier` classifier in this example.

```

from sklearn.linear_model import RidgeCV
from sklearn.ensemble import RandomForestClassifier
from sklearn.svm import SVC
from sklearn.linear_model import LogisticRegression
from sklearn.neighbors import KNeighborsClassifier
from sklearn.ensemble import GradientBoostingClassifier
from sklearn.ensemble import GradientBoostingRegressor
from sklearn.ensemble import HistGradientBoostingClassifier
from sklearn.linear_model import ElasticNet
from spotPython.hyperparameters.values import add_core_model_to_fun_control
from spotPython.data.sklearn_hyper_dict import SklearnHyperDict
from spotPython.fun.hypersklearn import HyperSklearn

```

```

# core_model = RidgeCV
# core_model = GradientBoostingRegressor
# core_model = ElasticNet
core_model = RandomForestClassifier
# core_model = SVC
# core_model = LogisticRegression
# core_model = KNeighborsClassifier
# core_model = GradientBoostingClassifier
core_model = HistGradientBoostingClassifier
fun_control = add_core_model_to_fun_control(core_model=core_model,
                                          fun_control=fun_control,
                                          hyper_dict=SklearnHyperDict,
                                          filename=None)

```

Now `fun_control` has the information from the JSON file. The available hyperparameters are:

```

print(*fun_control["core_model_hyper_dict"].keys(), sep="\n")

```

```

loss
learning_rate
max_iter
max_leaf_nodes
max_depth
min_samples_leaf
l2_regularization
max_bins
early_stopping

```

```
n_iter_no_change
tol
```

16.6 Step 6: Modify `hyper_dict` Hyperparameters for the Selected Algorithm aka `core_model`

16.6.1 Modify hyperparameter of type numeric and integer (boolean)

Numeric and boolean values can be modified using the `modify_hyper_parameter_bounds` method. For example, to change the `tol` hyperparameter of the `SVC` model to the interval `[1e-3, 1e-2]`, the following code can be used:

```
fun_control = modify_hyper_parameter_bounds(fun_control, "tol", bounds=[1e-3,
1e-2])
```

```
from spotPython.hyperparameters.values import modify_hyper_parameter_bounds
# fun_control = modify_hyper_parameter_bounds(fun_control, "tol", bounds=[1e-3, 1e-2])
# fun_control = modify_hyper_parameter_bounds(fun_control, "min_samples_split", bounds=[3,
# fun_control = modify_hyper_parameter_bounds(fun_control, "dual", bounds=[0, 0])
# fun_control = modify_hyper_parameter_bounds(fun_control, "probability", bounds=[1, 1])
# fun_control["core_model_hyper_dict"]["tol"]
# fun_control = modify_hyper_parameter_bounds(fun_control, "min_samples_leaf", bounds=[1,
# fun_control = modify_hyper_parameter_bounds(fun_control, "n_estimators", bounds=[5, 10])
```

16.6.2 Modify hyperparameter of type factor

`spotPython` provides functions for modifying the hyperparameters, their bounds and factors as well as for activating and de-activating hyperparameters without re-compilation of the Python source code. These functions were described in Section [14.6](#).

Factors can be modified with the `modify_hyper_parameter_levels` function. For example, to exclude the `sigmoid` kernel from the tuning, the `kernel` hyperparameter of the `SVC` model can be modified as follows:

```
fun_control = modify_hyper_parameter_levels(fun_control, "kernel", ["linear",
"rbf"])
```

The new setting can be controlled via:

```
fun_control["core_model_hyper_dict"]["kernel"]
```

```
from spotPython.hyperparameters.values import modify_hyper_parameter_levels
# XGBoost:
fun_control = modify_hyper_parameter_levels(fun_control, "loss", ["log_loss"])
```

16.6.3 Optimizers

Optimizers are described in Section [14.6.1](#).

16.7 Step 7: Selection of the Objective (Loss) Function

16.7.1 Evaluation

The evaluation procedure requires the specification of two elements:

1. the way how the data is split into a train and a test set and
2. the loss function (and a metric).

16.7.2 Selection of the Objective: Metric and Loss Functions

- Machine learning models are optimized with respect to a metric, for example, the accuracy function.
- Deep learning, e.g., neural networks are optimized with respect to a loss function, for example, the `cross_entropy` function and evaluated with respect to a metric, for example, the accuracy function.

16.7.3 Loss Function

The loss function, that is usually used in deep learning for optimizing the weights of the net, is stored in the `fun_control` dictionary as `"loss_function"`.

16.7.4 Metric Function

There are two different types of metrics in `spotPython`:

1. `"metric_river"` is used for the river based evaluation via `eval_oml_iter_progressive`.
2. `"metric_sklearn"` is used for the sklearn based evaluation.

We will consider multi-class classification metrics, e.g., `mapk_score` and `top_k_accuracy_score`.

i Predict Probabilities

In this multi-class classification example the machine learning algorithm should return the probabilities of the specific classes ("predict_proba") instead of the predicted values.

We set "predict_proba" to True in the `fun_control` dictionary.

16.7.4.1 The MAPK Metric

To select the MAPK metric, the following two entries can be added to the `fun_control` dictionary:

```
"metric_sklearn": mapk_score"  
"metric_params": {"k": 3}.
```

16.7.4.2 Other Metrics

Alternatively, other metrics for multi-class classification can be used, e.g.,: * `top_k_accuracy_score` or * `roc_auc_score`

The metric `roc_auc_score` requires the parameter "multi_class", e.g.,

```
"multi_class": "ovr".
```

This is set in the `fun_control` dictionary.

i Weights

spotPython performs a minimization, therefore, metrics that should be maximized have to be multiplied by -1. This is done by setting "weights" to -1.

- The complete setup for the metric in our example is:

```
from spotPython.utils.metrics import mapk_score  
fun_control.update({  
    "weights": -1,  
    "metric_sklearn": mapk_score,  
    "predict_proba": True,  
    "metric_params": {"k": 3},  
})
```

16.7.5 Evaluation on Hold-out Data

- The default method for computing the performance is "eval_holdout".
- Alternatively, cross-validation can be used for every machine learning model.
- Specifically for RandomForests, the OOB-score can be used.

```
fun_control.update({
    "eval": "train_hold_out",
})
```

16.7.5.1 Cross Validation

Instead of using the OOB-score, the classical cross validation can be used. The number of folds is set by the key "k_folds". For example, to use 5-fold cross validation, the key "k_folds" is set to 5. Uncomment the following line to use cross validation:

```
# fun_control.update({
#     "eval": "train_cv",
#     "k_folds": 10,
# })
```

16.8 Step 8: Calling the SPOT Function

16.8.1 Preparing the SPOT Call

- Get types and variable names as well as lower and upper bounds for the hyperparameters.

```
# extract the variable types, names, and bounds
from spotPython.hyperparameters.values import (get_bound_values,
    get_var_name,
    get_var_type,)
var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
fun_control.update({"var_type": var_type,
    "var_name": var_name})
lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")
```

```
from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))
```

name	type	default	lower	upper	transform
loss	factor	log_loss	0	0	None
learning_rate	float	-1.0	-5	0	transform_power_10
max_iter	int	7	3	10	transform_power_2_int
max_leaf_nodes	int	5	1	12	transform_power_2_int
max_depth	int	2	1	20	transform_power_2_int
min_samples_leaf	int	4	2	10	transform_power_2_int
l2_regularization	float	0.0	0	10	None
max_bins	int	255	127	255	None
early_stopping	factor	1	0	1	None
n_iter_no_change	int	10	5	20	None
tol	float	0.0001	1e-05	0.001	None

16.8.2 The Objective Function

The objective function is selected next. It implements an interface from sklearn's training, validation, and testing methods to spotPython.

```
from spotPython.fun.hypersklearn import HyperSklearn
fun = HyperSklearn().fun_sklearn
```

16.8.3 Run the Spot Optimizer

- Run SPOT for approx. x mins (max_time).
- Note: the run takes longer, because the evaluation time of initial design (here: initi_size, 20 points) is not considered.

```
from spotPython.hyperparameters.values import get_default_hyperparameters_as_array
hyper_dict=SklearnHyperDict().load()
X_start = get_default_hyperparameters_as_array(fun_control, hyper_dict)
X_start
```

```
array([[ 0.00e+00, -1.00e+00,  7.00e+00,  5.00e+00,  2.00e+00,  4.00e+00,
         0.00e+00,  2.55e+02,  1.00e+00,  1.00e+01,  1.00e-04]])
```

```

import numpy as np
from spotPython.spot import spot
from math import inf
spot_tuner = spot.Spot(fun=fun,
                       lower = lower,
                       upper = upper,
                       fun_evals = inf,
                       fun_repeats = 1,
                       max_time = MAX_TIME,
                       noise = False,
                       tolerance_x = np.sqrt(np.spacing(1)),
                       var_type = var_type,
                       var_name = var_name,
                       infill_criterion = "y",
                       n_points = 1,
                       seed=123,
                       log_level = 50,
                       show_models= False,
                       show_progress= True,
                       fun_control = fun_control,
                       design_control={"init_size": INIT_SIZE,
                                      "repeats": 1},
                       surrogate_control={"noise": True,
                                         "cod_type": "norm",
                                         "min_theta": -4,
                                         "max_theta": 3,
                                         "n_theta": len(var_name),
                                         "model_fun_evals": 10_000,
                                         "log_level": 50
                                         })

spot_tuner.run(X_start=X_start)

```

spotPython tuning: -0.3082706766917293 [-----] 4.50%

spotPython tuning: -0.32832080200501257 [##-----] 16.24%

spotPython tuning: -0.3446115288220551 [##-----] 22.25%

spotPython tuning: -0.3446115288220551 [##-----] 24.49%

spotPython tuning: -0.3446115288220551 [###-----] 26.75%

```

spotPython tuning: -0.3609022556390977 [####-----] 36.44%

spotPython tuning: -0.3609022556390977 [####-----] 44.05%

spotPython tuning: -0.3609022556390977 [#####----] 84.16%

spotPython tuning: -0.3609022556390977 [#####] 100.00% Done...

<spotPython.spot.spot.Spot at 0x28d365300>

```

16.9 Step 9: Tensorboard

The textual output shown in the console (or code cell) can be visualized with Tensorboard as described in Section 14.9, see also the description in the documentation: [Tensorboard](#).

16.10 Step 10: Results

After the hyperparameter tuning run is finished, the progress of the hyperparameter tuning can be visualized. The following code generates the progress plot from `?@fig-progress`.

```

spot_tuner.plot_progress(log_y=False,
    filename="./figures/" + experiment_name+"_progress.png")

```




Figure 16.1: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

- Print the results

```
print(gen_design_table(fun_control=fun_control,
                      spot=spot_tuner))
```

name	type	default	lower	upper	tuned	trans
loss	factor	log_loss	0.0	0.0	0.0	None
learning_rate	float	-1.0	-5.0	0.0	-1.8860361379529507	trans
max_iter	int	7	3.0	10.0	8.0	trans
max_leaf_nodes	int	5	1.0	12.0	2.0	trans
max_depth	int	2	1.0	20.0	20.0	trans
min_samples_leaf	int	4	2.0	10.0	2.0	trans
l2_regularization	float	0.0	0.0	10.0	5.0648633125794476	None
max_bins	int	255	127.0	255.0	130.0	None
early_stopping	factor	1	0.0	1.0	0.0	None
n_iter_no_change	int	10	5.0	20.0	12.0	None
tol	float	0.0001	1e-05	0.001	0.0009999411112607051	None

16.10.1 Show variable importance

```
spot_tuner.plot_importance(threshold=0.025, filename="./figures/" + experiment_name+"_impo
```




Figure 16.2: Variable importance plot, threshold 0.025.

16.10.2 Get Default Hyperparameters

```
from spotPython.hyperparameters.values import get_default_values, transform_hyper_parameter_values
values_default = get_default_values(fun_control)
values_default = transform_hyper_parameter_values(fun_control=fun_control, hyper_parameter_values=values_default)
```

```
{'loss': 'log_loss',
 'learning_rate': 0.1,
 'max_iter': 128,
 'max_leaf_nodes': 32,
 'max_depth': 4,
 'min_samples_leaf': 16,
 'l2_regularization': 0.0,
 'max_bins': 255,
 'early_stopping': 1,
 'n_iter_no_change': 10,
 'tol': 0.0001}
```

```
from sklearn.pipeline import make_pipeline
model_default = make_pipeline(fun_control["prep_model"], fun_control["core_model"](**values_default))
model_default
```

```
Pipeline(steps=[('nonetype', None),
                 ('histgradientboostingclassifier',
                  HistGradientBoostingClassifier(early_stopping=1, max_depth=4,
                                                  max_iter=128, max_leaf_nodes=32,
                                                  min_samples_leaf=16,
                                                  tol=0.0001))])
```

16.10.3 Get SPOT Results

```
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))
print(X)
```

```
[[ 0.00000000e+00 -1.88603614e+00  8.00000000e+00  2.00000000e+00
  2.00000000e+01  2.00000000e+00  5.06486331e+00  1.30000000e+02
  0.00000000e+00  1.20000000e+01  9.99941111e-04]]
```

```

from spotPython.hyperparameters.values import assign_values, return_conf_list_from_var_dict
v_dict = assign_values(X, fun_control["var_name"])
return_conf_list_from_var_dict(var_dict=v_dict, fun_control=fun_control)

```

```

[{'loss': 'log_loss',
  'learning_rate': 0.013000613944983763,
  'max_iter': 256,
  'max_leaf_nodes': 4,
  'max_depth': 1048576,
  'min_samples_leaf': 4,
  'l2_regularization': 5.0648633125794476,
  'max_bins': 130,
  'early_stopping': 0,
  'n_iter_no_change': 12,
  'tol': 0.0009999411112607051}]

```

```

from spotPython.hyperparameters.values import get_one_sklearn_model_from_X
model_spot = get_one_sklearn_model_from_X(X, fun_control)
model_spot

```

```

HistGradientBoostingClassifier(early_stopping=0,
                                l2_regularization=5.0648633125794476,
                                learning_rate=0.013000613944983763, max_bins=130,
                                max_depth=1048576, max_iter=256,
                                max_leaf_nodes=4, min_samples_leaf=4,
                                n_iter_no_change=12, tol=0.0009999411112607051)

```

16.10.4 Evaluate SPOT Results

- Fetch the data.

```

from spotPython.utils.convert import get_Xy_from_df
X_train, y_train = get_Xy_from_df(fun_control["train"], fun_control["target_column"])
X_test, y_test = get_Xy_from_df(fun_control["test"], fun_control["target_column"])
X_test.shape, y_test.shape

```

```
((177, 64), (177,))
```

- Fit the model with the tuned hyperparameters. This gives one result:

```
model_spot.fit(X_train, y_train)
y_pred = model_spot.predict_proba(X_test)
res = mapk_score(y_true=y_test, y_pred=y_pred, k=3)
res
```

0.3295668549905838

```
def repeated_eval(n, model):
    res_values = []
    for i in range(n):
        model.fit(X_train, y_train)
        y_pred = model.predict_proba(X_test)
        res = mapk_score(y_true=y_test, y_pred=y_pred, k=3)
        res_values.append(res)
    mean_res = np.mean(res_values)
    print(f"mean_res: {mean_res}")
    std_res = np.std(res_values)
    print(f"std_res: {std_res}")
    min_res = np.min(res_values)
    print(f"min_res: {min_res}")
    max_res = np.max(res_values)
    print(f"max_res: {max_res}")
    median_res = np.median(res_values)
    print(f"median_res: {median_res}")
    return mean_res, std_res, min_res, max_res, median_res
```

16.10.5 Handling Non-deterministic Results

- Because the model is non-deterministic, we perform $n = 30$ runs and calculate the mean and standard deviation of the performance metric.

```
_ = repeated_eval(30, model_spot)
```

```
mean_res: 0.3295668549905837
std_res: 1.1102230246251565e-16
min_res: 0.3295668549905838
max_res: 0.3295668549905838
median_res: 0.3295668549905838
```

16.10.6 Evaluation of the Default Hyperparameters

```
model_default.fit(X_train, y_train)["histgradientboostingclassifier"]
```

```
HistGradientBoostingClassifier(early_stopping=1, max_depth=4, max_iter=128,  
                                max_leaf_nodes=32, min_samples_leaf=16,  
                                tol=0.0001)
```

- One evaluation of the default hyperparameters is performed on the hold-out test set.

```
y_pred = model_default.predict_proba(X_test)  
mapk_score(y_true=y_test, y_pred=y_pred, k=3)
```

```
0.3389830508474576
```


Since one single evaluation is not meaningful, we perform, similar to the evaluation of the SPOT results, $n = 30$ runs of the default setting and calculate the mean and standard deviation of the performance metric.

```
_ = repeated_eval(30, model_default)
```


```
mean_res: 0.3375706214689266  
std_res: 0.012500864451367001  
min_res: 0.3107344632768362  
max_res: 0.3615819209039548  
median_res: 0.3408662900188324
```

16.10.7 Plot: Compare Predictions

```
from spotPython.plot.validation import plot_confusion_matrix  
plot_confusion_matrix(model_default, fun_control, title = "Default")
```



```
plot_confusion_matrix(model_spot, fun_control, title="SPOT")
```



```
min(spot_tuner.y), max(spot_tuner.y)
```

```
(-0.3609022556390977, -0.20927318295739344)
```

16.10.8 Cross-validated Evaluations

```
from spotPython.sklearn.traintest import evaluate_cv
fun_control.update({
    "eval": "train_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)
```

```
(0.35062893081761004, None)
```

```
fun_control.update({
    "eval": "test_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)
```

(0.30468409586056644, None)

- This is the evaluation that will be used in the comparison:







```
fun_control.update({
    "eval": "data_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)
```







(0.3542924211938296, None)







16.10.9 Detailed Hyperparameter Plots



```
filename = "./figures/" + experiment_name
spot_tuner.plot_important_hyperparameter_contour(filename=filename)
```

```
learning_rate: 22.139137766465417
min_samples_leaf: 7.571463764452644
l2_regularization: 12.511775110768742
early_stopping: 0.10772129685854573
tol: 100.0
```









16.10.10 Parallel Coordinates Plot

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Unable to display output for mime type(s): text/html

16.10.11 Plot all Combinations of Hyperparameters

- Warning: this may take a while.

```
PLOT_ALL = False
if PLOT_ALL:
    n = spot_tuner.k
    for i in range(n-1):
        for j in range(i+1, n):
            spot_tuner.plot_contour(i=i, j=j, min_z=min_z, max_z = max_z)
```

17 HPT: sklearn SVC VBDP Data

This document refers to the following software versions:

- python: 3.10.10

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

spotPython can be installed via pip. Alternatively, the source code can be downloaded from gitHub: <https://github.com/sequential-parameter-optimization/spotPython>.

```
!pip install spotPython
```

- Uncomment the following lines if you want to for (re-)installation the latest version of spotPython from gitHub.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

17.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time and the initial design size.

```
MAX_TIME = 1
INIT_SIZE = 5
ORIGINAL = False
```

```

import os
import copy
import socket
from datetime import datetime
from dateutil.tz import tzlocal
start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '18-svc-sklearn' + "_" + HOSTNAME + "_" + str(MAX_TIME) + "min_" + str(I
experiment_name = experiment_name.replace(':', '-')
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')

```

18-svc-sklearn_bartz09_1min_5init_2023-06-19_04-16-46

```

import warnings
warnings.filterwarnings("ignore")

```

17.2 Step 2: Initialization of the Empty fun_control Dictionary

 Caution: Tensorboard does not work under Windows

- Since tensorboard does not work under Windows, we recommend setting the parameter `tensorboard_path` to `None` if you are working under Windows.

```

from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="classification",
    tensorboard_path="runs/16_spot_hpt_sklearn_classification")

```

17.3 Step 3: PyTorch Data Loading

17.3.1 1. Load Data: Classification VBDP

```
import pandas as pd
if ORIGINAL == True:
    train_df = pd.read_csv('./data/VBDP/trainnn.csv')
    test_df = pd.read_csv('./data/VBDP/testtt.csv')
else:
    train_df = pd.read_csv('./data/VBDP/train.csv')
    # remove the id column
    train_df = train_df.drop(columns=['id'])

from sklearn.preprocessing import OrdinalEncoder
n_samples = train_df.shape[0]
n_features = train_df.shape[1] - 1
target_column = "prognosis"
# Encoder our prognosis labels as integers for easier decoding later
enc = OrdinalEncoder()
train_df[target_column] = enc.fit_transform(train_df[[target_column]])
train_df.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
print(train_df.shape)
train_df.head()
```

(707, 65)

	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	...	x56	x57	x58	x59	x60	x61	x62	x63
0	1.0	1.0	0.0	1.0	1.0	1.0	1.0	0.0	1.0	1.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	1.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0	1.0	1.0	1.0	0.0	1.0	1.0	1.0	1.0	1.0	...	1.0	1.0	1.0	1.0	1.0	0.0	1.0	1.0
3	0.0	0.0	1.0	1.0	1.0	1.0	0.0	1.0	0.0	1.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	...	0.0	1.0	0.0	0.0	1.0	1.0	1.0	0.0

The full data set `train_df` 64 features. The target column is labeled as `prognosis`.

17.3.2 Holdout Train and Test Data

We split out a hold-out test set (25% of the data) so we can calculate an example MAP@K

```

import numpy as np
from sklearn.model_selection import train_test_split
X_train, X_test, y_train, y_test = train_test_split(train_df.drop(target_column, axis=1),
                                                    random_state=42,
                                                    test_size=0.25,
                                                    stratify=train_df[target_column])
train = pd.DataFrame(np.hstack((X_train, np.array(y_train).reshape(-1, 1))))
test = pd.DataFrame(np.hstack((X_test, np.array(y_test).reshape(-1, 1))))
train.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
test.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
print(train.shape)
print(test.shape)
train.head()

```

(530, 65)

(177, 65)

	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	...	x56	x57	x58	x59	x60	x61	x62	x63
0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0	0.0	0.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	1.0	1.0	1.0	0.0
3	1.0	1.0	0.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	0.0	0.0	0.0	1.0	0.0	0.0	1.0	1.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

```

# add the dataset to the fun_control
fun_control.update({"data": train_df, # full dataset,
                  "train": train,
                  "test": test,
                  "n_samples": n_samples,
                  "target_column": target_column})

```

17.4 Step 4: Specification of the Preprocessing Model

Data preprocessing can be very simple, e.g., you can ignore it. Then you would choose the `prep_model` "None":

```

prep_model = None
fun_control.update({"prep_model": prep_model})

```

A default approach for numerical data is the `StandardScaler` (mean 0, variance 1). This can be selected as follows:

```
# prep_model = StandardScaler()
# fun_control.update({"prep_model": prep_model})
```

Even more complicated pre-processing steps are possible, e.g., the following pipeline:

```
# categorical_columns = []
# one_hot_encoder = OneHotEncoder(handle_unknown="ignore", sparse_output=False)
# prep_model = ColumnTransformer(
#     transformers=[
#         ("categorical", one_hot_encoder, categorical_columns),
#     ],
#     remainder=StandardScaler(),
# )
```

17.5 Step 5: Select Model (algorithm) and `core_model_hyper_dict`

The selection of the algorithm (ML model) that should be tuned is done by specifying the its name from the `sklearn` implementation. For example, the SVC support vector machine classifier is selected as follows:

```
fun_control = add_core_model_to_fun_control(SVC, fun_control, SklearnHyperDict)
```

Other `core_models` are, e.g.,:

- `RidgeCV`
- `GradientBoostingRegressor`
- `ElasticNet`
- `RandomForestClassifier`
- `LogisticRegression`
- `KNeighborsClassifier`
- `RandomForestClassifier`
- `GradientBoostingClassifier`
- `HistGradientBoostingClassifier`

We will use the `RandomForestClassifier` classifier in this example.

```

from sklearn.linear_model import RidgeCV
from sklearn.ensemble import RandomForestClassifier
from sklearn.svm import SVC
from sklearn.linear_model import LogisticRegression
from sklearn.neighbors import KNeighborsClassifier
from sklearn.ensemble import GradientBoostingClassifier
from sklearn.ensemble import GradientBoostingRegressor
from sklearn.ensemble import HistGradientBoostingClassifier
from sklearn.linear_model import ElasticNet
from spotPython.hyperparameters.values import add_core_model_to_fun_control
from spotPython.data.sklearn_hyper_dict import SklearnHyperDict
from spotPython.fun.hypersklearn import HyperSklearn

```

```

# core_model = RidgeCV
# core_model = GradientBoostingRegressor
# core_model = ElasticNet
# core_model = RandomForestClassifier
core_model = SVC
# core_model = LogisticRegression
# core_model = KNeighborsClassifier
# core_model = GradientBoostingClassifier
# core_model = HistGradientBoostingClassifier
fun_control = add_core_model_to_fun_control(core_model=core_model,
                                          fun_control=fun_control,
                                          hyper_dict=SklearnHyperDict,
                                          filename=None)

```

Now `fun_control` has the information from the JSON file. The available hyperparameters are:

```

print(*fun_control["core_model_hyper_dict"].keys(), sep="\n")

```

```

C
kernel
degree
gamma
coef0
shrinking
probability
tol
cache_size

```

```
break_ties
```

17.6 Step 6: Modify hyper_dict Hyperparameters for the Selected Algorithm aka core_model

17.6.1 Modify hyperparameter of type numeric and integer (boolean)

Numeric and boolean values can be modified using the `modify_hyper_parameter_bounds` method. For example, to change the `tol` hyperparameter of the SVC model to the interval `[1e-3, 1e-2]`, the following code can be used:

```
fun_control = modify_hyper_parameter_bounds(fun_control, "tol", bounds=[1e-3, 1e-2])
```

```
from spotPython.hyperparameters.values import modify_hyper_parameter_bounds
fun_control = modify_hyper_parameter_bounds(fun_control, "probability", bounds=[1, 1])
```

17.6.2 Modify hyperparameter of type factor

spotPython provides functions for modifying the hyperparameters, their bounds and factors as well as for activating and de-activating hyperparameters without re-compilation of the Python source code. These functions were described in Section [14.6](#).

Factors can be modified with the `modify_hyper_parameter_levels` function. For example, to exclude the `sigmoid` kernel from the tuning, the `kernel` hyperparameter of the SVC model can be modified as follows:

```
fun_control = modify_hyper_parameter_levels(fun_control, "kernel", ["linear", "rbf"])
```

The new setting can be controlled via:

```
fun_control["core_model_hyper_dict"]["kernel"]
```

```
from spotPython.hyperparameters.values import modify_hyper_parameter_levels
fun_control = modify_hyper_parameter_levels(fun_control, "kernel", ["rbf"])
```

17.6.3 Optimizers

Optimizers are described in Section [14.6.1](#).

17.6.4 Selection of the Objective: Metric and Loss Functions

- Machine learning models are optimized with respect to a metric, for example, the accuracy function.
- Deep learning, e.g., neural networks are optimized with respect to a loss function, for example, the `cross_entropy` function and evaluated with respect to a metric, for example, the accuracy function.

17.7 Step 7: Selection of the Objective (Loss) Function

The loss function, that is usually used in deep learning for optimizing the weights of the net, is stored in the `fun_control` dictionary as `"loss_function"`.

17.7.1 Metric Function

There are two different types of metrics in `spotPython`:

1. `"metric_river"` is used for the river based evaluation via `eval_oml_iter_progressive`.
2. `"metric_sklearn"` is used for the sklearn based evaluation.

We will consider multi-class classification metrics, e.g., `mapk_score` and `top_k_accuracy_score`.

i Predict Probabilities

In this multi-class classification example the machine learning algorithm should return the probabilities of the specific classes (`"predict_proba"`) instead of the predicted values.

We set `"predict_proba"` to `True` in the `fun_control` dictionary.

17.7.1.1 The MAPK Metric

To select the MAPK metric, the following two entries can be added to the `fun_control` dictionary:

```
"metric_sklearn": mapk_score"
```

```
"metric_params": {"k": 3}.
```

17.7.1.2 Other Metrics

Alternatively, other metrics for multi-class classification can be used, e.g.,: * `top_k_accuracy_score` or * `roc_auc_score`

The metric `roc_auc_score` requires the parameter `"multi_class"`, e.g.,

```
"multi_class": "ovr".
```

This is set in the `fun_control` dictionary.

i Weights

spotPython performs a minimization, therefore, metrics that should be maximized have to be multiplied by -1. This is done by setting `"weights"` to -1.

- The complete setup for the metric in our example is:

```
from spotPython.utils.metrics import mapk_score
fun_control.update({
    "weights": -1,
    "metric_sklearn": mapk_score,
    "predict_proba": True,
    "metric_params": {"k": 3},
})
```

17.7.2 Evaluation on Hold-out Data

- The default method for computing the performance is `"eval_holdout"`.
- Alternatively, cross-validation can be used for every machine learning model.
- Specifically for RandomForests, the OOB-score can be used.

```
fun_control.update({
    "eval": "train_hold_out",
})
```

17.7.2.1 Cross Validation

Instead of using the OOB-score, the classical cross validation can be used. The number of folds is set by the key `"k_folds"`. For example, to use 5-fold cross validation, the key `"k_folds"` is set to 5. Uncomment the following line to use cross validation:

```
# fun_control.update({
#     "eval": "train_cv",
#     "k_folds": 10,
# })
```

17.8 Step 8: Calling the SPOT Function

17.8.1 Preparing the SPOT Call

- Get types and variable names as well as lower and upper bounds for the hyperparameters.

```
# extract the variable types, names, and bounds
from spotPython.hyperparameters.values import (get_bound_values,
        get_var_name,
        get_var_type,)
var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
fun_control.update({"var_type": var_type,
                    "var_name": var_name})
lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")

from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))
```

name	type	default	lower	upper	transform
C	float	1.0	0.1	10	None
kernel	factor	rbf	0	0	None
degree	int	3	3	3	None
gamma	factor	scale	0	1	None
coef0	float	0.0	0	0	None
shrinking	factor	0	0	1	None
probability	factor	0	1	1	None
tol	float	0.001	0.0001	0.01	None
cache_size	float	200.0	100	400	None
break_ties	factor	0	0	1	None

17.8.2 The Objective Function

The objective function is selected next. It implements an interface from `sklearn`'s training, validation, and testing methods to `spotPython`.

```
from spotPython.fun.hypersklearn import HyperSklearn
fun = HyperSklearn().fun_sklearn
```

17.8.3 Run the Spot Optimizer

- Run SPOT for approx. x mins (`max_time`).
- Note: the run takes longer, because the evaluation time of initial design (here: `init_size`, 20 points) is not considered.

```
from spotPython.hyperparameters.values import get_default_hyperparameters_as_array
hyper_dict=SklearnHyperDict().load()
X_start = get_default_hyperparameters_as_array(fun_control, hyper_dict)
X_start
```

```
array([[1.e+00, 2.e+00, 3.e+00, 0.e+00, 0.e+00, 0.e+00, 0.e+00, 1.e-03,
        2.e+02, 0.e+00]])
```

```
import numpy as np
from spotPython.spot import spot
from math import inf
spot_tuner = spot.Spot(fun=fun,
                      lower = lower,
                      upper = upper,
                      fun_evals = inf,
                      fun_repeats = 1,
                      max_time = MAX_TIME,
                      noise = False,
                      tolerance_x = np.sqrt(np.spacing(1)),
                      var_type = var_type,
                      var_name = var_name,
                      infill_criterion = "y",
                      n_points = 1,
                      seed=123,
                      log_level = 50,
                      show_models= False,
```

```

show_progress= True,
fun_control = fun_control,
design_control={"init_size": INIT_SIZE,
               "repeats": 1},
surrogate_control={"noise": True,
                   "cod_type": "norm",
                   "min_theta": -4,
                   "max_theta": 3,
                   "n_theta": len(var_name),
                   "model_fun_evals": 10_000,
                   "log_level": 50
                  })

spot_tuner.run(X_start=X_start)

```

```

spotPython tuning: -0.3646616541353383 [-----] 0.33%

spotPython tuning: -0.3646616541353383 [-----] 0.79%

spotPython tuning: -0.3646616541353383 [-----] 1.38%

spotPython tuning: -0.3646616541353383 [-----] 1.90%

spotPython tuning: -0.3646616541353383 [-----] 2.52%

spotPython tuning: -0.3646616541353383 [-----] 3.03%

spotPython tuning: -0.3646616541353383 [-----] 3.51%

spotPython tuning: -0.3646616541353383 [-----] 3.94%

spotPython tuning: -0.3646616541353383 [-----] 4.36%

spotPython tuning: -0.3734335839598997 [-----] 4.73%

spotPython tuning: -0.3734335839598997 [#-----] 5.23%

spotPython tuning: -0.37844611528822053 [#-----] 5.69%

```

spotPython tuning: -0.3847117794486215 [#-----] 7.44%

spotPython tuning: -0.3847117794486215 [#-----] 8.99%

spotPython tuning: -0.3847117794486215 [#-----] 10.67%

spotPython tuning: -0.3847117794486215 [#-----] 12.53%

spotPython tuning: -0.3847117794486215 [#-----] 14.16%

spotPython tuning: -0.3847117794486215 [##-----] 15.51%

spotPython tuning: -0.3847117794486215 [##-----] 16.90%

spotPython tuning: -0.3847117794486215 [##-----] 18.88%

spotPython tuning: -0.3847117794486215 [##-----] 20.54%

spotPython tuning: -0.3847117794486215 [##-----] 21.79%

spotPython tuning: -0.3847117794486215 [##-----] 23.10%

spotPython tuning: -0.38596491228070173 [##-----] 24.48%

spotPython tuning: -0.38596491228070173 [###-----] 25.91%

spotPython tuning: -0.38596491228070173 [###-----] 27.08%

spotPython tuning: -0.38596491228070173 [###-----] 28.13%

spotPython tuning: -0.38721804511278196 [###-----] 29.35%

spotPython tuning: -0.3922305764411027 [###-----] 30.74%

spotPython tuning: -0.3922305764411027 [###-----] 32.61%

spotPython tuning: -0.3922305764411027 [###-----] 34.94%

spotPython tuning: -0.3922305764411027 [####-----] 37.51%

spotPython tuning: -0.3922305764411027 [####-----] 40.38%

spotPython tuning: -0.3922305764411027 [####-----] 43.00%

spotPython tuning: -0.3922305764411027 [#####-----] 45.57%

spotPython tuning: -0.3922305764411027 [#####-----] 48.59%

spotPython tuning: -0.3922305764411027 [#####-----] 51.19%

spotPython tuning: -0.3922305764411027 [#####-----] 53.60%

spotPython tuning: -0.3922305764411027 [#####-----] 56.12%

spotPython tuning: -0.3922305764411027 [#####-----] 58.45%

spotPython tuning: -0.3922305764411027 [#####-----] 61.00%

spotPython tuning: -0.3922305764411027 [#####-----] 63.76%

spotPython tuning: -0.3922305764411027 [#####-----] 66.76%

spotPython tuning: -0.3922305764411027 [#####-----] 70.26%

spotPython tuning: -0.3922305764411027 [#####-----] 73.44%

spotPython tuning: -0.3922305764411027 [#####-----] 76.50%

spotPython tuning: -0.3922305764411027 [#####-----] 78.98%

spotPython tuning: -0.3922305764411027 [#####-----] 82.50%

spotPython tuning: -0.3922305764411027 [#####-----] 85.12%

spotPython tuning: -0.3922305764411027 [#####-----] 87.79%

```

spotPython tuning: -0.39348370927318294 [#####-] 90.18%

spotPython tuning: -0.39348370927318294 [#####-] 93.68%

spotPython tuning: -0.39348370927318294 [#####] 97.03%

spotPython tuning: -0.39348370927318294 [#####] 99.79%

spotPython tuning: -0.39348370927318294 [#####] 100.00% Done...

<spotPython.spot.spot.Spot at 0x28f2ec580>

```

17.9 Step 9: Tensorboard

The textual output shown in the console (or code cell) can be visualized with Tensorboard as described in Section 14.9, see also the description in the documentation: [Tensorboard](#).

17.10 Step 10: Results

After the hyperparameter tuning run is finished, the progress of the hyperparameter tuning can be visualized. The following code generates the progress plot from `?@fig-progress`.

```

spot_tuner.plot_progress(log_y=False,
    filename="./figures/" + experiment_name+"_progress.png")

```




Figure 17.1: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

- Print the results

```
print(gen_design_table(fun_control=fun_control,
                      spot=spot_tuner))
```

name	type	default	lower	upper	tuned	transform
C	float	1.0	0.1	10.0	4.016160671342672	None
kernel	factor	rbf	0.0	0.0	0.0	None
degree	int	3	3.0	3.0	3.0	None
gamma	factor	scale	0.0	1.0	1.0	None
coef0	float	0.0	0.0	0.0	0.0	None
shrinking	factor	0	0.0	1.0	1.0	None
probability	factor	0	1.0	1.0	1.0	None
tol	float	0.001	0.0001	0.01	0.002210461951108952	None
cache_size	float	200.0	100.0	400.0	156.83495628776507	None
break_ties	factor	0	0.0	1.0	0.0	None

17.10.1 Show variable importance

```
spot_tuner.plot_importance(threshold=0.025, filename="./figures/" + experiment_name+"_impo
```




Figure 17.2: Variable importance plot, threshold 0.025.

17.10.2 Get Default Hyperparameters

```
from spotPython.hyperparameters.values import get_default_values, transform_hyper_parameter_values
values_default = get_default_values(fun_control)
values_default = transform_hyper_parameter_values(fun_control=fun_control, hyper_parameter=hyper_parameter)
values_default
```

```
{'C': 1.0,
 'kernel': 'rbf',
 'degree': 3,
 'gamma': 'scale',
 'coef0': 0.0,
 'shrinking': 0,
 'probability': 0,
 'tol': 0.001,
 'cache_size': 200.0,
 'break_ties': 0}
```

```
from sklearn.pipeline import make_pipeline
model_default = make_pipeline(fun_control["prep_model"], fun_control["core_model"](**values_default))
model_default
```

```
Pipeline(steps=[('nonetype', None),
                 ('svc',
                  SVC(break_ties=0, cache_size=200.0, probability=0,
                      shrinking=0))])
```

i Note

- Default value for “probability” is False, but we need it to be True for the metric “mapk_score”.

```
values_default.update({"probability": 1})
```

17.10.3 Get SPOT Results

```
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))
print(X)
```

```
[[4.01616067e+00 0.00000000e+00 3.00000000e+00 1.00000000e+00
 0.00000000e+00 1.00000000e+00 1.00000000e+00 2.21046195e-03
 1.56834956e+02 0.00000000e+00]]
```

```
from spotPython.hyperparameters.values import assign_values, return_conf_list_from_var_dict
v_dict = assign_values(X, fun_control["var_name"])
return_conf_list_from_var_dict(var_dict=v_dict, fun_control=fun_control)
```

```
[{'C': 4.016160671342672,
 'kernel': 'rbf',
 'degree': 3,
 'gamma': 'auto',
 'coef0': 0.0,
 'shrinking': 1,
 'probability': 1,
 'tol': 0.002210461951108952,
 'cache_size': 156.83495628776507,
 'break_ties': 0}]
```

```
from spotPython.hyperparameters.values import get_one_sklearn_model_from_X
model_spot = get_one_sklearn_model_from_X(X, fun_control)
model_spot
```

```
SVC(C=4.016160671342672, break_ties=0, cache_size=156.83495628776507,
    gamma='auto', probability=1, shrinking=1, tol=0.002210461951108952)
```

17.10.4 Evaluate SPOT Results

- Fetch the data.

```
from spotPython.utils.convert import get_Xy_from_df
X_train, y_train = get_Xy_from_df(fun_control["train"], fun_control["target_column"])
X_test, y_test = get_Xy_from_df(fun_control["test"], fun_control["target_column"])
```

```
X_test.shape, y_test.shape
```

```
((177, 64), (177,))
```

- Fit the model with the tuned hyperparameters. This gives one result:

```
model_spot.fit(X_train, y_train)
y_pred = model_spot.predict_proba(X_test)
res = mapk_score(y_true=y_test, y_pred=y_pred, k=3)
res
```

```
0.37947269303201503
```

```
def repeated_eval(n, model):
    res_values = []
    for i in range(n):
        model.fit(X_train, y_train)
        y_pred = model.predict_proba(X_test)
        res = mapk_score(y_true=y_test, y_pred=y_pred, k=3)
        res_values.append(res)
    mean_res = np.mean(res_values)
    print(f"mean_res: {mean_res}")
    std_res = np.std(res_values)
    print(f"std_res: {std_res}")
    min_res = np.min(res_values)
    print(f"min_res: {min_res}")
    max_res = np.max(res_values)
    print(f"max_res: {max_res}")
    median_res = np.median(res_values)
    print(f"median_res: {median_res}")
    return mean_res, std_res, min_res, max_res, median_res
```

17.10.5 Handling Non-deterministic Results

- Because the model is non-deterministic, we perform $n = 30$ runs and calculate the mean and standard deviation of the performance metric.

```
_ = repeated_eval(30, model_spot)
```

```
mean_res: 0.37790332705586943
std_res: 0.003721747359141646
min_res: 0.3700564971751412
max_res: 0.3851224105461393
median_res: 0.378060263653484
```

17.10.6 Evaluation of the Default Hyperparameters

```
model_default["svc"].probability = True
model_default.fit(X_train, y_train)["svc"]
```

```
SVC(break_ties=0, cache_size=200.0, probability=True, shrinking=0)
```

- One evaluation of the default hyperparameters is performed on the hold-out test set.

```
y_pred = model_default.predict_proba(X_test)
mapk_score(y_true=y_test, y_pred=y_pred, k=3)
```

```
0.3794726930320151
```


Since one single evaluation is not meaningful, we perform, similar to the evaluation of the SPOT results, $n = 30$ runs of the default setting and calculate the mean and standard deviation of the performance metric.

```
_ = repeated_eval(30, model_default)
```


```
mean_res: 0.3841180163214062
std_res: 0.004389289668904793
min_res: 0.37664783427495285
max_res: 0.39453860640301314
median_res: 0.38465160075329563
```

17.10.7 Plot: Compare Predictions

```
from spotPython.plot.validation import plot_confusion_matrix
plot_confusion_matrix(model_default, fun_control, title = "Default")
```



```
plot_confusion_matrix(model_spot, fun_control, title="SPOT")
```



```
min(spot_tuner.y), max(spot_tuner.y)
```

```
(-0.39348370927318294, -0.3345864661654135)
```

17.10.8 Cross-validated Evaluations

```
from spotPython.sklearn.traintest import evaluate_cv
fun_control.update({
    "eval": "train_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)
```

```
(0.35188679245283017, None)
```

```

fun_control.update({
    "eval": "test_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)

```

(0.34008714596949885, None)

- This is the evaluation that will be used in the comparison:

```

fun_control.update({
    "eval": "data_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)

```

(0.3612709590878605, None)



17.10.9 Detailed Hyperparameter Plots





```

filename = "./figures/" + experiment_name
spot_tuner.plot_important_hyperparameter_contour(filename=filename)

```

C: 1.1840993181254122
gamma: 100.0
tol: 10.866921286853268





17.10.10 Parallel Coordinates Plot

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Unable to display output for mime type(s): text/html

17.10.11 Plot all Combinations of Hyperparameters

- Warning: this may take a while.

```
PLOT_ALL = False
if PLOT_ALL:
    n = spot_tuner.k
    for i in range(n-1):
        for j in range(i+1, n):
            spot_tuner.plot_contour(i=i, j=j, min_z=min_z, max_z = max_z)
```

18 HPT: sklearn KNN Classifier VBDP Data

This document refers to the following software versions:

- python: 3.10.10

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

spotPython can be installed via pip. Alternatively, the source code can be downloaded from gitHub: <https://github.com/sequential-parameter-optimization/spotPython>.

```
!pip install spotPython
```

- Uncomment the following lines if you want to for (re-)installation the latest version of spotPython from gitHub.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

18.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time and the initial design size.

```
MAX_TIME = 1
INIT_SIZE = 5
ORIGINAL = False
```

```

import os
import copy
import socket
from datetime import datetime
from dateutil.tz import tzlocal
start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '19-knn-sklearn' + "_" + HOSTNAME + "_" + str(MAX_TIME) + "min_" + str(I
experiment_name = experiment_name.replace(':', '-')
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')

```

19-knn-sklearn_bartz09_1min_5init_2023-06-19_04-19-17

```

import warnings
warnings.filterwarnings("ignore")

```

18.2 Step 2: Initialization of the Empty `fun_control` Dictionary

 Caution: Tensorboard does not work under Windows

- Since tensorboard does not work under Windows, we recommend setting the parameter `tensorboard_path` to `None` if you are working under Windows.

```

from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="classification",
    tensorboard_path="runs/16_spot_hpt_sklearn_classification")

```

18.2.1 Load Data: Classification VBDP

```

import pandas as pd
if ORIGINAL == True:
    train_df = pd.read_csv('./data/VBDP/trainn.csv')
    test_df = pd.read_csv('./data/VBDP/testt.csv')
else:
    train_df = pd.read_csv('./data/VBDP/train.csv')

```

```

# remove the id column
train_df = train_df.drop(columns=['id'])

from sklearn.preprocessing import OrdinalEncoder
n_samples = train_df.shape[0]
n_features = train_df.shape[1] - 1
target_column = "prognosis"
# Encoder our prognosis labels as integers for easier decoding later
enc = OrdinalEncoder()
train_df[target_column] = enc.fit_transform(train_df[[target_column]])
train_df.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
print(train_df.shape)
train_df.head()

```

(707, 65)

	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	...	x56	x57	x58	x59	x60	x61	x62	x63
0	1.0	1.0	0.0	1.0	1.0	1.0	1.0	0.0	1.0	1.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	1.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0	1.0	1.0	1.0	0.0	1.0	1.0	1.0	1.0	1.0	...	1.0	1.0	1.0	1.0	1.0	0.0	1.0	1.0
3	0.0	0.0	1.0	1.0	1.0	1.0	0.0	1.0	0.0	1.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	0.0	...	0.0	1.0	0.0	0.0	1.0	1.0	1.0	0.0

The full data set `train_df` 64 features. The target column is labeled as `prognosis`.

18.2.2 Holdout Train and Test Data

We split out a hold-out test set (25% of the data) so we can calculate an example MAP@K

```

import numpy as np
from sklearn.model_selection import train_test_split
X_train, X_test, y_train, y_test = train_test_split(train_df.drop(target_column, axis=1),
                                                    random_state=42,
                                                    test_size=0.25,
                                                    stratify=train_df[target_column])
train = pd.DataFrame(np.hstack((X_train, np.array(y_train).reshape(-1, 1))))
test = pd.DataFrame(np.hstack((X_test, np.array(y_test).reshape(-1, 1))))
train.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]

```

```

test.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
print(train.shape)
print(test.shape)
train.head()

```

(530, 65)

(177, 65)

	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	...	x56	x57	x58	x59	x60	x61	x62	x63	
0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	0.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.0	0.0	0.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	1.0	1.0	1.0	0.0	0.0
3	1.0	1.0	0.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	0.0	0.0	0.0	1.0	0.0	0.0	1.0	1.0	0.0	0.0	...	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

```

# add the dataset to the fun_control
fun_control.update({"data": train_df, # full dataset,
                  "train": train,
                  "test": test,
                  "n_samples": n_samples,
                  "target_column": target_column})

```

18.3 Step 4: Specification of the Preprocessing Model

Data preprocessing can be very simple, e.g., you can ignore it. Then you would choose the `prep_model` "None":

```

prep_model = None
fun_control.update({"prep_model": prep_model})

```

A default approach for numerical data is the `StandardScaler` (mean 0, variance 1). This can be selected as follows:

```

# prep_model = StandardScaler()
# fun_control.update({"prep_model": prep_model})

```

Even more complicated pre-processing steps are possible, e.g., the following pipeline:

```

# categorical_columns = []
# one_hot_encoder = OneHotEncoder(handle_unknown="ignore", sparse_output=False)
# prep_model = ColumnTransformer(
#     transformers=[
#         ("categorical", one_hot_encoder, categorical_columns),
#     ],
#     remainder=StandardScaler(),
# )

```

18.4 Step 5: Select Model (algorithm) and core_model_hyper_dict

The selection of the algorithm (ML model) that should be tuned is done by specifying the its name from the `sklearn` implementation. For example, the SVC support vector machine classifier is selected as follows:

```
fun_control = add_core_model_to_fun_control(SVC, fun_control, SklearnHyperDict)
```

Other core_models are, e.g.,:

- RidgeCV
- GradientBoostingRegressor
- ElasticNet
- RandomForestClassifier
- LogisticRegression
- KNeighborsClassifier
- RandomForestClassifier
- GradientBoostingClassifier
- HistGradientBoostingClassifier

We will use the `RandomForestClassifier` classifier in this example.

```

from sklearn.linear_model import RidgeCV
from sklearn.ensemble import RandomForestClassifier
from sklearn.svm import SVC
from sklearn.linear_model import LogisticRegression
from sklearn.neighbors import KNeighborsClassifier
from sklearn.ensemble import GradientBoostingClassifier
from sklearn.ensemble import GradientBoostingRegressor
from sklearn.ensemble import HistGradientBoostingClassifier
from sklearn.linear_model import ElasticNet

```

```

from spotPython.hyperparameters.values import add_core_model_to_fun_control
from spotPython.data.sklearn_hyper_dict import SklearnHyperDict
from spotPython.fun.hypersklearn import HyperSklearn

# core_model = RidgeCV
# core_model = GradientBoostingRegressor
# core_model = ElasticNet
# core_model = RandomForestClassifier
core_model = KNeighborsClassifier
# core_model = LogisticRegression
# core_model = KNeighborsClassifier
# core_model = GradientBoostingClassifier
# core_model = HistGradientBoostingClassifier
fun_control = add_core_model_to_fun_control(core_model=core_model,
                                          fun_control=fun_control,
                                          hyper_dict=SklearnHyperDict,
                                          filename=None)

```

Now `fun_control` has the information from the JSON file. The available hyperparameters are:

```
print(*fun_control["core_model_hyper_dict"].keys(), sep="\n")
```

```

n_neighbors
weights
algorithm
leaf_size
p

```

18.5 Step 6: Modify `hyper_dict` Hyperparameters for the Selected Algorithm aka `core_model`

18.5.1 Modify hyperparameter of type numeric and integer (boolean)

Numeric and boolean values can be modified using the `modify_hyper_parameter_bounds` method. For example, to change the `tol` hyperparameter of the `SVC` model to the interval $[1e-3, 1e-2]$, the following code can be used:

```
fun_control = modify_hyper_parameter_bounds(fun_control, "tol", bounds=[1e-3, 1e-2])
```

```
# from spotPython.hyperparameters.values import modify_hyper_parameter_bounds
# fun_control = modify_hyper_parameter_bounds(fun_control, "probability", bounds=[1, 1])
```

18.5.2 Modify hyperparameter of type factor

spotPython provides functions for modifying the hyperparameters, their bounds and factors as well as for activating and de-activating hyperparameters without re-compilation of the Python source code. These functions were described in Section 14.6.

Factors can be modified with the `modify_hyper_parameter_levels` function. For example, to exclude the `sigmoid` kernel from the tuning, the `kernel` hyperparameter of the SVC model can be modified as follows:

```
fun_control = modify_hyper_parameter_levels(fun_control, "kernel", ["linear",
"rbf"])
```

The new setting can be controlled via:

```
fun_control["core_model_hyper_dict"]["kernel"]
```

```
# from spotPython.hyperparameters.values import modify_hyper_parameter_levels
# fun_control = modify_hyper_parameter_levels(fun_control, "kernel", ["rbf"])
```

18.5.3 Optimizers

Optimizers are described in Section 14.6.1.

18.5.4 Selection of the Objective: Metric and Loss Functions

- Machine learning models are optimized with respect to a metric, for example, the accuracy function.
- Deep learning, e.g., neural networks are optimized with respect to a loss function, for example, the `cross_entropy` function and evaluated with respect to a metric, for example, the accuracy function.

18.6 Step 7: Selection of the Objective (Loss) Function

The loss function, that is usually used in deep learning for optimizing the weights of the net, is stored in the `fun_control` dictionary as `"loss_function"`.

18.6.1 Metric Function

There are two different types of metrics in `spotPython`:

1. "metric_river" is used for the river based evaluation via `eval_oml_iter_progressive`.
2. "metric_sklearn" is used for the sklearn based evaluation.

We will consider multi-class classification metrics, e.g., `mapk_score` and `top_k_accuracy_score`.

i Predict Probabilities

In this multi-class classification example the machine learning algorithm should return the probabilities of the specific classes ("`predict_proba`") instead of the predicted values.

We set "`predict_proba`" to `True` in the `fun_control` dictionary.

18.6.1.1 The MAPK Metric

To select the MAPK metric, the following two entries can be added to the `fun_control` dictionary:

```
"metric_sklearn": mapk_score"
```

```
"metric_params": {"k": 3}.
```

18.6.1.2 Other Metrics

Alternatively, other metrics for multi-class classification can be used, e.g., * `top_k_accuracy_score` or * `roc_auc_score`

The metric `roc_auc_score` requires the parameter "`multi_class`", e.g.,

```
"multi_class": "ovr".
```

This is set in the `fun_control` dictionary.

i Weights

`spotPython` performs a minimization, therefore, metrics that should be maximized have to be multiplied by -1. This is done by setting "`weights`" to -1.

- The complete setup for the metric in our example is:

```

from spotPython.utils.metrics import mapk_score
fun_control.update({
    "weights": -1,
    "metric_sklearn": mapk_score,
    "predict_proba": True,
    "metric_params": {"k": 3},
})

```

18.6.2 Evaluation on Hold-out Data

- The default method for computing the performance is "eval_holdout".
- Alternatively, cross-validation can be used for every machine learning model.
- Specifically for RandomForests, the OOB-score can be used.

```

fun_control.update({
    "eval": "train_hold_out",
})

```

18.6.2.1 Cross Validation

Instead of using the OOB-score, the classical cross validation can be used. The number of folds is set by the key "k_folds". For example, to use 5-fold cross validation, the key "k_folds" is set to 5. Uncomment the following line to use cross validation:

```

# fun_control.update({
#     "eval": "train_cv",
#     "k_folds": 10,
# })

```

18.7 Step 8: Calling the SPOT Function

18.7.1 Preparing the SPOT Call

- Get types and variable names as well as lower and upper bounds for the hyperparameters.

```

# extract the variable types, names, and bounds
from spotPython.hyperparameters.values import (get_bound_values,
    get_var_name,

```

```

    get_var_type,)
var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
fun_control.update({"var_type": var_type,
                   "var_name": var_name})
lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")

from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))

```

name	type	default	lower	upper	transform
n_neighbors	int	2	1	7	transform_power_2_int
weights	factor	uniform	0	1	None
algorithm	factor	auto	0	3	None
leaf_size	int	5	2	7	transform_power_2_int
p	int	2	1	2	None

18.7.2 The Objective Function

The objective function is selected next. It implements an interface from sklearn's training, validation, and testing methods to spotPython.

```

from spotPython.fun.hyper sklearn import HyperSklearn
fun = HyperSklearn().fun_sklearn

```

18.7.3 Run the Spot Optimizer

- Run SPOT for approx. x mins (max_time).
- Note: the run takes longer, because the evaluation time of initial design (here: initi_size, 20 points) is not considered.

```

from spotPython.hyperparameters.values import get_default_hyperparameters_as_array
hyper_dict=SklearnHyperDict().load()
X_start = get_default_hyperparameters_as_array(fun_control, hyper_dict)
X_start

```

```
array([[2, 0, 0, 5, 2]])
```

```

import numpy as np
from spotPython.spot import spot
from math import inf
spot_tuner = spot.Spot(fun=fun,
                      lower = lower,
                      upper = upper,
                      fun_evals = inf,
                      fun_repeats = 1,
                      max_time = MAX_TIME,
                      noise = False,
                      tolerance_x = np.sqrt(np.spacing(1)),
                      var_type = var_type,
                      var_name = var_name,
                      infill_criterion = "y",
                      n_points = 1,
                      seed=123,
                      log_level = 50,
                      show_models= False,
                      show_progress= True,
                      fun_control = fun_control,
                      design_control={"init_size": INIT_SIZE,
                                     "repeats": 1},
                      surrogate_control={"noise": True,
                                       "cod_type": "norm",
                                       "min_theta": -4,
                                       "max_theta": 3,
                                       "n_theta": len(var_name),
                                       "model_fun_evals": 10_000,
                                       "log_level": 50
                                       })

spot_tuner.run(X_start=X_start)

```

spotPython tuning: -0.3107769423558897 [-----] 0.24%

spotPython tuning: -0.3107769423558897 [-----] 0.51%

spotPython tuning: -0.3107769423558897 [-----] 0.77%

spotPython tuning: -0.3107769423558897 [-----] 1.02%

spotPython tuning: -0.3107769423558897 [-----] 1.27%

spotPython tuning: -0.3107769423558897 [-----] 1.59%

spotPython tuning: -0.3107769423558897 [-----] 1.97%

spotPython tuning: -0.3107769423558897 [-----] 2.30%

spotPython tuning: -0.3107769423558897 [-----] 2.62%

spotPython tuning: -0.3107769423558897 [-----] 2.91%

spotPython tuning: -0.3107769423558897 [-----] 3.20%

spotPython tuning: -0.3107769423558897 [-----] 4.18%

spotPython tuning: -0.3107769423558897 [#-----] 5.15%

spotPython tuning: -0.3107769423558897 [#-----] 6.33%

spotPython tuning: -0.3107769423558897 [#-----] 7.52%

spotPython tuning: -0.3107769423558897 [#-----] 8.76%

spotPython tuning: -0.3107769423558897 [#-----] 10.20%

spotPython tuning: -0.3107769423558897 [#-----] 11.32%

spotPython tuning: -0.3107769423558897 [#-----] 13.27%

spotPython tuning: -0.3107769423558897 [#-----] 14.52%

spotPython tuning: -0.3107769423558897 [##-----] 15.56%

spotPython tuning: -0.3107769423558897 [##-----] 16.43%

spotPython tuning: -0.3107769423558897 [##-----] 17.43%

spotPython tuning: -0.3107769423558897 [##-----] 18.65%

spotPython tuning: -0.3107769423558897 [##-----] 19.76%

spotPython tuning: -0.3107769423558897 [##-----] 20.98%

spotPython tuning: -0.3107769423558897 [##-----] 22.40%

spotPython tuning: -0.3107769423558897 [##-----] 23.99%

spotPython tuning: -0.3107769423558897 [###-----] 25.74%

spotPython tuning: -0.3107769423558897 [###-----] 27.68%

spotPython tuning: -0.3107769423558897 [###-----] 29.34%

spotPython tuning: -0.3107769423558897 [###-----] 31.78%

spotPython tuning: -0.3107769423558897 [###-----] 33.87%

spotPython tuning: -0.3107769423558897 [####-----] 36.42%

spotPython tuning: -0.3107769423558897 [####-----] 40.27%

spotPython tuning: -0.3107769423558897 [####-----] 43.10%

spotPython tuning: -0.3107769423558897 [#####-----] 45.80%

spotPython tuning: -0.3107769423558897 [#####-----] 48.82%

spotPython tuning: -0.3107769423558897 [#####-----] 51.67%

spotPython tuning: -0.3107769423558897 [#####-----] 54.63%

spotPython tuning: -0.3107769423558897 [#####-----] 57.81%

spotPython tuning: -0.3107769423558897 [#####-----] 60.86%

spotPython tuning: -0.3107769423558897 [#####-----] 64.07%

```
spotPython tuning: -0.3107769423558897 [#####---] 67.24%
spotPython tuning: -0.3107769423558897 [#####---] 70.04%
spotPython tuning: -0.3107769423558897 [#####---] 73.05%
spotPython tuning: -0.3107769423558897 [#####--] 77.21%
spotPython tuning: -0.3107769423558897 [#####--] 80.49%
spotPython tuning: -0.3107769423558897 [#####--] 84.07%
spotPython tuning: -0.3107769423558897 [#####-] 87.87%
spotPython tuning: -0.3107769423558897 [#####-] 91.85%
spotPython tuning: -0.3107769423558897 [#####] 95.14%
spotPython tuning: -0.3107769423558897 [#####] 97.58%
spotPython tuning: -0.3107769423558897 [#####] 100.00% Done...

<spotPython.spot.spot.Spot at 0x14f8e8790>
```

18.8 Step 9: Tensorboard

The textual output shown in the console (or code cell) can be visualized with Tensorboard as described in Section 14.9, see also the description in the documentation: [Tensorboard](#).

18.9 Step 10: Results

After the hyperparameter tuning run is finished, the progress of the hyperparameter tuning can be visualized. The following code generates the progress plot from `?@fig-progress`.

```
spot_tuner.plot_progress(log_y=False,  
                        filename="./figures/" + experiment_name+"_progress.png")
```




Figure 18.1: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

- Print the results

```
print(gen_design_table(fun_control=fun_control,  
                      spot=spot_tuner))
```

name	type	default	lower	upper	tuned	transform
n_neighbors	int	2	1	7	4.0	transform_power_2_int
weights	factor	uniform	0	1	1.0	None
algorithm	factor	auto	0	3	2.0	None
leaf_size	int	5	2	7	6.0	transform_power_2_int
p	int	2	1	2	1.0	None

18.9.1 Show variable importance

```
spot_tuner.plot_importance(threshold=0.025, filename="./figures/" + experiment_name+"_impo
```




Figure 18.2: Variable importance plot, threshold 0.025.

18.9.2 Get Default Hyperparameters

```
from spotPython.hyperparameters.values import get_default_values, transform_hyper_parameter_values
values_default = get_default_values(fun_control)
values_default = transform_hyper_parameter_values(fun_control=fun_control, hyper_parameter_values_default=values_default)
```

```
{'n_neighbors': 4,
 'weights': 'uniform',
 'algorithm': 'auto',
 'leaf_size': 32,
 'p': 2}
```

```
from sklearn.pipeline import make_pipeline
model_default = make_pipeline(fun_control["prep_model"], fun_control["core_model"](**value_default))
model_default
```

```
Pipeline(steps=[('nonetype', None),
                 ('kneighborsclassifier',
                  KNeighborsClassifier(leaf_size=32, n_neighbors=4))])
```

18.9.3 Get SPOT Results

```
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))
print(X)
```

```
[[4. 1. 2. 6. 1.]]
```

```
from spotPython.hyperparameters.values import assign_values, return_conf_list_from_var_dict
v_dict = assign_values(X, fun_control["var_name"])
return_conf_list_from_var_dict(var_dict=v_dict, fun_control=fun_control)
```

```
[{'n_neighbors': 16,
  'weights': 'distance',
  'algorithm': 'kd_tree',
  'leaf_size': 64,
  'p': 1}]
```

```
from spotPython.hyperparameters.values import get_one_sklearn_model_from_X
model_spot = get_one_sklearn_model_from_X(X, fun_control)
model_spot
```

```
KNeighborsClassifier(algorithm='kd_tree', leaf_size=64, n_neighbors=16, p=1,
                    weights='distance')
```

18.9.4 Evaluate SPOT Results

- Fetch the data.

```
from spotPython.utils.convert import get_Xy_from_df
X_train, y_train = get_Xy_from_df(fun_control["train"], fun_control["target_column"])
X_test, y_test = get_Xy_from_df(fun_control["test"], fun_control["target_column"])
X_test.shape, y_test.shape
```

```
((177, 64), (177,))
```

- Fit the model with the tuned hyperparameters. This gives one result:

```
model_spot.fit(X_train, y_train)
y_pred = model_spot.predict_proba(X_test)
res = mapk_score(y_true=y_test, y_pred=y_pred, k=3)
res
```

0.3267419962335216

```
def repeated_eval(n, model):
    res_values = []
    for i in range(n):
        model.fit(X_train, y_train)
        y_pred = model.predict_proba(X_test)
        res = mapk_score(y_true=y_test, y_pred=y_pred, k=3)
        res_values.append(res)
    mean_res = np.mean(res_values)
    print(f"mean_res: {mean_res}")
    std_res = np.std(res_values)
    print(f"std_res: {std_res}")
    min_res = np.min(res_values)
    print(f"min_res: {min_res}")
    max_res = np.max(res_values)
    print(f"max_res: {max_res}")
    median_res = np.median(res_values)
    print(f"median_res: {median_res}")
    return mean_res, std_res, min_res, max_res, median_res
```

18.9.5 Handling Non-deterministic Results

- Because the model is non-deterministic, we perform $n = 30$ runs and calculate the mean and standard deviation of the performance metric.

```
_ = repeated_eval(30, model_spot)
```

```
mean_res: 0.3267419962335218
std_res: 1.6653345369377348e-16
min_res: 0.3267419962335216
max_res: 0.3267419962335216
median_res: 0.3267419962335216
```

18.9.6 Evolution of the Default Hyperparameters

```
model_default.fit(X_train, y_train)["kneighborsclassifier"]
```

```
KNeighborsClassifier(leaf_size=32, n_neighbors=4)
```

- One evaluation of the default hyperparameters is performed on the hold-out test set.

```
y_pred = model_default.predict_proba(X_test)  
mapk_score(y_true=y_test, y_pred=y_pred, k=3)
```

```
0.2768361581920904
```


Since one single evaluation is not meaningful, we perform, similar to the evaluation of the SPOT results, $n = 30$ runs of the default setting and calculate the mean and standard deviation of the performance metric.

```
_ = repeated_eval(30, model_default)
```


```
mean_res: 0.2768361581920903  
std_res: 1.1102230246251565e-16  
min_res: 0.2768361581920904  
max_res: 0.2768361581920904  
median_res: 0.2768361581920904
```

18.9.7 Plot: Compare Predictions

```
from spotPython.plot.validation import plot_confusion_matrix  
plot_confusion_matrix(model_default, fun_control, title = "Default")
```



```
plot_confusion_matrix(model_spot, fun_control, title="SPOT")
```



```
min(spot_tuner.y), max(spot_tuner.y)
```

```
(-0.3107769423558897, -0.23558897243107768)
```

18.9.8 Cross-validated Evaluations

```
from spotPython.sklearn.traintest import evaluate_cv
fun_control.update({
    "eval": "train_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)
```

```
(0.3157232704402516, None)
```

```

fun_control.update({
    "eval": "test_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)

```

(0.2832788671023965, None)

- This is the evaluation that will be used in the comparison:

```

fun_control.update({
    "eval": "data_cv",
    "k_folds": 10,
})
evaluate_cv(model=model_spot, fun_control=fun_control, verbose=0)

```

(0.3061904761904762, None)


18.9.9 Detailed Hyperparameter Plots

```

filename = "./figures/" + experiment_name
spot_tuner.plot_important_hyperparameter_contour(filename=filename)

```

n_neighbors: 7.659298853276286
weights: 100.0



18.9.10 Parallel Coordinates Plot

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Unable to display output for mime type(s): text/html

18.9.11 Plot all Combinations of Hyperparameters

- Warning: this may take a while.

```
PLOT_ALL = False
if PLOT_ALL:
    n = spot_tuner.k
    for i in range(n-1):
        for j in range(i+1, n):
            spot_tuner.plot_contour(i=i, j=j, min_z=min_z, max_z = max_z)
```

19 HPT PyTorch: Regression

In this tutorial, we will show how `spotPython` can be integrated into the PyTorch training workflow for regression tasks.

This document refers to the following software versions:

- python: 3.10.10
- torch: 2.0.1
- torchvision: 0.15.0

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

`spotPython` can be installed via `pip`. Alternatively, the source code can be downloaded from `gitHub`: <https://github.com/sequential-parameter-optimization/spotPython>.

```
!pip install spotPython
```

- Uncomment the following lines if you want to for (re-)installation the latest version of `spotPython` from `gitHub`.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

19.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time, initial design size and the device that is used.

 Caution: Run time and initial design size should be increased for real experiments

- MAX_TIME is set to one minute for demonstration purposes. For real experiments, this should be increased to at least 1 hour.
- INIT_SIZE is set to 5 for demonstration purposes. For real experiments, this should be increased to at least 10.

 Note: Device selection

- The device can be selected by setting the variable DEVICE.
- Since we are using a simple neural net, the setting "cpu" is preferred (on Mac).
- If you have a GPU, you can use "cuda:0" instead.
- If DEVICE is set to None, spotPython will automatically select the device.
 - This might result in "mps" on Macs, which is not the best choice for simple neural nets.

```
MAX_TIME = 1
INIT_SIZE = 5
DEVICE = "cpu" # "cuda:0"
```

```
from spotPython.utils.device import getDevice
DEVICE = getDevice(DEVICE)
print(DEVICE)
```

cpu

```
import os
import copy
import socket
from datetime import datetime
from dateutil.tz import tzlocal
start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '24-torch' + "_" + HOSTNAME + "_" + str(MAX_TIME) + "min_" + str(INIT_SIZE)
experiment_name = experiment_name.replace(':', '-')
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')
```

24-torch_bartz09_1min_5init_2023-06-19_04-24-41

19.2 Step 2: Initialization of the fun_control Dictionary

 Caution: Tensorboard does not work under Windows

- Since tensorboard does not work under Windows, we recommend setting the parameter `tensorboard_path` to `None` if you are working under Windows.

spotPython uses a Python dictionary for storing the information required for the hyperparameter tuning process, which was described in Section 14.2.

```
from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="regression",
    tensorboard_path="runs/24_spot_torch_regression",
    device=DEVICE)
```

19.3 Step 3: PyTorch Data Loading

```
# Create dataset
import pandas as pd
import numpy as np
from sklearn import datasets as sklearn_datasets
from sklearn.preprocessing import MinMaxScaler
from sklearn.model_selection import train_test_split
X, y = sklearn_datasets.make_regression(
    n_samples=1000, n_features=10, noise=1, random_state=123)
y = y.reshape(-1, 1)

# Normalize the data
X_scaler = MinMaxScaler()
X_scaled = X_scaler.fit_transform(X)
y_scaler = MinMaxScaler()
y_scaled = y_scaler.fit_transform(y)

# combine the features and target into a single dataframe named train_df
train_df = pd.DataFrame(np.hstack((X_scaled, y_scaled)))

target_column = "y"
n_samples = train_df.shape[0]
n_features = train_df.shape[1] - 1
```

```

train_df.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
X_train, X_test, y_train, y_test = train_test_split(train_df.drop(target_column,
axis=1),
train_df[target_column],
random_state=42,
test_size=0.25)
trainset = pd.DataFrame(np.hstack((X_train, np.array(y_train).reshape(-1, 1))))
testset = pd.DataFrame(np.hstack((X_test, np.array(y_test).reshape(-1, 1))))
trainset.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
testset.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
print(train_df.shape)
print(trainset.shape)
print(testset.shape)

```

(1000, 11)

(750, 11)

(250, 11)

```

import torch
from spotPython.torch.dataframedataset import DataFrameDataset
dtype_x = torch.float32
dtype_y = torch.float32
train_df = DataFrameDataset(train_df, target_column=target_column,
dtype_x=dtype_x, dtype_y=dtype_y)
train = DataFrameDataset(trainset, target_column=target_column,
dtype_x=dtype_x, dtype_y=dtype_y)
test = DataFrameDataset(testset, target_column=target_column,
dtype_x=dtype_x, dtype_y=dtype_y)
n_samples = len(train)

```

- Now we can test the data loading:

```

from spotPython.torch.traintest import create_train_val_data_loaders
trainloader, testloader = create_train_val_data_loaders(train, 2, True, 0)
for i, data in enumerate(trainloader, 0):
    inputs, labels = data
    print(inputs.shape)
    print(labels.shape)
    print(inputs)
    print(labels)
    break

```

```

torch.Size([2, 10])
torch.Size([2])
tensor([[0.2893, 0.2240, 0.6973, 0.5556, 0.5022, 0.3423, 0.4093, 0.6143, 0.3057,
         0.6452],
        [0.4557, 0.3300, 0.7537, 0.6591, 0.2341, 0.3079, 0.6961, 0.1191, 0.5170,
         0.4020]])
tensor([0.2696, 0.1006])

```

- Since this works fine, we can add the data loading to the `fun_control` dictionary:

```

# add the dataset to the fun_control
fun_control.update({"data": train_df, # full dataset,
                  "train": train,
                  "test": test,
                  "n_samples": n_samples,
                  "target_column": target_column,})

```

19.4 Step 4: Specification of the Preprocessing Model

After the training and test data are specified and added to the `fun_control` dictionary, `spotPython` allows the specification of a data preprocessing pipeline, e.g., for the scaling of the data or for the one-hot encoding of categorical variables, see Section 14.4. This feature is not used here, so we do not change the default value (which is `None`).

19.5 Step 5: Select Model (algorithm) and `core_model_hyper_dict`

19.5.1 Implementing a Configurable Neural Network With `spotPython`

`spotPython` includes the `Net_lin_reg` class which is implemented in the file `netregression.py`.

This class inherits from the class `Net_Core` which is implemented in the file `netcore.py`, see Section 14.5.1.

```

from torch import nn
import spotPython.torch.netcore as netcore

class Net_lin_reg(netcore.Net_Core):
    def __init__(

```

```

self, _L_in, _L_out, l1, dropout_prob, lr_mult,
batch_size, epochs, k_folds, patience, optimizer,
sgd_momentum
):
    super(Net_lin_reg, self).__init__(
        lr_mult=lr_mult,
        batch_size=batch_size,
        epochs=epochs,
        k_folds=k_folds,
        patience=patience,
        optimizer=optimizer,
        sgd_momentum=sgd_momentum,
    )
    l2 = max(l1 // 2, 4)
    self.fc1 = nn.Linear(_L_in, l1)
    self.fc2 = nn.Linear(l1, l2)
    self.fc3 = nn.Linear(l2, _L_out)
    self.relu = nn.ReLU()
    self.softmax = nn.Softmax(dim=1)
    self.dropout1 = nn.Dropout(p=dropout_prob)
    self.dropout2 = nn.Dropout(p=dropout_prob / 2)

def forward(self, x):
    x = self.fc1(x)
    x = self.relu(x)
    x = self.dropout1(x)
    x = self.fc2(x)
    x = self.relu(x)
    x = self.dropout2(x)
    x = self.fc3(x)
    return x

```

19.5.1.1 The Net_Core class

Net_lin_reg inherits from the class Net_Core which is implemented in the file netcore.py. This class was described in Section 14.5.1.

```

from spotPython.torch.netregression import Net_lin_reg
from spotPython.data.torch_hyper_dict import TorchHyperDict
from spotPython.hyperparameters.values import add_core_model_to_fun_control
fun_control = add_core_model_to_fun_control(core_model=Net_lin_reg,

```

```
fun_control=fun_control,  
hyper_dict=TorchHyperDict,  
filename=None)
```

19.5.2 The Search Space

19.5.3 Configuring the Search Space With spotPython

19.5.3.1 The hyper_dict Hyperparameters for the Selected Algorithm

spotPython uses JSON files for the specification of the hyperparameters, which were described in Section 14.5.5.

The corresponding entries for the `core_model` class are shown below.

```
fun_control['core_model_hyper_dict']  
  
{'_L_in': {'type': 'int',  
  'default': 10,  
  'transform': 'None',  
  'lower': 10,  
  'upper': 10},  
 '_L_out': {'type': 'int',  
  'default': 1,  
  'transform': 'None',  
  'lower': 1,  
  'upper': 1},  
 'l1': {'type': 'int',  
  'default': 3,  
  'transform': 'transform_power_2_int',  
  'lower': 3,  
  'upper': 8},  
 'dropout_prob': {'type': 'float',  
  'default': 0.01,  
  'transform': 'None',  
  'lower': 0.0,  
  'upper': 0.9},  
 'lr_mult': {'type': 'float',  
  'default': 1.0,  
  'transform': 'None',  
  'lower': 0.1,
```

```

    'upper': 10.0},
'batch_size': {'type': 'int',
  'default': 4,
  'transform': 'transform_power_2_int',
  'lower': 1,
  'upper': 4},
'epochs': {'type': 'int',
  'default': 4,
  'transform': 'transform_power_2_int',
  'lower': 4,
  'upper': 9},
'k_folds': {'type': 'int',
  'default': 1,
  'transform': 'None',
  'lower': 1,
  'upper': 1},
'patience': {'type': 'int',
  'default': 2,
  'transform': 'transform_power_2_int',
  'lower': 1,
  'upper': 5},
'optimizer': {'levels': ['Adadelta',
  'Adagrad',
  'Adam',
  'AdamW',
  'SparseAdam',
  'Adamax',
  'ASGD',
  'NAdam',
  'RAdam',
  'RMSprop',
  'Rprop',
  'SGD'],
  'type': 'factor',
  'default': 'SGD',
  'transform': 'None',
  'class_name': 'torch.optim',
  'core_model_parameter_type': 'str',
  'lower': 0,
  'upper': 12},
'sgd_momentum': {'type': 'float',
  'default': 0.0,
  'transform': 'None',

```

```
'lower': 0.0,  
'upper': 1.0}}
```

19.6 Step 6: Modify hyper_dict Hyperparameters for the Selected Algorithm aka core_model

spotPython provides functions for modifying the hyperparameters, their bounds and factors as well as for activating and de-activating hyperparameters without re-compilation of the Python source code. These functions were described in Section [14.6](#).

19.6.1 Modify hyper_dict Hyperparameters for the Selected Algorithm aka core_model

19.6.1.1 Modify Hyperparameters of Type numeric and integer (boolean)

```
from spotPython.hyperparameters.values import modify_hyper_parameter_bounds  
  
fun_control = modify_hyper_parameter_bounds(fun_control, "epochs", bounds=[2, 16])  
fun_control = modify_hyper_parameter_bounds(fun_control, "patience", bounds=[3, 7])
```

19.6.1.2 Modify Hyperparameter of Type factor

```
from spotPython.hyperparameters.values import modify_hyper_parameter_levels  
fun_control = modify_hyper_parameter_levels(fun_control, "optimizer",  
                                           ["Adadelta", "Adagrad", "Adam", "AdamW", "Adamax", "ASGD", "NAdam"])  
  
fun_control.update({  
    "_L_in": n_features,  
    "_L_out": 1,})
```

19.6.2 Optimizers

Optimizers are described in Section [14.6.1](#).

19.7 Step 7: Selection of the Objective (Loss) Function

19.7.1 Evaluation

The evaluation procedure requires the specification of two elements:

1. the way how the data is split into a train and a test set (see Section [14.7.1](#))
2. the loss function (and a metric).

19.7.2 Loss Functions and Metrics

The key "loss_function" specifies the loss function which is used during the optimization, see Section [14.7.5](#).

We will use MSE loss for the regression task.

```
from torch.nn import MSELoss
loss_torch = MSELoss()
fun_control.update({"loss_function": loss_torch})
```

19.7.3 Metric

```
from torchmetrics import MeanAbsoluteError
metric_torch = MeanAbsoluteError().to(fun_control["device"])
fun_control.update({"metric_torch": metric_torch})
```

19.8 Step 8: Calling the SPOT Function

19.8.1 Preparing the SPOT Call

The following code passes the information about the parameter ranges and bounds to `spot`.

```
# extract the variable types, names, and bounds
from spotPython.hyperparameters.values import (get_bound_values,
        get_var_name,
        get_var_type,)
var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
fun_control.update({"var_type": var_type,
```

```

        "var_name": var_name})
lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")

from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))

```

name	type	default	lower	upper	transform
_L_in	int	10	10	10	None
_L_out	int	1	1	1	None
l1	int	3	3	8	transform_power_2_int
dropout_prob	float	0.01	0	0.9	None
lr_mult	float	1.0	0.1	10	None
batch_size	int	4	1	4	transform_power_2_int
epochs	int	4	2	16	transform_power_2_int
k_folds	int	1	1	1	None
patience	int	2	3	7	transform_power_2_int
optimizer	factor	SGD	0	6	None
sgd_momentum	float	0.0	0	1	None

19.8.2 The Objective Function `fun_torch`

The objective function `fun_torch` is selected next. It implements an interface from PyTorch's training, validation, and testing methods to `spotPython`.

```

from spotPython.fun.hypertorch import HyperTorch
fun = HyperTorch().fun_torch

from spotPython.hyperparameters.values import get_default_hyperparameters_as_array
hyper_dict=TorchHyperDict().load()
X_start = get_default_hyperparameters_as_array(fun_control, hyper_dict)

```

19.8.3 Starting the Hyperparameter Tuning

The `spotPython` hyperparameter tuning is started by calling the `Spot` function as described in Section [14.8.4](#).

```

from spotPython.spot import spot
from math import inf
spot_tuner = spot.Spot(fun=fun,
                      lower = lower,
                      upper = upper,
                      fun_evals = inf,
                      fun_repeats = 1,
                      max_time = MAX_TIME,
                      noise = False,
                      tolerance_x = np.sqrt(np.spacing(1)),
                      var_type = var_type,
                      var_name = var_name,
                      infill_criterion = "y",
                      n_points = 1,
                      seed=123,
                      log_level = 50,
                      show_models= False,
                      show_progress= True,
                      fun_control = fun_control,
                      design_control={"init_size": INIT_SIZE,
                                     "repeats": 1},
                      surrogate_control={"noise": True,
                                       "cod_type": "norm",
                                       "min_theta": -4,
                                       "max_theta": 3,
                                       "n_theta": len(var_name),
                                       "model_fun_evals": 10_000,
                                       "log_level": 50
                                       })

spot_tuner.run(X_start=X_start)

```

```

config: {'_L_in': 10, '_L_out': 1, 'l1': 64, 'dropout_prob': 0.7103122166156, 'lr_mult': 3.6}
Epoch: 1 |

```

```

MeanAbsoluteError: 0.1601102948188782 | Loss: 0.0404565179548962 | Epoch: 2 | MeanAbsoluteError:

```

```

MeanAbsoluteError: 0.1188025027513504 | Loss: 0.0226461006156904 | Epoch: 10 |

```

```

MeanAbsoluteError: 0.1162779778242111 | Loss: 0.0218368963902130 | Epoch: 11 | MeanAbsoluteError:

```

MeanAbsoluteError: 0.1110387668013573 | Loss: 0.0201997655002694 | Epoch: 13 | MeanAbsoluteE

MeanAbsoluteError: 0.0965183302760124 | Loss: 0.0158265021592869 | Epoch: 21 |

MeanAbsoluteError: 0.0922852382063866 | Loss: 0.0147978182220341 | Epoch: 22 | MeanAbsoluteE

MeanAbsoluteError: 0.0933773964643478 | Loss: 0.0145014572128850 | Epoch: 24 | MeanAbsoluteE

MeanAbsoluteError: 0.0840223059058189 | Loss: 0.0132114910459342 | Epoch: 32 |

MeanAbsoluteError: 0.0798487067222595 | Loss: 0.0115149722550996 | Epoch: 33 | MeanAbsoluteE

MeanAbsoluteError: 0.0771048516035080 | Loss: 0.0100658434354333 | Epoch: 35 | MeanAbsoluteE

MeanAbsoluteError: 0.0751999095082283 | Loss: 0.0100256902668135 | Epoch: 43 |

MeanAbsoluteError: 0.0762479156255722 | Loss: 0.0095815105167659 | Epoch: 44 | MeanAbsoluteE

MeanAbsoluteError: 0.0837002620100975 | Loss: 0.0120471901701469 | Epoch: 46 | MeanAbsoluteE

MeanAbsoluteError: 0.0691955089569092 | Loss: 0.0084583916131554 | Epoch: 54 |

MeanAbsoluteError: 0.0699153617024422 | Loss: 0.0088047922227385 | Epoch: 55 | MeanAbsoluteE

MeanAbsoluteError: 0.0690940245985985 | Loss: 0.0087129603060404 | Epoch: 57 | MeanAbsoluteE

MeanAbsoluteError: 0.0657618269324303 | Loss: 0.0092513991525936 | Epoch: 65 |

MeanAbsoluteError: 0.0739136412739754 | Loss: 0.0103740964366711 | Epoch: 66 | MeanAbsoluteE

MeanAbsoluteError: 0.0754689499735832 | Loss: 0.0108573954295073 | Epoch: 68 | MeanAbsoluteE

MeanAbsoluteError: 0.0825603902339935 | Loss: 0.0114418036292160 | Early stopping at epoch 74
Returned to Spot: Validation loss: 0.011441803629216003

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.19981931523998656, 'lr_mult':
Epoch: 1 |

MeanAbsoluteError: 0.1285235434770584 | Loss: 0.0255406690075209 | Epoch: 2 | MeanAbsoluteEr
MeanAbsoluteError: 0.1041038855910301 | Loss: 0.0166600543613497 | Epoch: 6 | MeanAbsoluteEr
MeanAbsoluteError: 0.0694405138492584 | Loss: 0.0078225458334935 | Epoch: 21 |
MeanAbsoluteError: 0.0613898709416389 | Loss: 0.0064352994392577 | Epoch: 22 | MeanAbsoluteE
MeanAbsoluteError: 0.0749706923961639 | Loss: 0.0081256309793772 | Epoch: 26 | MeanAbsoluteE
MeanAbsoluteError: 0.1428901404142380 | Loss: 0.0232064912193700 | Epoch: 41 |
MeanAbsoluteError: 0.0527426935732365 | Loss: 0.0045689562102780 | Epoch: 42 | MeanAbsoluteE
MeanAbsoluteError: 0.0561683401465416 | Loss: 0.0051905093101883 | Epoch: 46 | MeanAbsoluteE
MeanAbsoluteError: 0.1025368124246597 | Loss: 0.0124958825150603 | Epoch: 61 |
MeanAbsoluteError: 0.0518786795437336 | Loss: 0.0044410075061023 | Epoch: 62 | MeanAbsoluteE
MeanAbsoluteError: 0.0467758290469646 | Loss: 0.0036402315783658 | Epoch: 66 | MeanAbsoluteE
MeanAbsoluteError: 0.0589869506657124 | Loss: 0.0048729497571721 | Epoch: 81 |
MeanAbsoluteError: 0.0840588286519051 | Loss: 0.0111870647800204 | Epoch: 82 | MeanAbsoluteE
MeanAbsoluteError: 0.0776453688740730 | Loss: 0.0082958192368479 | Epoch: 86 | MeanAbsoluteE
MeanAbsoluteError: 0.0356966890394688 | Loss: 0.0028338138516502 | Epoch: 101 |
MeanAbsoluteError: 0.0442608445882797 | Loss: 0.0031996979776427 | Epoch: 102 | MeanAbsoluteE
MeanAbsoluteError: 0.0576076693832874 | Loss: 0.0049367160107450 | Epoch: 106 | MeanAbsoluteE
MeanAbsoluteError: 0.0624696351587772 | Loss: 0.0056513458073727 | Epoch: 121 |
MeanAbsoluteError: 0.0531211905181408 | Loss: 0.0037555727237639 | Epoch: 122 | MeanAbsoluteE

MeanAbsoluteError: 0.1433090120553970 | Loss: 0.0322113928193054 | Epoch: 27 |

MeanAbsoluteError: 0.1434110254049301 | Loss: 0.0326007635802186 | Epoch: 28 | MeanAbsoluteE

MeanAbsoluteError: 0.1439524888992310 | Loss: 0.0322994961411071 | Epoch: 30 |

MeanAbsoluteError: 0.1396969556808472 | Loss: 0.0315996324743173 | Epoch: 31 | MeanAbsoluteE

MeanAbsoluteError: 0.1406234800815582 | Loss: 0.0324784485347724 | Epoch: 33 |

MeanAbsoluteError: 0.1395697146654129 | Loss: 0.0311478909936947 | Epoch: 34 | MeanAbsoluteE

MeanAbsoluteError: 0.1376837790012360 | Loss: 0.0306010694957513 | Epoch: 36 |

MeanAbsoluteError: 0.1391707360744476 | Loss: 0.0314540806531052 | Epoch: 37 | MeanAbsoluteE

MeanAbsoluteError: 0.1370310485363007 | Loss: 0.0316148921005273 | Epoch: 39 |

MeanAbsoluteError: 0.1375934034585953 | Loss: 0.0306637983043523 | Epoch: 40 | MeanAbsoluteE

MeanAbsoluteError: 0.1430599540472031 | Loss: 0.0326819322009881 | Epoch: 42 |

MeanAbsoluteError: 0.1441847681999207 | Loss: 0.0320939676389874 | Epoch: 43 | MeanAbsoluteE

MeanAbsoluteError: 0.1366764158010483 | Loss: 0.0301272337008656 | Epoch: 45 |

MeanAbsoluteError: 0.1420582085847855 | Loss: 0.0311487587107695 | Epoch: 46 | MeanAbsoluteE

MeanAbsoluteError: 0.1374262273311615 | Loss: 0.0302541757219781 | Epoch: 48 |

MeanAbsoluteError: 0.1379495412111282 | Loss: 0.0314791454951046 | Epoch: 49 | MeanAbsoluteE

MeanAbsoluteError: 0.1411153525114059 | Loss: 0.0321870837764194 | Epoch: 51 |

MeanAbsoluteError: 0.1407373547554016 | Loss: 0.0313053134295236 | Epoch: 52 | MeanAbsoluteE

MeanAbsoluteError: 0.1387081593275070 | Loss: 0.0312055863823237 | Epoch: 54 |

MeanAbsoluteError: 0.1368575096130371 | Loss: 0.0310989914258759 | Epoch: 55 | MeanAbsoluteE

MeanAbsoluteError: 0.1411862522363663 | Loss: 0.0317030826597572 | Epoch: 57 |

MeanAbsoluteError: 0.1407575458288193 | Loss: 0.0315680355561199 | Epoch: 58 | MeanAbsoluteE

MeanAbsoluteError: 0.1384860724210739 | Loss: 0.0312285285633940 | Epoch: 60 |

MeanAbsoluteError: 0.1383236199617386 | Loss: 0.0306956717501938 | Epoch: 61 | MeanAbsoluteE

MeanAbsoluteError: 0.1386741548776627 | Loss: 0.0310652788671238 | Epoch: 63 |

MeanAbsoluteError: 0.1377692073583603 | Loss: 0.0299367827040260 | Epoch: 64 | MeanAbsoluteE

MeanAbsoluteError: 0.1375640481710434 | Loss: 0.0307705067817005 | Epoch: 66 |

MeanAbsoluteError: 0.1389013975858688 | Loss: 0.0303686141961953 | Epoch: 67 | MeanAbsoluteE

MeanAbsoluteError: 0.1394073814153671 | Loss: 0.0311817009362130 | Epoch: 69 |

MeanAbsoluteError: 0.1403766572475433 | Loss: 0.0321219426812604 | Epoch: 70 | MeanAbsoluteE

MeanAbsoluteError: 0.1426118165254593 | Loss: 0.0326408725848887 | Epoch: 72 |

MeanAbsoluteError: 0.1407147198915482 | Loss: 0.0313990318420110 | Epoch: 73 | MeanAbsoluteE

MeanAbsoluteError: 0.1356390565633774 | Loss: 0.0296504653904897 | Epoch: 75 |

MeanAbsoluteError: 0.1393146961927414 | Loss: 0.0309464088677972 | Epoch: 76 | MeanAbsoluteE

MeanAbsoluteError: 0.1393281817436218 | Loss: 0.0310093173607796 | Epoch: 78 |

MeanAbsoluteError: 0.1332419961690903 | Loss: 0.0280423711391631 | Epoch: 79 | MeanAbsoluteE

MeanAbsoluteError: 0.1365984827280045 | Loss: 0.0307244451430840 | Epoch: 81 |

MeanAbsoluteError: 0.1341423243284225 | Loss: 0.0293213327607373 | Epoch: 82 | MeanAbsoluteE

MeanAbsoluteError: 0.1411965787410736 | Loss: 0.0325048757477877 | Epoch: 84 |

MeanAbsoluteError: 0.1401728242635727 | Loss: 0.0313650476891780 | Epoch: 85 | MeanAbsoluteE

MeanAbsoluteError: 0.1399764865636826 | Loss: 0.0314035333051773 | Epoch: 87 |

MeanAbsoluteError: 0.1397745758295059 | Loss: 0.0307054841263744 | Epoch: 88 | MeanAbsoluteE

MeanAbsoluteError: 0.1414524018764496 | Loss: 0.0313294751920330 | Epoch: 90 |

MeanAbsoluteError: 0.1355191171169281 | Loss: 0.0299577264483863 | Epoch: 91 | MeanAbsoluteE

MeanAbsoluteError: 0.1376853883266449 | Loss: 0.0305649406720477 | Epoch: 93 |

MeanAbsoluteError: 0.1361009925603867 | Loss: 0.0306828916974579 | Epoch: 94 | MeanAbsoluteE

MeanAbsoluteError: 0.1377325505018234 | Loss: 0.0293978021023213 | Epoch: 96 |

MeanAbsoluteError: 0.1370553672313690 | Loss: 0.0297081933851587 | Epoch: 97 | MeanAbsoluteE

MeanAbsoluteError: 0.1370528787374496 | Loss: 0.0302792378686718 | Epoch: 99 |

MeanAbsoluteError: 0.1371410638093948 | Loss: 0.0295009834085795 | Epoch: 100 | MeanAbsoluteE

MeanAbsoluteError: 0.1330109089612961 | Loss: 0.0282311750476947 | Epoch: 102 |

MeanAbsoluteError: 0.1352411508560181 | Loss: 0.0298016918178958 | Epoch: 103 | MeanAbsoluteE

MeanAbsoluteError: 0.1395914554595947 | Loss: 0.0307571879188375 | Epoch: 105 |

MeanAbsoluteError: 0.1355815231800079 | Loss: 0.0294904962637277 | Epoch: 106 | MeanAbsoluteE

MeanAbsoluteError: 0.1371858566999435 | Loss: 0.0304561324491321 | Epoch: 108 |

MeanAbsoluteError: 0.1363332420587540 | Loss: 0.0303062791479169 | Epoch: 109 | MeanAbsoluteE

MeanAbsoluteError: 0.1385216712951660 | Loss: 0.0310823449473052 | Epoch: 111 |

MeanAbsoluteError: 0.1359178870916367 | Loss: 0.0295285313763695 | Epoch: 141 |

MeanAbsoluteError: 0.1375299841165543 | Loss: 0.0295349924256637 | Epoch: 142 | MeanAbsolutel

MeanAbsoluteError: 0.1359470337629318 | Loss: 0.0292667949958316 | Epoch: 144 |

MeanAbsoluteError: 0.1372626274824142 | Loss: 0.0304680362900884 | Epoch: 145 | MeanAbsolutel

MeanAbsoluteError: 0.1382762640714645 | Loss: 0.0313304909632037 | Epoch: 147 |

MeanAbsoluteError: 0.1380169540643692 | Loss: 0.0306622618690017 | Epoch: 148 | MeanAbsolutel

MeanAbsoluteError: 0.1341437995433807 | Loss: 0.0289952681551222 | Epoch: 150 |

MeanAbsoluteError: 0.1374948471784592 | Loss: 0.0292981952896419 | Epoch: 151 | MeanAbsolutel

MeanAbsoluteError: 0.1318342536687851 | Loss: 0.0287973995880263 | Epoch: 153 |

MeanAbsoluteError: 0.1358292549848557 | Loss: 0.0293529462134272 | Epoch: 154 | MeanAbsolutel

MeanAbsoluteError: 0.1403776556253433 | Loss: 0.0321658388456854 | Epoch: 156 |

MeanAbsoluteError: 0.1334628164768219 | Loss: 0.0289739706391507 | Epoch: 157 | MeanAbsolutel

MeanAbsoluteError: 0.1343490183353424 | Loss: 0.0297320331781521 | Epoch: 159 |

MeanAbsoluteError: 0.1348469704389572 | Loss: 0.0298085070737094 | Epoch: 160 | MeanAbsolutel

MeanAbsoluteError: 0.1361219733953476 | Loss: 0.0294330167223234 | Epoch: 162 |

MeanAbsoluteError: 0.1370345205068588 | Loss: 0.0296588205251707 | Epoch: 163 | MeanAbsolutel

MeanAbsoluteError: 0.1315025389194489 | Loss: 0.0273800269841255 | Epoch: 165 |

MeanAbsoluteError: 0.1362742930650711 | Loss: 0.0291855764850334 | Epoch: 166 | MeanAbsolutel

MeanAbsoluteError: 0.1375095993280411 | Loss: 0.0288208475947613 | Epoch: 168 |

MeanAbsoluteError: 0.1362537741661072 | Loss: 0.0296346070171664 | Epoch: 169 | MeanAbsoluteError: 0.1379045397043228 | Loss: 0.0297404186014319 | Epoch: 171 |

MeanAbsoluteError: 0.1325749009847641 | Loss: 0.0281810333162624 | Epoch: 172 | MeanAbsoluteError: 0.1361813247203827 | Loss: 0.0293832107112636 | Epoch: 174 |

MeanAbsoluteError: 0.1341451704502106 | Loss: 0.0299626988438346 | Epoch: 175 | MeanAbsoluteError: 0.1327986866235733 | Loss: 0.0284976897125792 | Epoch: 177 |

MeanAbsoluteError: 0.1338888406753540 | Loss: 0.0286528620581764 | Epoch: 178 | MeanAbsoluteError: 0.1393194198608398 | Loss: 0.0308004588026476 | Epoch: 180 |

MeanAbsoluteError: 0.1356994509696960 | Loss: 0.0290231276371439 | Epoch: 181 | MeanAbsoluteError: 0.1368722468614578 | Loss: 0.0294705174944829 | Epoch: 183 |

MeanAbsoluteError: 0.1321402639150620 | Loss: 0.0283420373903997 | Epoch: 184 | MeanAbsoluteError: 0.1341343075037003 | Loss: 0.0293422745988937 | Epoch: 186 |

MeanAbsoluteError: 0.1370597183704376 | Loss: 0.0304942178827696 | Epoch: 187 | MeanAbsoluteError: 0.1282214671373367 | Loss: 0.0258546913651905 | Epoch: 189 |

MeanAbsoluteError: 0.1370003819465637 | Loss: 0.0295187026314670 | Epoch: 190 | MeanAbsoluteError: 0.1304422914981842 | Loss: 0.0272914119136597 | Epoch: 192 |

MeanAbsoluteError: 0.1307411491870880 | Loss: 0.0270123030998123 | Epoch: 193 | MeanAbsoluteError: 0.1334327161312103 | Loss: 0.0276915281801485 | Epoch: 195 |

MeanAbsoluteError: 0.1282100379467010 | Loss: 0.0266633031415404 | Epoch: 196 | MeanAbsoluteError: 0.1304422914981842 | Loss: 0.0272914119136597 | Epoch: 192 |

MeanAbsoluteError: 0.1317851245403290 | Loss: 0.0284395202235707 | Epoch: 198 |

MeanAbsoluteError: 0.1323168873786926 | Loss: 0.0283415334556776 | Epoch: 199 | MeanAbsoluteError: 0.1323168873786926 | Loss: 0.0283415334556776 | Epoch: 199 |

MeanAbsoluteError: 0.1376602351665497 | Loss: 0.0295546540190602 | Epoch: 201 |

MeanAbsoluteError: 0.1315214931964874 | Loss: 0.0282499767783641 | Epoch: 202 | MeanAbsoluteError: 0.1315214931964874 | Loss: 0.0282499767783641 | Epoch: 202 |

MeanAbsoluteError: 0.1323757618665695 | Loss: 0.0276949746489602 | Epoch: 204 |

MeanAbsoluteError: 0.1323472559452057 | Loss: 0.0281917585924566 | Epoch: 205 | MeanAbsoluteError: 0.1323472559452057 | Loss: 0.0281917585924566 | Epoch: 205 |

MeanAbsoluteError: 0.1380791813135147 | Loss: 0.0296551289246903 | Epoch: 207 |

MeanAbsoluteError: 0.1341572552919388 | Loss: 0.0288272898017021 | Epoch: 208 | MeanAbsoluteError: 0.1341572552919388 | Loss: 0.0288272898017021 | Epoch: 208 |

MeanAbsoluteError: 0.1288682073354721 | Loss: 0.0261408846603202 | Epoch: 210 |

MeanAbsoluteError: 0.1330207437276840 | Loss: 0.0291963440484930 | Epoch: 211 | MeanAbsoluteError: 0.1330207437276840 | Loss: 0.0291963440484930 | Epoch: 211 |

MeanAbsoluteError: 0.1394057422876358 | Loss: 0.0313721590387286 | Epoch: 213 |

MeanAbsoluteError: 0.1295985132455826 | Loss: 0.0271830963300999 | Epoch: 214 | MeanAbsoluteError: 0.1295985132455826 | Loss: 0.0271830963300999 | Epoch: 214 |

MeanAbsoluteError: 0.1319358348846436 | Loss: 0.0280005775883910 | Epoch: 216 |

MeanAbsoluteError: 0.1315614879131317 | Loss: 0.0290409765935813 | Epoch: 217 | MeanAbsoluteError: 0.1315614879131317 | Loss: 0.0290409765935813 | Epoch: 217 |

MeanAbsoluteError: 0.1287194788455963 | Loss: 0.0277105604157744 | Epoch: 219 |

MeanAbsoluteError: 0.1299667060375214 | Loss: 0.0268688624852120 | Epoch: 220 | MeanAbsoluteError: 0.1299667060375214 | Loss: 0.0268688624852120 | Epoch: 220 |

MeanAbsoluteError: 0.1351269334554672 | Loss: 0.0291309571649375 | Epoch: 222 |

MeanAbsoluteError: 0.1320157051086426 | Loss: 0.0274007470364450 | Epoch: 223 | MeanAbsoluteError: 0.1320157051086426 | Loss: 0.0274007470364450 | Epoch: 223 |

MeanAbsoluteError: 0.1311803758144379 | Loss: 0.0286482639482711 | Epoch: 225 |

MeanAbsoluteError: 0.1318035423755646 | Loss: 0.0271831327560358 | Epoch: 255 |

MeanAbsoluteError: 0.1255322843790054 | Loss: 0.0255165518723273 | Epoch: 256 | MeanAbsoluteError: 0.1255322843790054 | Loss: 0.0255165518723273 | Epoch: 256 |

MeanAbsoluteError: 0.1328963041305542 | Loss: 0.0274179629157637 | Epoch: 258 |

MeanAbsoluteError: 0.1299600154161453 | Loss: 0.0273879893561631 | Epoch: 259 | MeanAbsoluteError: 0.1299600154161453 | Loss: 0.0273879893561631 | Epoch: 259 |

MeanAbsoluteError: 0.1256195157766342 | Loss: 0.0258862118818797 | Epoch: 261 |

MeanAbsoluteError: 0.1267097890377045 | Loss: 0.0265098322766426 | Epoch: 262 | MeanAbsoluteError: 0.1267097890377045 | Loss: 0.0265098322766426 | Epoch: 262 |

MeanAbsoluteError: 0.1292830854654312 | Loss: 0.0275549091245436 | Epoch: 264 |

MeanAbsoluteError: 0.1313143223524094 | Loss: 0.0275548951422873 | Epoch: 265 | MeanAbsoluteError: 0.1313143223524094 | Loss: 0.0275548951422873 | Epoch: 265 |

MeanAbsoluteError: 0.1296031177043915 | Loss: 0.0263879111133671 | Epoch: 267 |

MeanAbsoluteError: 0.1284626722335815 | Loss: 0.0269588652646780 | Epoch: 268 | MeanAbsoluteError: 0.1284626722335815 | Loss: 0.0269588652646780 | Epoch: 268 |

MeanAbsoluteError: 0.1312844902276993 | Loss: 0.0270620739020039 | Epoch: 270 |

MeanAbsoluteError: 0.1317726969718933 | Loss: 0.0271841016062535 | Epoch: 271 | MeanAbsoluteError: 0.1317726969718933 | Loss: 0.0271841016062535 | Epoch: 271 |

MeanAbsoluteError: 0.1309649348258972 | Loss: 0.0267362869348532 | Epoch: 273 |

MeanAbsoluteError: 0.1332586705684662 | Loss: 0.0282610271194911 | Epoch: 274 | MeanAbsoluteError: 0.1332586705684662 | Loss: 0.0282610271194911 | Epoch: 274 |

MeanAbsoluteError: 0.1304628401994705 | Loss: 0.0270747860043775 | Epoch: 276 |

MeanAbsoluteError: 0.1254433989524841 | Loss: 0.0260165572548431 | Epoch: 277 | MeanAbsoluteError: 0.1254433989524841 | Loss: 0.0260165572548431 | Epoch: 277 |

MeanAbsoluteError: 0.1350724697113037 | Loss: 0.0293320919935165 | Epoch: 279 |

MeanAbsoluteError: 0.1309100687503815 | Loss: 0.0262592869275250 | Epoch: 280 | MeanAbsoluteError: 0.1309100687503815 | Loss: 0.0262592869275250 | Epoch: 280 |

MeanAbsoluteError: 0.1316926330327988 | Loss: 0.0276744631583279 | Epoch: 282 |

MeanAbsoluteError: 0.1315719336271286 | Loss: 0.0277478296079789 | Epoch: 283 | MeanAbsoluteError: 0.1280479878187180 | Loss: 0.0267350972954349 | Epoch: 285 |

MeanAbsoluteError: 0.1314600408077240 | Loss: 0.0274624155480221 | Epoch: 286 | MeanAbsoluteError: 0.1300276815891266 | Loss: 0.0263229235740922 | Epoch: 288 |

MeanAbsoluteError: 0.1381802111864090 | Loss: 0.0299661014220328 | Epoch: 289 | MeanAbsoluteError: 0.1249421164393425 | Loss: 0.0250683379721401 | Epoch: 291 |

MeanAbsoluteError: 0.1254620999097824 | Loss: 0.0254232222107445 | Epoch: 292 | MeanAbsoluteError: 0.1251504570245743 | Loss: 0.0258731536175280 | Epoch: 294 |

MeanAbsoluteError: 0.1259124428033829 | Loss: 0.0254536174224146 | Epoch: 295 | MeanAbsoluteError: 0.1232094839215279 | Loss: 0.0242145087626219 | Epoch: 297 |

MeanAbsoluteError: 0.1261363178491592 | Loss: 0.0256611447573232 | Epoch: 298 | MeanAbsoluteError: 0.1285180002450943 | Loss: 0.0257367642097718 | Epoch: 300 |

MeanAbsoluteError: 0.1268004775047302 | Loss: 0.0258037206227891 | Epoch: 301 | MeanAbsoluteError: 0.1314958930015564 | Loss: 0.0274887814225319 | Epoch: 303 |

MeanAbsoluteError: 0.1253576278686523 | Loss: 0.0249695983647446 | Epoch: 304 | MeanAbsoluteError: 0.1239471212029457 | Loss: 0.0253455219474078 | Epoch: 306 |

MeanAbsoluteError: 0.1211020424962044 | Loss: 0.0245878614609440 | Epoch: 307 | MeanAbsoluteError: 0.1241138577461243 | Loss: 0.0248863854034183 | Epoch: 309 |

MeanAbsoluteError: 0.1298381537199020 | Loss: 0.0276404545996532 | Epoch: 310 | MeanAbsoluteError: 0.1280479878187180 | Loss: 0.0267350972954349 | Epoch: 285 |

MeanAbsoluteError: 0.1256634593009949 | Loss: 0.0256056599076692 | Epoch: 312 |

MeanAbsoluteError: 0.1284182965755463 | Loss: 0.0264935907151084 | Epoch: 313 | MeanAbsoluteError: 0.1284182965755463 | Loss: 0.0264935907151084 | Epoch: 313 |

MeanAbsoluteError: 0.1286122053861618 | Loss: 0.0274394082310270 | Epoch: 315 |

MeanAbsoluteError: 0.1247081086039543 | Loss: 0.0254868176849171 | Epoch: 316 | MeanAbsoluteError: 0.1247081086039543 | Loss: 0.0254868176849171 | Epoch: 316 |

MeanAbsoluteError: 0.1299386918544769 | Loss: 0.0270193143154029 | Epoch: 318 |

MeanAbsoluteError: 0.1267980188131332 | Loss: 0.0251817280432927 | Epoch: 319 | MeanAbsoluteError: 0.1267980188131332 | Loss: 0.0251817280432927 | Epoch: 319 |

MeanAbsoluteError: 0.1229846850037575 | Loss: 0.0253037055706939 | Epoch: 321 |

MeanAbsoluteError: 0.1249031946063042 | Loss: 0.0253790452082952 | Epoch: 322 | MeanAbsoluteError: 0.1249031946063042 | Loss: 0.0253790452082952 | Epoch: 322 |

MeanAbsoluteError: 0.1288567483425140 | Loss: 0.0266044645042408 | Epoch: 324 |

MeanAbsoluteError: 0.1293512433767319 | Loss: 0.0264796367173161 | Epoch: 325 | MeanAbsoluteError: 0.1293512433767319 | Loss: 0.0264796367173161 | Epoch: 325 |

MeanAbsoluteError: 0.1269568651914597 | Loss: 0.0258674473984865 | Epoch: 327 |

MeanAbsoluteError: 0.1332146078348160 | Loss: 0.0276444321005450 | Epoch: 328 | MeanAbsoluteError: 0.1332146078348160 | Loss: 0.0276444321005450 | Epoch: 328 |

MeanAbsoluteError: 0.1234438866376877 | Loss: 0.0253599361762948 | Epoch: 330 |

MeanAbsoluteError: 0.1226711645722389 | Loss: 0.0243985538248671 | Epoch: 331 | MeanAbsoluteError: 0.1226711645722389 | Loss: 0.0243985538248671 | Epoch: 331 |

MeanAbsoluteError: 0.1247135922312737 | Loss: 0.0256391348833373 | Epoch: 333 |

MeanAbsoluteError: 0.1309696286916733 | Loss: 0.0273355852990547 | Epoch: 334 | MeanAbsoluteError: 0.1309696286916733 | Loss: 0.0273355852990547 | Epoch: 334 |

MeanAbsoluteError: 0.1262930780649185 | Loss: 0.0254899970364446 | Epoch: 336 |

MeanAbsoluteError: 0.1201877966523170 | Loss: 0.0241141836998819 | Epoch: 337 | MeanAbsoluteError: 0.1201877966523170 | Loss: 0.0241141836998819 | Epoch: 337 |

MeanAbsoluteError: 0.1213838085532188 | Loss: 0.0242508924591918 | Epoch: 339 |

MeanAbsoluteError: 0.1241269186139107 | Loss: 0.0246884177208024 | Epoch: 369 |

MeanAbsoluteError: 0.1225882247090340 | Loss: 0.0251055830847084 | Epoch: 370 | MeanAbsoluteError: 0.1225882247090340 | Loss: 0.0251055830847084 | Epoch: 370 |

MeanAbsoluteError: 0.1205846369266510 | Loss: 0.0230266395974225 | Epoch: 372 |

MeanAbsoluteError: 0.1239507645368576 | Loss: 0.0249557320701812 | Epoch: 373 | MeanAbsoluteError: 0.1239507645368576 | Loss: 0.0249557320701812 | Epoch: 373 |

MeanAbsoluteError: 0.1271646916866302 | Loss: 0.0253866011677746 | Epoch: 375 |

MeanAbsoluteError: 0.1277868300676346 | Loss: 0.0258335400522143 | Epoch: 376 | MeanAbsoluteError: 0.1277868300676346 | Loss: 0.0258335400522143 | Epoch: 376 |

MeanAbsoluteError: 0.1266208142042160 | Loss: 0.0250559270012551 | Epoch: 378 |

MeanAbsoluteError: 0.1174176037311554 | Loss: 0.0225082880275052 | Epoch: 379 | MeanAbsoluteError: 0.1174176037311554 | Loss: 0.0225082880275052 | Epoch: 379 |

MeanAbsoluteError: 0.1211943849921227 | Loss: 0.0238845143471068 | Epoch: 381 |

MeanAbsoluteError: 0.1179179772734642 | Loss: 0.0236015078611672 | Epoch: 382 | MeanAbsoluteError: 0.1179179772734642 | Loss: 0.0236015078611672 | Epoch: 382 |

MeanAbsoluteError: 0.1212617084383965 | Loss: 0.0238460082990544 | Epoch: 384 |

MeanAbsoluteError: 0.1202387139201164 | Loss: 0.0232610355365129 | Epoch: 385 | MeanAbsoluteError: 0.1202387139201164 | Loss: 0.0232610355365129 | Epoch: 385 |

MeanAbsoluteError: 0.1243932992219925 | Loss: 0.0249200518497188 | Epoch: 387 |

MeanAbsoluteError: 0.1273380070924759 | Loss: 0.0249749634814604 | Epoch: 388 | MeanAbsoluteError: 0.1273380070924759 | Loss: 0.0249749634814604 | Epoch: 388 |

MeanAbsoluteError: 0.1197026222944260 | Loss: 0.0234155171562452 | Epoch: 390 |

MeanAbsoluteError: 0.1275382786989212 | Loss: 0.0258577291603433 | Epoch: 391 | MeanAbsoluteError: 0.1275382786989212 | Loss: 0.0258577291603433 | Epoch: 391 |

MeanAbsoluteError: 0.1248457729816437 | Loss: 0.0259199293009199 | Epoch: 393 |

MeanAbsoluteError: 0.1225762590765953 | Loss: 0.0242783517483622 | Epoch: 394 | MeanAbsoluteError: 0.1225762590765953 | Loss: 0.0242783517483622 | Epoch: 394 |

MeanAbsoluteError: 0.1203770339488983 | Loss: 0.0237145033135675 | Epoch: 396 |

MeanAbsoluteError: 0.1231833770871162 | Loss: 0.0239502648225365 | Epoch: 397 | MeanAbsoluteError: 0.1197540313005447 | Loss: 0.0232518437416487 | Epoch: 399 |

MeanAbsoluteError: 0.1147084459662437 | Loss: 0.0210196565250468 | Epoch: 400 | MeanAbsoluteError: 0.1219423189759254 | Loss: 0.0234856793665737 | Epoch: 402 |

MeanAbsoluteError: 0.1215109899640083 | Loss: 0.0237356624977353 | Epoch: 403 | MeanAbsoluteError: 0.1224004849791527 | Loss: 0.0249684639938490 | Epoch: 405 |

MeanAbsoluteError: 0.1222252920269966 | Loss: 0.0230587176613820 | Epoch: 406 | MeanAbsoluteError: 0.1236640512943268 | Loss: 0.0232954801152421 | Epoch: 408 |

MeanAbsoluteError: 0.1214794665575027 | Loss: 0.0234388577082912 | Epoch: 409 | MeanAbsoluteError: 0.1243082210421562 | Loss: 0.0240326059500997 | Epoch: 411 |

MeanAbsoluteError: 0.1214008703827858 | Loss: 0.0243904422617440 | Epoch: 412 | MeanAbsoluteError: 0.1260180622339249 | Loss: 0.0247075726637073 | Epoch: 414 |

MeanAbsoluteError: 0.1239298880100250 | Loss: 0.0251472057976450 | Epoch: 415 | MeanAbsoluteError: 0.1141055896878242 | Loss: 0.0213864299199971 | Epoch: 417 |

MeanAbsoluteError: 0.1251353621482849 | Loss: 0.0255780191432374 | Epoch: 418 | MeanAbsoluteError: 0.1235620751976967 | Loss: 0.0251959325410523 | Epoch: 420 |

MeanAbsoluteError: 0.1236843392252922 | Loss: 0.0246359554579249 | Epoch: 421 | MeanAbsoluteError: 0.1166776418685913 | Loss: 0.0222318655347287 | Epoch: 423 |

MeanAbsoluteError: 0.1253994405269623 | Loss: 0.0244706834396250 | Epoch: 424 | MeanAbsoluteError: 0.1231833770871162 | Loss: 0.0239502648225365 | Epoch: 397 | MeanAbsoluteError: 0.1197540313005447 | Loss: 0.0232518437416487 | Epoch: 399 |

MeanAbsoluteError: 0.1124109104275703 | Loss: 0.0210668649721871 | Epoch: 426 |

MeanAbsoluteError: 0.1169683933258057 | Loss: 0.0217766026812994 | Epoch: 427 | MeanAbsoluteError: 0.1175745502114296 | Loss: 0.0223663000104473 | Epoch: 429 |

MeanAbsoluteError: 0.1198257803916931 | Loss: 0.0233815158873635 | Epoch: 430 | MeanAbsoluteError: 0.1224837005138397 | Loss: 0.0244661066971215 | Epoch: 432 |

MeanAbsoluteError: 0.1169906109571457 | Loss: 0.0230295199779600 | Epoch: 433 | MeanAbsoluteError: 0.1208820864558220 | Loss: 0.0230436882792371 | Epoch: 435 |

MeanAbsoluteError: 0.1160492449998856 | Loss: 0.0219817942931938 | Epoch: 436 | MeanAbsoluteError: 0.1205234378576279 | Loss: 0.0233987209044668 | Epoch: 438 |

MeanAbsoluteError: 0.1157085150480270 | Loss: 0.0224138965919459 | Epoch: 439 | MeanAbsoluteError: 0.1180680692195892 | Loss: 0.0230636127900652 | Epoch: 441 |

MeanAbsoluteError: 0.1110676825046539 | Loss: 0.0205952375293418 | Epoch: 442 | MeanAbsoluteError: 0.1220044419169426 | Loss: 0.0233308512288689 | Epoch: 444 |

MeanAbsoluteError: 0.1163577288389206 | Loss: 0.0222096639110896 | Epoch: 445 | MeanAbsoluteError: 0.1211250424385071 | Loss: 0.0233090447994376 | Epoch: 447 |

MeanAbsoluteError: 0.1097105666995049 | Loss: 0.0211428390926934 | Epoch: 448 | MeanAbsoluteError: 0.1137713715434074 | Loss: 0.0216073687860141 | Epoch: 450 |

MeanAbsoluteError: 0.1199487522244453 | Loss: 0.0232041026372462 | Epoch: 451 | MeanAbsoluteError: 0.1215065792202950 | Loss: 0.0226742944635286 | Epoch: 453 |

MeanAbsoluteError: 0.1186540573835373 | Loss: 0.0223569689330179 | Epoch: 483 |

MeanAbsoluteError: 0.1158484518527985 | Loss: 0.0218686015723991 | Epoch: 484 | MeanAbsoluteError: 0.1158484518527985 | Loss: 0.0218686015723991 | Epoch: 484 |

MeanAbsoluteError: 0.1174416095018387 | Loss: 0.0216584962303750 | Epoch: 486 |

MeanAbsoluteError: 0.1179016008973122 | Loss: 0.0216957217076318 | Epoch: 487 | MeanAbsoluteError: 0.1179016008973122 | Loss: 0.0216957217076318 | Epoch: 487 |

MeanAbsoluteError: 0.1148181036114693 | Loss: 0.0222993561596377 | Epoch: 489 |

MeanAbsoluteError: 0.1182194277644157 | Loss: 0.0229138464643620 | Epoch: 490 | MeanAbsoluteError: 0.1182194277644157 | Loss: 0.0229138464643620 | Epoch: 490 |

MeanAbsoluteError: 0.1144624426960945 | Loss: 0.0220397393820167 | Epoch: 492 |

MeanAbsoluteError: 0.1142407879233360 | Loss: 0.0208252670606210 | Epoch: 493 | MeanAbsoluteError: 0.1142407879233360 | Loss: 0.0208252670606210 | Epoch: 493 |

MeanAbsoluteError: 0.1178809255361557 | Loss: 0.0227220525057055 | Epoch: 495 |

MeanAbsoluteError: 0.1102326363325119 | Loss: 0.0194642630000696 | Epoch: 496 | MeanAbsoluteError: 0.1102326363325119 | Loss: 0.0194642630000696 | Epoch: 496 |

MeanAbsoluteError: 0.1165632009506226 | Loss: 0.0236518804397686 | Epoch: 498 |

MeanAbsoluteError: 0.1067558303475380 | Loss: 0.0199006765635083 | Epoch: 499 | MeanAbsoluteError: 0.1067558303475380 | Loss: 0.0199006765635083 | Epoch: 499 |

MeanAbsoluteError: 0.1152180582284927 | Loss: 0.0212733052367306 | Epoch: 501 |

MeanAbsoluteError: 0.1103544905781746 | Loss: 0.0206822357676962 | Epoch: 502 | MeanAbsoluteError: 0.1103544905781746 | Loss: 0.0206822357676962 | Epoch: 502 |

MeanAbsoluteError: 0.1166188716888428 | Loss: 0.0218166421200052 | Epoch: 504 |

MeanAbsoluteError: 0.1115140244364738 | Loss: 0.0209139007193153 | Epoch: 505 | MeanAbsoluteError: 0.1115140244364738 | Loss: 0.0209139007193153 | Epoch: 505 |

MeanAbsoluteError: 0.1135153099894524 | Loss: 0.0209780812450723 | Epoch: 507 |

MeanAbsoluteError: 0.1120294854044914 | Loss: 0.0208749749840414 | Epoch: 508 | MeanAbsoluteError: 0.1120294854044914 | Loss: 0.0208749749840414 | Epoch: 508 |

MeanAbsoluteError: 0.1179017126560211 | Loss: 0.0223987717437558 | Epoch: 510 |

MeanAbsoluteError: 0.1198519393801689 | Loss: 0.0231759745894245 | Epoch: 511 | MeanAbsoluteError: 0.1168532446026802 | Loss: 0.0223068611091003 | Epoch: 513 |

MeanAbsoluteError: 0.1136586368083954 | Loss: 0.0215863588681289 | Epoch: 514 | MeanAbsoluteError: 0.1145715937018394 | Loss: 0.0214914524084694 | Epoch: 516 |

MeanAbsoluteError: 0.1103968918323517 | Loss: 0.0200588133686567 | Epoch: 517 | MeanAbsoluteError: 0.1097104772925377 | Loss: 0.0197886418581766 | Epoch: 519 |

MeanAbsoluteError: 0.1132775768637657 | Loss: 0.0202840849188700 | Epoch: 520 | MeanAbsoluteError: 0.1077289357781410 | Loss: 0.0185025214020546 | Epoch: 522 |

MeanAbsoluteError: 0.1115570440888405 | Loss: 0.0208279374519285 | Epoch: 523 | MeanAbsoluteError: 0.1081727445125580 | Loss: 0.0187162002424399 | Epoch: 525 |

MeanAbsoluteError: 0.1167383193969727 | Loss: 0.0217735167617987 | Epoch: 526 | MeanAbsoluteError: 0.1076447963714600 | Loss: 0.0195692195441613 | Epoch: 528 |

MeanAbsoluteError: 0.1133123412728310 | Loss: 0.0206649380481880 | Epoch: 529 | MeanAbsoluteError: 0.1105562075972557 | Loss: 0.0200950857458277 | Epoch: 531 |

MeanAbsoluteError: 0.1109908074140549 | Loss: 0.0205207560964239 | Epoch: 532 | MeanAbsoluteError: 0.1099294275045395 | Loss: 0.0194292617273944 | Epoch: 534 |

MeanAbsoluteError: 0.1186062842607498 | Loss: 0.0234090660795240 | Epoch: 535 | MeanAbsoluteError: 0.1165904998779297 | Loss: 0.0211985721188466 | Epoch: 537 |

MeanAbsoluteError: 0.1186190769076347 | Loss: 0.0226738280033169 | Epoch: 538 | MeanAbsoluteError: 0.1186190769076347 | Loss: 0.0226738280033169 | Epoch: 538 |

MeanAbsoluteError: 0.1132856756448746 | Loss: 0.0205531572037338 | Epoch: 540 |

MeanAbsoluteError: 0.1063591241836548 | Loss: 0.0191815092804124 | Epoch: 541 | MeanAbsoluteError: 0.1118664294481277 | Loss: 0.0204878853089758 | Epoch: 543 |

MeanAbsoluteError: 0.1171410977840424 | Loss: 0.0214114476777604 | Epoch: 544 | MeanAbsoluteError: 0.1146660074591637 | Loss: 0.0212221684833154 | Epoch: 546 |

MeanAbsoluteError: 0.1113823354244232 | Loss: 0.0194008214226536 | Epoch: 547 | MeanAbsoluteError: 0.1066673249006271 | Loss: 0.0184199820691720 | Epoch: 549 |

MeanAbsoluteError: 0.1075967550277710 | Loss: 0.0190857432628157 | Epoch: 550 | MeanAbsoluteError: 0.1089843139052391 | Loss: 0.0198139162301231 | Epoch: 552 |

MeanAbsoluteError: 0.1101785302162170 | Loss: 0.0197380124394840 | Epoch: 553 | MeanAbsoluteError: 0.1123454421758652 | Loss: 0.0205291473852897 | Epoch: 555 |

MeanAbsoluteError: 0.1074836999177933 | Loss: 0.0183233360032318 | Epoch: 556 | MeanAbsoluteError: 0.1071446761488914 | Loss: 0.0183753717834285 | Epoch: 558 |

MeanAbsoluteError: 0.1124401316046715 | Loss: 0.0211839635806973 | Epoch: 559 | MeanAbsoluteError: 0.1136745214462280 | Loss: 0.0210684328707066 | Epoch: 561 |

MeanAbsoluteError: 0.1146416068077087 | Loss: 0.0218512381568932 | Epoch: 562 | MeanAbsoluteError: 0.1137218847870827 | Loss: 0.0215744075819384 | Epoch: 564 |

MeanAbsoluteError: 0.1077500283718109 | Loss: 0.0187633363569815 | Epoch: 565 | MeanAbsoluteError: 0.1111449822783470 | Loss: 0.0205696323991287 | Epoch: 567 |

MeanAbsoluteError: 0.1107379421591759 | Loss: 0.0199199419089321 | Epoch: 597 |
MeanAbsoluteError: 0.1108448132872581 | Loss: 0.0197538779096309 | Epoch: 598 | MeanAbsoluteError:
MeanAbsoluteError: 0.1058530360460281 | Loss: 0.0195769453446458 | Epoch: 600 |
MeanAbsoluteError: 0.1091854348778725 | Loss: 0.0190390770450904 | Epoch: 601 | MeanAbsoluteError:
MeanAbsoluteError: 0.1001044288277626 | Loss: 0.0169095355674654 | Epoch: 603 |
MeanAbsoluteError: 0.1083107218146324 | Loss: 0.0201551477289710 | Epoch: 604 | MeanAbsoluteError:
MeanAbsoluteError: 0.1097362264990807 | Loss: 0.0194977458344268 | Epoch: 606 |
MeanAbsoluteError: 0.1061251461505890 | Loss: 0.0193286593284089 | Epoch: 607 | MeanAbsoluteError:
MeanAbsoluteError: 0.1115806475281715 | Loss: 0.0198378574630381 | Epoch: 609 |
MeanAbsoluteError: 0.1048089042305946 | Loss: 0.0187611381627115 | Epoch: 610 | MeanAbsoluteError:
MeanAbsoluteError: 0.1038298234343529 | Loss: 0.0178828501722698 | Epoch: 612 |
MeanAbsoluteError: 0.1069902032613754 | Loss: 0.0192174725189883 | Epoch: 613 | MeanAbsoluteError:
MeanAbsoluteError: 0.1012466326355934 | Loss: 0.0174716246015547 | Epoch: 615 |
MeanAbsoluteError: 0.1072184219956398 | Loss: 0.0188205216240992 | Epoch: 616 | MeanAbsoluteError:
MeanAbsoluteError: 0.1019762456417084 | Loss: 0.0176599779674628 | Epoch: 618 |
MeanAbsoluteError: 0.1050189062952995 | Loss: 0.0183090117695125 | Epoch: 619 | MeanAbsoluteError:
MeanAbsoluteError: 0.1125763207674026 | Loss: 0.0202363356459803 | Epoch: 621 |
MeanAbsoluteError: 0.1075809597969055 | Loss: 0.0189568840382465 | Epoch: 622 | MeanAbsoluteError:
MeanAbsoluteError: 0.1011569127440453 | Loss: 0.0169694015076925 | Epoch: 624 |

MeanAbsoluteError: 0.1049950271844864 | Loss: 0.0176354743842118 | Epoch: 654 |

MeanAbsoluteError: 0.1074968576431274 | Loss: 0.0191161751476708 | Epoch: 655 | MeanAbsoluteError: 0.1074968576431274 | Loss: 0.0191161751476708 | Epoch: 655 |

MeanAbsoluteError: 0.1003220826387405 | Loss: 0.0168693769249088 | Epoch: 657 |

MeanAbsoluteError: 0.1029541715979576 | Loss: 0.0175482233861476 | Epoch: 658 | MeanAbsoluteError: 0.1029541715979576 | Loss: 0.0175482233861476 | Epoch: 658 |

MeanAbsoluteError: 0.1050288230180740 | Loss: 0.0186305167778240 | Epoch: 660 |

MeanAbsoluteError: 0.1016950160264969 | Loss: 0.0173676276338938 | Epoch: 661 | MeanAbsoluteError: 0.1016950160264969 | Loss: 0.0173676276338938 | Epoch: 661 |

MeanAbsoluteError: 0.1012259349226952 | Loss: 0.0170701968579427 | Epoch: 663 |

MeanAbsoluteError: 0.1030073463916779 | Loss: 0.0169529348081172 | Epoch: 664 | MeanAbsoluteError: 0.1030073463916779 | Loss: 0.0169529348081172 | Epoch: 664 |

MeanAbsoluteError: 0.1036048084497452 | Loss: 0.0177447978372220 | Epoch: 666 |

MeanAbsoluteError: 0.0985486954450607 | Loss: 0.0171744482333058 | Epoch: 667 | MeanAbsoluteError: 0.0985486954450607 | Loss: 0.0171744482333058 | Epoch: 667 |

MeanAbsoluteError: 0.1026226207613945 | Loss: 0.0171081862532689 | Epoch: 669 |

MeanAbsoluteError: 0.0955241248011589 | Loss: 0.0157172781186334 | Epoch: 670 | MeanAbsoluteError: 0.0955241248011589 | Loss: 0.0157172781186334 | Epoch: 670 |

MeanAbsoluteError: 0.1011879220604897 | Loss: 0.0177575510546255 | Epoch: 672 |

MeanAbsoluteError: 0.1066883876919746 | Loss: 0.0187175919121364 | Epoch: 673 | MeanAbsoluteError: 0.1066883876919746 | Loss: 0.0187175919121364 | Epoch: 673 |

MeanAbsoluteError: 0.1027997061610222 | Loss: 0.0176676001956306 | Epoch: 675 |

MeanAbsoluteError: 0.0996896773576736 | Loss: 0.0164347479746599 | Epoch: 676 | MeanAbsoluteError: 0.0996896773576736 | Loss: 0.0164347479746599 | Epoch: 676 |

MeanAbsoluteError: 0.1013355106115341 | Loss: 0.0174178135933835 | Epoch: 678 |

MeanAbsoluteError: 0.1007369384169579 | Loss: 0.0176884642694980 | Epoch: 679 | MeanAbsoluteError: 0.1007369384169579 | Loss: 0.0176884642694980 | Epoch: 679 |

MeanAbsoluteError: 0.1016495674848557 | Loss: 0.0171614808255981 | Epoch: 681 |

MeanAbsoluteError: 0.0981377288699150 | Loss: 0.0164491149381502 | Epoch: 711 |

MeanAbsoluteError: 0.0985470041632652 | Loss: 0.0160858718905365 | Epoch: 712 | MeanAbsoluteError: 0.0981759205460548 | Loss: 0.0161107653316382 | Epoch: 714 |

MeanAbsoluteError: 0.0983947291970253 | Loss: 0.0158231249045881 | Epoch: 715 | MeanAbsoluteError: 0.1015727892518044 | Loss: 0.0179424921590059 | Epoch: 717 |

MeanAbsoluteError: 0.1003616675734520 | Loss: 0.0173221626306380 | Epoch: 718 | MeanAbsoluteError: 0.0986120998859406 | Loss: 0.0167913342037355 | Epoch: 720 |

MeanAbsoluteError: 0.1007067710161209 | Loss: 0.0163432166783605 | Epoch: 721 | MeanAbsoluteError: 0.0973814800381660 | Loss: 0.0161536684884535 | Epoch: 723 |

MeanAbsoluteError: 0.0959409549832344 | Loss: 0.0150623327889965 | Epoch: 724 | MeanAbsoluteError: 0.1061602383852005 | Loss: 0.0184779917397585 | Epoch: 726 |

MeanAbsoluteError: 0.0928905904293060 | Loss: 0.0144650431418268 | Epoch: 727 | MeanAbsoluteError: 0.0985680371522903 | Loss: 0.0157471253839321 | Epoch: 729 |

MeanAbsoluteError: 0.0994159728288651 | Loss: 0.0165386747511608 | Epoch: 730 | MeanAbsoluteError: 0.1083010360598564 | Loss: 0.0195118837318538 | Epoch: 732 |

MeanAbsoluteError: 0.0970344915986061 | Loss: 0.0157529007344662 | Epoch: 733 | MeanAbsoluteError: 0.1006526723504066 | Loss: 0.0173769718848598 | Epoch: 735 |

MeanAbsoluteError: 0.1050557345151901 | Loss: 0.0182304356609772 | Epoch: 736 | MeanAbsoluteError: 0.0997761338949203 | Loss: 0.0160831342991150 | Epoch: 738 |

MeanAbsoluteError: 0.0992054119706154 | Loss: 0.0159764312126208 | Epoch: 768 |

MeanAbsoluteError: 0.1011943370103836 | Loss: 0.0162104140326846 | Epoch: 769 | MeanAbsoluteError: 0.1000064015388489 | Loss: 0.0168411061903741 | Epoch: 771 |

MeanAbsoluteError: 0.0942044034600258 | Loss: 0.0152744118848386 | Epoch: 772 | MeanAbsoluteError: 0.0910591483116150 | Loss: 0.0145060059351575 | Epoch: 774 |

MeanAbsoluteError: 0.0967089906334877 | Loss: 0.0158893699822026 | Epoch: 775 | MeanAbsoluteError: 0.0923384875059128 | Loss: 0.0140206549184101 | Epoch: 777 |

MeanAbsoluteError: 0.0971237048506737 | Loss: 0.0152784394326348 | Epoch: 778 | MeanAbsoluteError: 0.0973426476120949 | Loss: 0.0159019126486965 | Epoch: 780 |

MeanAbsoluteError: 0.0956884622573853 | Loss: 0.0150283993591438 | Epoch: 781 | MeanAbsoluteError: 0.0971700102090836 | Loss: 0.0152078277585739 | Epoch: 783 |

MeanAbsoluteError: 0.0984471291303635 | Loss: 0.0167226721541192 | Epoch: 784 | MeanAbsoluteError: 0.0923895761370659 | Loss: 0.0136245046240704 | Epoch: 786 |

MeanAbsoluteError: 0.0922588482499123 | Loss: 0.0141494800593258 | Epoch: 787 | MeanAbsoluteError: 0.0932257175445557 | Loss: 0.0157888659615370 | Epoch: 789 |

MeanAbsoluteError: 0.0932410731911659 | Loss: 0.0146960425603902 | Epoch: 790 | MeanAbsoluteError: 0.0923603475093842 | Loss: 0.0141003478299050 | Epoch: 792 |

MeanAbsoluteError: 0.0985226184129715 | Loss: 0.0166834699600198 | Epoch: 793 | MeanAbsoluteError: 0.0943809673190117 | Loss: 0.0158968969879788 | Epoch: 795 |

MeanAbsoluteError: 0.0955378711223602 | Loss: 0.0154935410584480 | Epoch: 825 |

MeanAbsoluteError: 0.0931416898965836 | Loss: 0.0146562517306302 | Epoch: 826 | MeanAbsoluteError: 0.0925639942288399 | Loss: 0.0142314712998874 | Epoch: 828 |

MeanAbsoluteError: 0.0947838127613068 | Loss: 0.0153073989096144 | Epoch: 829 | MeanAbsoluteError: 0.0949512645602226 | Loss: 0.0151505220945789 | Epoch: 831 |

MeanAbsoluteError: 0.0945107862353325 | Loss: 0.0157217176605870 | Epoch: 832 | MeanAbsoluteError: 0.0943966805934906 | Loss: 0.0143498473948299 | Epoch: 834 |

MeanAbsoluteError: 0.0950543656945229 | Loss: 0.0149650085644680 | Epoch: 835 | MeanAbsoluteError: 0.0919704139232635 | Loss: 0.0146635710523030 | Epoch: 837 |

MeanAbsoluteError: 0.0866925716400146 | Loss: 0.0137410020935446 | Epoch: 838 | MeanAbsoluteError: 0.0888545066118240 | Loss: 0.0134990676906940 | Epoch: 840 |

MeanAbsoluteError: 0.0897268876433372 | Loss: 0.0138668577935702 | Epoch: 841 | MeanAbsoluteError: 0.0967288091778755 | Loss: 0.0149478258004334 | Epoch: 843 |

MeanAbsoluteError: 0.0940672084689140 | Loss: 0.0148848106863928 | Epoch: 844 | MeanAbsoluteError: 0.0923726409673691 | Loss: 0.0149564693983605 | Epoch: 846 |

MeanAbsoluteError: 0.0972796082496643 | Loss: 0.0154710923142072 | Epoch: 847 | MeanAbsoluteError: 0.0898391231894493 | Loss: 0.0137364220361633 | Epoch: 849 |

MeanAbsoluteError: 0.0972773209214211 | Loss: 0.0157828888055277 | Epoch: 850 | MeanAbsoluteError: 0.0885674580931664 | Loss: 0.0137219115074864 | Epoch: 852 |

MeanAbsoluteError: 0.0922941491007805 | Loss: 0.0145518563371540 | Epoch: 882 |

MeanAbsoluteError: 0.0920543670654297 | Loss: 0.0157285554742460 | Epoch: 883 | MeanAbsoluteError: 0.0920543670654297 | Loss: 0.0157285554742460 | Epoch: 883 |

MeanAbsoluteError: 0.0938300266861916 | Loss: 0.0155638820174499 | Epoch: 885 |

MeanAbsoluteError: 0.0908783748745918 | Loss: 0.0139381371456655 | Epoch: 886 | MeanAbsoluteError: 0.0908783748745918 | Loss: 0.0139381371456655 | Epoch: 886 |

MeanAbsoluteError: 0.0919426381587982 | Loss: 0.0154362994761323 | Epoch: 888 |

MeanAbsoluteError: 0.0883331522345543 | Loss: 0.0148875595483211 | Epoch: 889 | MeanAbsoluteError: 0.0883331522345543 | Loss: 0.0148875595483211 | Epoch: 889 |

MeanAbsoluteError: 0.0919242650270462 | Loss: 0.0148638810523456 | Epoch: 891 |

MeanAbsoluteError: 0.0907325893640518 | Loss: 0.0146195812678222 | Epoch: 892 | MeanAbsoluteError: 0.0907325893640518 | Loss: 0.0146195812678222 | Epoch: 892 |

MeanAbsoluteError: 0.0907133892178535 | Loss: 0.0141848694054352 | Epoch: 894 |

MeanAbsoluteError: 0.0898244380950928 | Loss: 0.0141555894480067 | Epoch: 895 | MeanAbsoluteError: 0.0898244380950928 | Loss: 0.0141555894480067 | Epoch: 895 |

MeanAbsoluteError: 0.0918782129883766 | Loss: 0.0144951475769631 | Epoch: 897 |

MeanAbsoluteError: 0.0886125043034554 | Loss: 0.0137304951814197 | Epoch: 898 | MeanAbsoluteError: 0.0886125043034554 | Loss: 0.0137304951814197 | Epoch: 898 |

MeanAbsoluteError: 0.0889383032917976 | Loss: 0.0141911212831231 | Epoch: 900 |

MeanAbsoluteError: 0.0896413102746010 | Loss: 0.0141183183339551 | Epoch: 901 | MeanAbsoluteError: 0.0896413102746010 | Loss: 0.0141183183339551 | Epoch: 901 |

MeanAbsoluteError: 0.0889725089073181 | Loss: 0.0144346220001413 | Epoch: 903 |

MeanAbsoluteError: 0.0853799879550934 | Loss: 0.0132699511687436 | Epoch: 904 | MeanAbsoluteError: 0.0853799879550934 | Loss: 0.0132699511687436 | Epoch: 904 |

MeanAbsoluteError: 0.0912100449204445 | Loss: 0.0145358693122398 | Epoch: 906 |

MeanAbsoluteError: 0.0913767591118813 | Loss: 0.0137612145014767 | Epoch: 907 | MeanAbsoluteError: 0.0913767591118813 | Loss: 0.0137612145014767 | Epoch: 907 |

MeanAbsoluteError: 0.0873437002301216 | Loss: 0.0127499121374103 | Epoch: 909 |

MeanAbsoluteError: 0.0961063355207443 | Loss: 0.0165980820448992 | Epoch: 939 |

MeanAbsoluteError: 0.0928683057427406 | Loss: 0.0141827037453671 | Epoch: 940 | MeanAbsoluteError:
Returned to Spot: Validation loss: 0.015143395054760731

config: {'_L_in': 10, '_L_out': 1, 'l1': 16, 'dropout_prob': 0.1773189149831582, 'lr_mult': 1.0}
Epoch: 1 | MeanAbsoluteError: 0.1589266955852509 | Loss: 0.0396107463864610 | Epoch: 2 |

MeanAbsoluteError: 0.1385613828897476 | Loss: 0.0305150567926466 | Epoch: 3 | MeanAbsoluteError:

MeanAbsoluteError: 0.1253505945205688 | Loss: 0.0253665243213375 | Returned to Spot: Validation loss:

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.3840970624671163, 'lr_mult': 1.0}
Epoch: 1 | MeanAbsoluteError: 0.1594958007335663 | Loss: 0.0378252000694996 | Epoch: 2 | MeanAbsoluteError:

MeanAbsoluteError: 0.1188225299119949 | Loss: 0.0216705332338614 | Epoch: 9 | MeanAbsoluteError:

MeanAbsoluteError: 0.0922567695379257 | Loss: 0.0130999888673327 | Epoch: 12 | MeanAbsoluteError:

MeanAbsoluteError: 0.0632167607545853 | Loss: 0.0072452452180809 | Epoch: 20 | MeanAbsoluteError:

MeanAbsoluteError: 0.0552488006651402 | Loss: 0.0051966664246529 | Epoch: 23 | MeanAbsoluteError:

MeanAbsoluteError: 0.0509727448225021 | Loss: 0.0046219598949201 | Epoch: 31 | MeanAbsoluteError:

MeanAbsoluteError: 0.0481102131307125 | Loss: 0.0042923886442287 | Epoch: 34 | MeanAbsoluteError:

MeanAbsoluteError: 0.0462850928306580 | Loss: 0.0041914032343285 | Epoch: 42 | MeanAbsoluteError:

MeanAbsoluteError: 0.0526083037257195 | Loss: 0.0052131456537417 | Epoch: 45 | MeanAbsoluteError:

MeanAbsoluteError: 0.0546875186264515 | Loss: 0.0051224719178495 | Epoch: 53 | MeanAbsoluteError:

MeanAbsoluteError: 0.0587547682225704 | Loss: 0.0052781668154670 | Epoch: 56 | MeanAbsoluteError:

MeanAbsoluteError: 0.0486871413886547 | Loss: 0.0048688383172511 | Epoch: 64 | MeanAbsoluteError:

MeanAbsoluteError: 0.0678964033722878 | Loss: 0.0065304376246212 | Epoch: 67 | MeanAbsoluteError:

MeanAbsoluteError: 0.0430374108254910 | Loss: 0.0036049185276313 | Epoch: 75 | MeanAbsoluteE
MeanAbsoluteError: 0.0412441343069077 | Loss: 0.0035073819009612 | Epoch: 78 | MeanAbsoluteE
MeanAbsoluteError: 0.0375385470688343 | Loss: 0.0025915994100567 | Epoch: 86 | MeanAbsoluteE
MeanAbsoluteError: 0.0486825928092003 | Loss: 0.0046143816863677 | Epoch: 89 | MeanAbsoluteE
Epoch: 97 | MeanAbsoluteError: 0.0512637495994568 | Loss: 0.0048938488731724 | Epoch: 98 | M
MeanAbsoluteError: 0.0390989370644093 | Loss: 0.0029917780300725 | Epoch: 100 | MeanAbsoluteE
MeanAbsoluteError: 0.0464583635330200 | Loss: 0.0047108315185668 | Epoch: 108 | MeanAbsoluteE
MeanAbsoluteError: 0.0462678819894791 | Loss: 0.0040446847928443 | Epoch: 111 | MeanAbsoluteE
MeanAbsoluteError: 0.0449468269944191 | Loss: 0.0037452920690779 | Epoch: 119 | MeanAbsoluteE
MeanAbsoluteError: 0.0385937988758087 | Loss: 0.0031782205945714 | Epoch: 122 | MeanAbsoluteE
MeanAbsoluteError: 0.0397592559456825 | Loss: 0.0027568183768859 | Epoch: 130 | MeanAbsoluteE
MeanAbsoluteError: 0.0441173948347569 | Loss: 0.0040700652645778 | Epoch: 133 | MeanAbsoluteE
Epoch: 141 | MeanAbsoluteError: 0.0624889247119427 | Loss: 0.0069354968038265 | Epoch: 142 |
Epoch: 144 | MeanAbsoluteError: 0.0457574538886547 | Loss: 0.0038846386477847 | Epoch: 145 |
Returned to Spot: Validation loss: 0.0036308488557360284

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.24580168789011042, 'lr_mult':
Epoch: 1 | MeanAbsoluteError: 0.1775597184896469 | Loss: 0.0445753564371874 | Epoch: 2 | Mean
MeanAbsoluteError: 0.1454614251852036 | Loss: 0.0260739028453827 | Epoch: 21 | MeanAbsoluteE
MeanAbsoluteError: 0.1459318995475769 | Loss: 0.0252722677235541 | Epoch: 41 | MeanAbsoluteE

MeanAbsoluteError: 0.1268258541822433 | Loss: 0.0187740137212371 | Epoch: 61 | MeanAbsoluteError: 0.1268258541822433

MeanAbsoluteError: 0.0641660168766975 | Loss: 0.0054286807211802 | Epoch: 81 | MeanAbsoluteError: 0.0641660168766975

MeanAbsoluteError: 0.0529515109956264 | Loss: 0.0046637564542164 | Epoch: 101 | MeanAbsoluteError: 0.0529515109956264

MeanAbsoluteError: 0.0443784222006798 | Loss: 0.0040601667500787 | Epoch: 121 | MeanAbsoluteError: 0.0443784222006798

Returned to Spot: Validation loss: 0.01517733024727357

spotPython tuning: 0.002616552669726508 [-----] 2.71%

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.22837518654402147, 'lr_mult': 1.0}
Epoch: 1 | MeanAbsoluteError: 0.1308252066373825 | Loss: 0.0265597031990948 | Epoch: 2 | MeanAbsoluteError: 0.1308252066373825

MeanAbsoluteError: 0.1069270074367523 | Loss: 0.0178842768563252 | Epoch: 21 | MeanAbsoluteError: 0.1069270074367523

MeanAbsoluteError: 0.0597703754901886 | Loss: 0.0062055303592627 | Epoch: 41 | MeanAbsoluteError: 0.0597703754901886

MeanAbsoluteError: 0.1089378520846367 | Loss: 0.0150901309067481 | Epoch: 61 | MeanAbsoluteError: 0.1089378520846367

MeanAbsoluteError: 0.0719462558627129 | Loss: 0.0067318116902913 | Epoch: 81 | MeanAbsoluteError: 0.0719462558627129

MeanAbsoluteError: 0.0696906223893166 | Loss: 0.0081204373487516 | Epoch: 101 | MeanAbsoluteError: 0.0696906223893166

MeanAbsoluteError: 0.0611213296651840 | Loss: 0.0048668051849266 | Epoch: 121 | MeanAbsoluteError: 0.0611213296651840

MeanAbsoluteError: 0.0380189679563046 | Loss: 0.0024404489300459 | Epoch: 141 | MeanAbsoluteError: 0.0380189679563046

Returned to Spot: Validation loss: 0.005170660289494615

spotPython tuning: 0.002616552669726508 [#-----] 5.96%

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.19678865158046985, 'lr_mult': 1.0}
Epoch: 1 | MeanAbsoluteError: 0.1275790929794312 | Loss: 0.0253350573778152 | Epoch: 2 | MeanAbsoluteError: 0.1275790929794312

MeanAbsoluteError: 0.0661873370409012 | Loss: 0.0069111135570953 | Epoch: 7 | MeanAbsoluteError: 0.0661873370409012

MeanAbsoluteError: 0.0467432662844658 | Loss: 0.0043522030753472 | Epoch: 13 | MeanAbsoluteError: 0.0467432662844658

MeanAbsoluteError: 0.0466980598866940 | Loss: 0.0039003083331045 | Epoch: 19 | MeanAbsoluteE
MeanAbsoluteError: 0.0457979887723923 | Loss: 0.0040228050011986 | Epoch: 25 | MeanAbsoluteE
MeanAbsoluteError: 0.0347631834447384 | Loss: 0.0025733338040300 | Epoch: 31 | MeanAbsoluteE
MeanAbsoluteError: 0.0444794781506062 | Loss: 0.0032055134675466 | Epoch: 37 | MeanAbsoluteE
MeanAbsoluteError: 0.0444412045180798 | Loss: 0.0046647486839599 | Epoch: 43 | MeanAbsoluteE
MeanAbsoluteError: 0.0509976670145988 | Loss: 0.0046998962290430 | Epoch: 49 | MeanAbsoluteE
MeanAbsoluteError: 0.0509128123521805 | Loss: 0.0050164787824421 | Epoch: 55 | MeanAbsoluteE
MeanAbsoluteError: 0.0405981652438641 | Loss: 0.0033934029686498 | Epoch: 61 | MeanAbsoluteE
MeanAbsoluteError: 0.0533174723386765 | Loss: 0.0074600999536536 | Epoch: 67 | MeanAbsoluteE
MeanAbsoluteError: 0.0753298401832581 | Loss: 0.0079268644641464 | Epoch: 73 | MeanAbsoluteE
MeanAbsoluteError: 0.0492283552885056 | Loss: 0.0050330607480525 | Epoch: 79 | MeanAbsoluteE
MeanAbsoluteError: 0.0493206232786179 | Loss: 0.0053274197876453 | Epoch: 85 | MeanAbsoluteE
MeanAbsoluteError: 0.0561037883162498 | Loss: 0.0059418327745516 | Epoch: 91 | MeanAbsoluteE
MeanAbsoluteError: 0.0489932931959629 | Loss: 0.0058350499216855 | Epoch: 97 | MeanAbsoluteE
MeanAbsoluteError: 0.0501342155039310 | Loss: 0.0051888407611599 | Epoch: 103 | MeanAbsoluteE
Returned to Spot: Validation loss: 0.005212400178036963

spotPython tuning: 0.002616552669726508 [#-----] 12.69%

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.30307433204932804, 'lr_mult':
Epoch: 1 | MeanAbsoluteError: 0.1706100106239319 | Loss: 0.0451106128135794 | Epoch: 2 | Mean

MeanAbsoluteError: 0.1653519570827484 | Loss: 0.0355563676474910 | Epoch: 13 | MeanAbsoluteE

MeanAbsoluteError: 0.0747667104005814 | Loss: 0.0089444822974895 | Epoch: 25 | MeanAbsoluteE
Returned to Spot: Validation loss: 0.016290426266526704

spotPython tuning: 0.002616552669726508 [#-----] 14.32%

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.8628392823470216, 'lr_mult': 1
Epoch: 1 | MeanAbsoluteError: 0.2146806865930557 | Loss: 0.0723293547175432 | Epoch: 2 | Mean

MeanAbsoluteError: 0.1322571486234665 | Loss: 0.0282383693480178 | Epoch: 20 | MeanAbsoluteE

MeanAbsoluteError: 0.1324374079704285 | Loss: 0.0281580863030333 | Epoch: 39 | MeanAbsoluteE

MeanAbsoluteError: 0.1271349340677261 | Loss: 0.0275305566426955 | Epoch: 59 | MeanAbsoluteE

MeanAbsoluteError: 0.1154948696494102 | Loss: 0.0212981859222054 | Epoch: 78 | MeanAbsoluteE

MeanAbsoluteError: 0.1144222542643547 | Loss: 0.0212672250345349 | Epoch: 97 | MeanAbsoluteE

MeanAbsoluteError: 0.1010147854685783 | Loss: 0.0171388577130672 | Epoch: 117 | MeanAbsolutel

MeanAbsoluteError: 0.0958815738558769 | Loss: 0.0176301640703490 | Epoch: 136 | MeanAbsolutel

MeanAbsoluteError: 0.0934851467609406 | Loss: 0.0156441889703274 | Epoch: 155 | MeanAbsolutel

MeanAbsoluteError: 0.0991537570953369 | Loss: 0.0176874779066757 | Epoch: 175 | MeanAbsolutel

MeanAbsoluteError: 0.1081068441271782 | Loss: 0.0213530002848098 | Epoch: 194 | MeanAbsolutel
Returned to Spot: Validation loss: 0.018607519567012787

spotPython tuning: 0.002616552669726508 [##-----] 18.60%

config: {'_L_in': 10, '_L_out': 1, 'l1': 64, 'dropout_prob': 0.7103107631285819, 'lr_mult': 1
Epoch: 1 | MeanAbsoluteError: 0.1484457105398178 | Loss: 0.0354797716128329 | Epoch: 2 | Mean

MeanAbsoluteError: 0.1385626494884491 | Loss: 0.0296473820817967 | Epoch: 6 |

MeanAbsoluteError: 0.1356103718280792 | Loss: 0.0286731227394193 | Epoch: 7 | MeanAbsoluteError: 0.1356103718280792 | Loss: 0.0286731227394193 | Epoch: 7 |

MeanAbsoluteError: 0.1379132717847824 | Loss: 0.0301302732372036 | Epoch: 12 |

MeanAbsoluteError: 0.1361361891031265 | Loss: 0.0297206921176985 | Epoch: 13 | MeanAbsoluteError: 0.1361361891031265 | Loss: 0.0297206921176985 | Epoch: 13 |

MeanAbsoluteError: 0.1287064403295517 | Loss: 0.0272098044585437 | Epoch: 18 |

MeanAbsoluteError: 0.1286418288946152 | Loss: 0.0275937212367232 | Epoch: 19 | MeanAbsoluteError: 0.1286418288946152 | Loss: 0.0275937212367232 | Epoch: 19 |

MeanAbsoluteError: 0.1244635879993439 | Loss: 0.0267346769458770 | Epoch: 24 |

MeanAbsoluteError: 0.1428149044513702 | Loss: 0.0322177020693198 | Epoch: 25 | MeanAbsoluteError: 0.1428149044513702 | Loss: 0.0322177020693198 | Epoch: 25 |

MeanAbsoluteError: 0.1264004707336426 | Loss: 0.0267952924503091 | Epoch: 30 |

MeanAbsoluteError: 0.1339374482631683 | Loss: 0.0290738498543700 | Epoch: 31 | MeanAbsoluteError: 0.1339374482631683 | Loss: 0.0290738498543700 | Epoch: 31 |

Returned to Spot: Validation loss: 0.03217077505231525
spotPython tuning: 0.002616552669726508 [##-----] 20.64%

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.1890993228956366, 'lr_mult': 1.0}

Epoch: 1 | MeanAbsoluteError: 0.1773824244737625 | Loss: 0.0453626172813146 | Epoch: 2 | MeanAbsoluteError: 0.1773824244737625 | Loss: 0.0453626172813146 | Epoch: 2 |

MeanAbsoluteError: 0.0970599725842476 | Loss: 0.0139431930637281 | Epoch: 13 | MeanAbsoluteError: 0.0970599725842476 | Loss: 0.0139431930637281 | Epoch: 13 |

MeanAbsoluteError: 0.0822291150689125 | Loss: 0.0091155101422613 | Epoch: 25 | MeanAbsoluteError: 0.0822291150689125 | Loss: 0.0091155101422613 | Epoch: 25 |

Returned to Spot: Validation loss: 0.018819046893010016
spotPython tuning: 0.002616552669726508 [##-----] 22.45%

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.3835145831979707, 'lr_mult': 1.0}

Epoch: 1 | MeanAbsoluteError: 0.1489920169115067 | Loss: 0.0356704042244114 | Epoch: 2 | MeanAbsoluteError: 0.1489920169115067 | Loss: 0.0356704042244114 | Epoch: 2 |

MeanAbsoluteError: 0.0820037648081779 | Loss: 0.0115389609035399 | Epoch: 12 | MeanAbsoluteE
MeanAbsoluteError: 0.0596814677119255 | Loss: 0.0058988091094713 | Epoch: 23 | MeanAbsoluteE
MeanAbsoluteError: 0.0582116134464741 | Loss: 0.0058222524974936 | Epoch: 34 | MeanAbsoluteE
MeanAbsoluteError: 0.0492109656333923 | Loss: 0.0048178747268761 | Epoch: 45 | MeanAbsoluteE
MeanAbsoluteError: 0.0504262782633305 | Loss: 0.0045588772550936 | Epoch: 56 | MeanAbsoluteE
MeanAbsoluteError: 0.0598067790269852 | Loss: 0.0063105148606394 | Epoch: 67 | MeanAbsoluteE
MeanAbsoluteError: 0.0794244483113289 | Loss: 0.0094187081619901 | Epoch: 78 | MeanAbsoluteE
MeanAbsoluteError: 0.0437474623322487 | Loss: 0.0036801710972979 | Epoch: 89 | MeanAbsoluteE
Epoch: 100 | MeanAbsoluteError: 0.0562689900398254 | Loss: 0.0059406774388136 | Epoch: 101 |
MeanAbsoluteError: 0.0607690699398518 | Loss: 0.0068627441865637 | Epoch: 111 | MeanAbsoluteE
Epoch: 122 | MeanAbsoluteError: 0.0598649866878986 | Loss: 0.0066692082266519 | Epoch: 123 |
MeanAbsoluteError: 0.0535027496516705 | Loss: 0.0052996922596848 | Epoch: 133 | MeanAbsoluteE
Epoch: 144 | MeanAbsoluteError: 0.0468924976885319 | Loss: 0.0040765977326756 | Epoch: 145 |
MeanAbsoluteError: 0.0695816949009895 | Loss: 0.0074796419235339 | Epoch: 155 | MeanAbsoluteE
Returned to Spot: Validation loss: 0.004756120258733924

spotPython tuning: 0.002616552669726508 [###-----] 28.08%

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.40135685211463074, 'lr_mult':
Epoch: 1 | MeanAbsoluteError: 0.1695955097675323 | Loss: 0.0437562819815388 | Epoch: 2 | Mean

MeanAbsoluteError: 0.1079661324620247 | Loss: 0.0187867664683022 | Epoch: 12 | MeanAbsoluteE

MeanAbsoluteError: 0.0875907763838768 | Loss: 0.0126785195898265 | Epoch: 23 | MeanAbsoluteE
MeanAbsoluteError: 0.0613874047994614 | Loss: 0.0065921421199547 | Epoch: 34 | MeanAbsoluteE
MeanAbsoluteError: 0.0597726888954639 | Loss: 0.0061596736042319 | Epoch: 46 | MeanAbsoluteE
MeanAbsoluteError: 0.0547145195305347 | Loss: 0.0053882033113807 | Epoch: 57 | MeanAbsoluteE
MeanAbsoluteError: 0.0495417602360249 | Loss: 0.0046079174897290 | Epoch: 68 | MeanAbsoluteE
Epoch: 79 | MeanAbsoluteError: 0.0474483631551266 | Loss: 0.0038625205098039 | Epoch: 80 | M
MeanAbsoluteError: 0.0476360954344273 | Loss: 0.0038943909457885 | Epoch: 90 | MeanAbsoluteE
MeanAbsoluteError: 0.0600732751190662 | Loss: 0.0066696898092663 | Epoch: 101 | MeanAbsoluteE
MeanAbsoluteError: 0.0483831837773323 | Loss: 0.0044231331220975 | Epoch: 112 | MeanAbsoluteE
MeanAbsoluteError: 0.0410999022424221 | Loss: 0.0030726504728112 | Epoch: 123 | MeanAbsoluteE
MeanAbsoluteError: 0.0428813584148884 | Loss: 0.0037233148445682 | Epoch: 134 | MeanAbsoluteE
MeanAbsoluteError: 0.0580237768590450 | Loss: 0.0047887946956938 | Epoch: 145 | MeanAbsoluteE
MeanAbsoluteError: 0.0470833890140057 | Loss: 0.0042149910219259 | Epoch: 156 | MeanAbsoluteE
Returned to Spot: Validation loss: 0.0032871504985163675

spotPython tuning: 0.002616552669726508 [###-----] 34.18%

config: {'_L_in': 10, '_L_out': 1, 'l1': 32, 'dropout_prob': 0.5036868011192788, 'lr_mult': 0
Epoch: 1 | MeanAbsoluteError: 0.5931276679039001 | Loss: 0.3876352882698962 | Epoch: 2 | Mean

MeanAbsoluteError: 0.1885185986757278 | Loss: 0.0561530161894074 | Epoch: 12 | MeanAbsoluteE

MeanAbsoluteError: 0.1695781201124191 | Loss: 0.0453501213587036 | Epoch: 23 | MeanAbsoluteE

MeanAbsoluteError: 0.1609495580196381 | Loss: 0.0389755930457460 | Epoch: 34 | MeanAbsoluteE

MeanAbsoluteError: 0.1499064862728119 | Loss: 0.0370555248748707 | Epoch: 45 | MeanAbsoluteE

MeanAbsoluteError: 0.1495599895715714 | Loss: 0.0351669098691721 | Epoch: 56 | MeanAbsoluteE

MeanAbsoluteError: 0.1450883597135544 | Loss: 0.0322472212934180 | Epoch: 67 | MeanAbsoluteE

MeanAbsoluteError: 0.1369211524724960 | Loss: 0.0286052595995563 | Epoch: 78 | MeanAbsoluteE

MeanAbsoluteError: 0.1317282915115356 | Loss: 0.0284721241700218 | Epoch: 89 | MeanAbsoluteE

MeanAbsoluteError: 0.1357117295265198 | Loss: 0.0278955287986288 | Epoch: 100 | MeanAbsoluteE

MeanAbsoluteError: 0.1209888681769371 | Loss: 0.0217036396436589 | Epoch: 111 | MeanAbsoluteE

MeanAbsoluteError: 0.1195110455155373 | Loss: 0.0221455399388153 | Epoch: 122 | MeanAbsoluteE

MeanAbsoluteError: 0.1228398233652115 | Loss: 0.0228623643296918 | Epoch: 133 | MeanAbsoluteE

MeanAbsoluteError: 0.1220626831054688 | Loss: 0.0222875854417093 | Epoch: 144 | MeanAbsoluteE

MeanAbsoluteError: 0.1157317608594894 | Loss: 0.0212904203585104 | Epoch: 155 | MeanAbsoluteE

MeanAbsoluteError: 0.1132365539669991 | Loss: 0.0207484713265378 | Epoch: 166 | MeanAbsoluteE

MeanAbsoluteError: 0.1124493032693863 | Loss: 0.0203929552195692 | Epoch: 177 | MeanAbsoluteE

MeanAbsoluteError: 0.1103010177612305 | Loss: 0.0189801586032110 | Epoch: 188 | MeanAbsoluteE

MeanAbsoluteError: 0.1113350316882133 | Loss: 0.0202908302461238 | Epoch: 199 | MeanAbsoluteE

MeanAbsoluteError: 0.0997997745871544 | Loss: 0.0160152249920525 | Epoch: 210 | MeanAbsoluteE

MeanAbsoluteError: 0.1015908494591713 | Loss: 0.0158760220939784 | Epoch: 221 | MeanAbsoluteE

MeanAbsoluteError: 0.1028972342610359 | Loss: 0.0177104644060723 | Epoch: 232 | MeanAbsoluteE

MeanAbsoluteError: 0.0580024980008602 | Loss: 0.0054147510211139 | Epoch: 49 | MeanAbsoluteE

MeanAbsoluteError: 0.0366057083010674 | Loss: 0.0026266111431566 | Epoch: 55 | MeanAbsoluteE

MeanAbsoluteError: 0.0440691933035851 | Loss: 0.0033013638947159 | Epoch: 61 | MeanAbsoluteE

MeanAbsoluteError: 0.0355960801243782 | Loss: 0.0024697162634887 | Epoch: 67 | MeanAbsoluteE

MeanAbsoluteError: 0.0441784635186195 | Loss: 0.0032939165762218 | Epoch: 73 | MeanAbsoluteE

MeanAbsoluteError: 0.0371045731008053 | Loss: 0.0029758054268314 | Epoch: 79 | MeanAbsoluteE

MeanAbsoluteError: 0.0345577895641327 | Loss: 0.0023971278136984 | Epoch: 85 | MeanAbsoluteE

MeanAbsoluteError: 0.0387997254729271 | Loss: 0.0026114425817893 | Epoch: 91 | MeanAbsoluteE

MeanAbsoluteError: 0.0357546694576740 | Loss: 0.0025868095501210 | Epoch: 97 | MeanAbsoluteE

MeanAbsoluteError: 0.0376084260642529 | Loss: 0.0025542132544797 | Epoch: 103 | MeanAbsoluteE

MeanAbsoluteError: 0.0363630689680576 | Loss: 0.0025113605616692 | Epoch: 109 | MeanAbsoluteE

MeanAbsoluteError: 0.0347551181912422 | Loss: 0.0022541738994914 | Epoch: 115 | MeanAbsoluteE

MeanAbsoluteError: 0.0326102934777737 | Loss: 0.0018003056561186 | Epoch: 121 | MeanAbsoluteE

MeanAbsoluteError: 0.0354748331010342 | Loss: 0.0028137335800178 | Epoch: 127 | MeanAbsoluteE

MeanAbsoluteError: 0.0365852825343609 | Loss: 0.0028655014280230 | Epoch: 133 | MeanAbsoluteE

MeanAbsoluteError: 0.0372179709374905 | Loss: 0.0028832547731387 | Epoch: 139 | MeanAbsoluteE

MeanAbsoluteError: 0.0375973507761955 | Loss: 0.0025292277737753 | Epoch: 145 | MeanAbsoluteE

MeanAbsoluteError: 0.0362293124198914 | Loss: 0.0024698337898008 | Epoch: 151 | MeanAbsoluteE

MeanAbsoluteError: 0.0353367440402508 | Loss: 0.0027586833432239 | Epoch: 157 | MeanAbsoluteE

19.10 Step 10: Results

After the hyperparameter tuning run is finished, the results can be analyzed as described in Section 14.10.

```
spot_tuner.plot_progress(log_y=False,  
                        filename="./figures/" + experiment_name+"_progress.png")
```




Figure 19.1: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

```
print(gen_design_table(fun_control=fun_control, spot=spot_tuner))
```

name	type	default	lower	upper	tuned	transform
_L_in	int	10	10.0	10.0	10.0	None
_L_out	int	1	1.0	1.0	1.0	None
l1	int	3	3.0	8.0	5.0	transform_po
dropout_prob	float	0.01	0.0	0.9	0.19981931523998656	None
lr_mult	float	1.0	0.1	10.0	7.004318498645526	None
batch_size	int	4	1.0	4.0	4.0	transform_po
epochs	int	4	2.0	16.0	11.0	transform_po
k_folds	int	1	1.0	1.0	1.0	None
patience	int	2	3.0	7.0	5.0	transform_po
optimizer	factor	SGD	0.0	6.0	0.0	None
sgd_momentum	float	0.0	0.0	1.0	0.07401195908206384	None

```
spot_tuner.plot_importance(threshold=0.025,
                           filename="./figures/" + experiment_name+"_importance.png")
```




Figure 19.2: Variable importance plot, threshold 0.025.

19.10.1 Get the Tuned Architecture (SPOT Results)

```
from spotPython.hyperparameters.values import get_one_core_model_from_X
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))
model_spot = get_one_core_model_from_X(X, fun_control)
model_spot
```

```
Net_lin_reg(
  (fc1): Linear(in_features=10, out_features=32, bias=True)
  (fc2): Linear(in_features=32, out_features=16, bias=True)
  (fc3): Linear(in_features=16, out_features=1, bias=True)
  (relu): ReLU()
  (softmax): Softmax(dim=1)
  (dropout1): Dropout(p=0.19981931523998656, inplace=False)
  (dropout2): Dropout(p=0.09990965761999328, inplace=False)
)
```

19.10.2 Evaluation of the Tuned Architecture

```
from spotPython.torch.traintest import (
    train_tuned,
    test_tuned,
)

train_tuned(net=model_spot, train_dataset=train,
            loss_function=fun_control["loss_function"],
            metric=fun_control["metric_torch"],
            shuffle=True,
            device = fun_control["device"],
            path=None,
            task=fun_control["task"],)
```

```
Epoch: 1 | MeanAbsoluteError: 0.1370587348937988 | Loss: 0.0296040675357768 | Epoch: 2 | Mean
MeanAbsoluteError: 0.1273256391286850 | Loss: 0.0254505154137549 | Epoch: 6 | MeanAbsoluteEr
MeanAbsoluteError: 0.1084648370742798 | Loss: 0.0162495498692519 | Epoch: 21 | MeanAbsoluteE
MeanAbsoluteError: 0.0921800211071968 | Loss: 0.0120782979862078 | Epoch: 26 | MeanAbsoluteE
MeanAbsoluteError: 0.0608531460165977 | Loss: 0.0063501899130642 | Epoch: 41 | MeanAbsoluteE
MeanAbsoluteError: 0.0506178848445415 | Loss: 0.0047787399411103 | Epoch: 46 | MeanAbsoluteE
MeanAbsoluteError: 0.0487709417939186 | Loss: 0.0039907744901843 | Epoch: 61 | MeanAbsoluteE
MeanAbsoluteError: 0.1488473564386368 | Loss: 0.0273869378786338 | Epoch: 66 | MeanAbsoluteE
MeanAbsoluteError: 0.0829284191131592 | Loss: 0.0092043102424788 | Epoch: 81 | MeanAbsoluteE
MeanAbsoluteError: 0.0774341747164726 | Loss: 0.0094372643108823 | Epoch: 86 | MeanAbsoluteE
MeanAbsoluteError: 0.0726040303707123 | Loss: 0.0073204866708501 | Epoch: 101 | MeanAbsolutel
MeanAbsoluteError: 0.0559390708804131 | Loss: 0.0048706169289194 | Epoch: 106 | MeanAbsolutel
```

```

MeanAbsoluteError: 0.0840886309742928 | Loss: 0.0093243038281798 | Epoch: 121 | MeanAbsoluteError: 0.0840886309742928
MeanAbsoluteError: 0.0962029844522476 | Loss: 0.0123873427755346 | Epoch: 126 | MeanAbsoluteError: 0.0962029844522476
MeanAbsoluteError: 0.0664651840925217 | Loss: 0.0069836770723525 | Epoch: 142 | MeanAbsoluteError: 0.0664651840925217
MeanAbsoluteError: 0.0932655557990074 | Loss: 0.0104182182419065 | Epoch: 147 | MeanAbsoluteError: 0.0932655557990074
MeanAbsoluteError: 0.1201086938381195 | Loss: 0.0178982440480276 | Epoch: 162 | MeanAbsoluteError: 0.1201086938381195
MeanAbsoluteError: 0.0302110612392426 | Loss: 0.0021190302918273 | Epoch: 167 | MeanAbsoluteError: 0.0302110612392426
MeanAbsoluteError: 0.0590001121163368 | Loss: 0.0049689801556892 | Epoch: 182 | MeanAbsoluteError: 0.0590001121163368
Returned to Spot: Validation loss: 0.008977211006966076

```

If path is set to a filename, e.g., path = "model_spot_trained.pt", the weights of the trained model will be loaded from this file.

```

test_tuned(net=model_spot, test_dataset=test,
           shuffle=False,
           loss_function=fun_control["loss_function"],
           metric=fun_control["metric_torch"],
           device = fun_control["device"],
           task=fun_control["task"],)

```

```

MeanAbsoluteError: 0.0897793695330620 | Loss: 0.0096923084929585 | Final evaluation: Validation loss: 0.008977211006966076
Final evaluation: Validation metric: 0.08977936953306198

```

```

(0.009692308492958546, nan, tensor(0.0898))

```

19.10.3 Cross-validated Evaluations

- This is the evaluation that will be used in the comparison (evaluatecv has to be updated before, to get metric vlaues!):

```

from spotPython.torch.traintest import evaluate_cv
# modify k-kolds:

```

```
setattr(model_spot, "k_folds", 10)
df_eval, df_preds, df_metrics = evaluate_cv(net=model_spot,
      dataset=fun_control["data"],
      loss_function=fun_control["loss_function"],
      metric=fun_control["metric_torch"],
      task=fun_control["task"],
      writer=fun_control["writer"],
      writerId="model_spot_cv",
      device = fun_control["device"])
```

Fold: 1

Epoch: 1 | MeanAbsoluteError: 0.2074535191059113 | Loss: 0.0568870306015015 | Epoch: 2 |

MeanAbsoluteError: 0.1662064939737320 | Loss: 0.0420318836612361 | Epoch: 3 | MeanAbsoluteError:

MeanAbsoluteError: 0.1947877258062363 | Loss: 0.0447941528899329 | Epoch: 11 | MeanAbsoluteError:

MeanAbsoluteError: 0.0690206214785576 | Loss: 0.0081713101826608 | Epoch: 13 | MeanAbsoluteError:

MeanAbsoluteError: 0.0666046813130379 | Loss: 0.0085620973591826 | Epoch: 15 | MeanAbsoluteError:

MeanAbsoluteError: 0.0844514369964600 | Loss: 0.0123406526898699 | Epoch: 23 | MeanAbsoluteError:

MeanAbsoluteError: 0.1196470633149147 | Loss: 0.0246764271120940 | Epoch: 25 | MeanAbsoluteError:

MeanAbsoluteError: 0.0736473277211189 | Loss: 0.0080379972766553 | Epoch: 27 | MeanAbsoluteError:

MeanAbsoluteError: 0.1055080518126488 | Loss: 0.0147056539676019 | Epoch: 35 | MeanAbsoluteError:

MeanAbsoluteError: 0.0430901236832142 | Loss: 0.0027375899787460 | Epoch: 37 | MeanAbsoluteError:

MeanAbsoluteError: 0.0461486764252186 | Loss: 0.0035476277282994 | Epoch: 39 | MeanAbsoluteError:

MeanAbsoluteError: 0.0895985960960388 | Loss: 0.0097864660035287 | Epoch: 47 | MeanAbsoluteError:

MeanAbsoluteError: 0.0772603824734688 | Loss: 0.0082105003031237 | Epoch: 49 | MeanAbsoluteError:

MeanAbsoluteError: 0.0399408563971519 | Loss: 0.0022265247368653 | Epoch: 68 | MeanAbsoluteE

MeanAbsoluteError: 0.0321230478584766 | Loss: 0.0018570320348122 | Epoch: 76 | MeanAbsoluteE

MeanAbsoluteError: 0.0608917176723480 | Loss: 0.0044224276207387 | Epoch: 78 | MeanAbsoluteE

MeanAbsoluteError: 0.0570819079875946 | Loss: 0.0038611046371183 | Epoch: 80 | MeanAbsoluteE

MeanAbsoluteError: 0.0795265883207321 | Loss: 0.0078631915551211 | Epoch: 88 | MeanAbsoluteE

MeanAbsoluteError: 0.0495218485593796 | Loss: 0.0045100579570447 | Epoch: 90 | MeanAbsoluteE

MeanAbsoluteError: 0.0427144579589367 | Loss: 0.0032473812046062 | Epoch: 92 | MeanAbsoluteE

MeanAbsoluteError: 0.0529210604727268 | Loss: 0.0037669243922989 | Epoch: 100 | MeanAbsoluteE

MeanAbsoluteError: 0.0308412592858076 | Loss: 0.0016081157872187 | Epoch: 102 | MeanAbsoluteE

MeanAbsoluteError: 0.1251624524593353 | Loss: 0.0226234257487314 | Epoch: 104 | MeanAbsoluteE

MeanAbsoluteError: 0.0232830904424191 | Loss: 0.0011768913037875 | Epoch: 112 | MeanAbsoluteE

MeanAbsoluteError: 0.0373995900154114 | Loss: 0.0022113513467567 | Epoch: 114 | MeanAbsoluteE

MeanAbsoluteError: 0.1019664257764816 | Loss: 0.0121278472776924 | Epoch: 116 | MeanAbsoluteE

MeanAbsoluteError: 0.1250427663326263 | Loss: 0.0195915512740612 | Epoch: 124 | MeanAbsoluteE

MeanAbsoluteError: 0.0343714617192745 | Loss: 0.0019450686273298 | Epoch: 126 | MeanAbsoluteE

MeanAbsoluteError: 0.0393057651817799 | Loss: 0.0027623526818518 | Epoch: 128 | MeanAbsoluteE

MeanAbsoluteError: 0.0860252305865288 | Loss: 0.0084457891727132 | Epoch: 136 | MeanAbsoluteE

MeanAbsoluteError: 0.0419947430491447 | Loss: 0.0028372788801789 | Epoch: 138 | MeanAbsoluteE

MeanAbsoluteError: 0.0890257805585861 | Loss: 0.0109283778417323 | Epoch: 71 | MeanAbsoluteE
MeanAbsoluteError: 0.0636503174901009 | Loss: 0.0062934905290604 | Epoch: 79 | MeanAbsoluteE
MeanAbsoluteError: 0.1296603381633759 | Loss: 0.0189358596024769 | Epoch: 81 | MeanAbsoluteE
MeanAbsoluteError: 0.0400524809956551 | Loss: 0.0027643695274102 | Epoch: 83 | MeanAbsoluteE
MeanAbsoluteError: 0.0320393815636635 | Loss: 0.0021322620949442 | Epoch: 91 | MeanAbsoluteE
MeanAbsoluteError: 0.0714936405420303 | Loss: 0.0094864719680377 | Epoch: 93 | MeanAbsoluteE
MeanAbsoluteError: 0.0424944646656513 | Loss: 0.0035145867482892 | Epoch: 95 | MeanAbsoluteE
Fold: 5
Epoch: 1 | MeanAbsoluteError: 0.2128690779209137 | Loss: 0.0633987880178860 | Epoch: 2 | Mean
MeanAbsoluteError: 0.0684989914298058 | Loss: 0.0069009853926088 | Epoch: 8 | MeanAbsoluteE
MeanAbsoluteError: 0.2181337475776672 | Loss: 0.0560905140425478 | Epoch: 10 | MeanAbsoluteE
MeanAbsoluteError: 0.1143022254109383 | Loss: 0.0160619865304657 | Epoch: 12 | MeanAbsoluteE
MeanAbsoluteError: 0.0461497195065022 | Loss: 0.0038473213623677 | Epoch: 20 | MeanAbsoluteE
MeanAbsoluteError: 0.1638076454401016 | Loss: 0.0321640063609396 | Epoch: 22 | MeanAbsoluteE
MeanAbsoluteError: 0.0675982013344765 | Loss: 0.0071500979496964 | Epoch: 24 | MeanAbsoluteE
MeanAbsoluteError: 0.0577786304056644 | Loss: 0.0046593038007684 | Epoch: 32 | MeanAbsoluteE
MeanAbsoluteError: 0.1285535097122192 | Loss: 0.0190455727279186 | Epoch: 34 | MeanAbsoluteE
MeanAbsoluteError: 0.0576459839940071 | Loss: 0.0054684286005795 | Epoch: 36 | MeanAbsoluteE
MeanAbsoluteError: 0.0375327728688717 | Loss: 0.0020682091791449 | Epoch: 44 | MeanAbsoluteE
MeanAbsoluteError: 0.0545713864266872 | Loss: 0.0041839025049870 | Epoch: 46 | MeanAbsoluteE

MeanAbsoluteError: 0.1498713642358780 | Loss: 0.0254791323095560 | Epoch: 48 | MeanAbsoluteE

MeanAbsoluteError: 0.0420306362211704 | Loss: 0.0034979840607515 | Epoch: 56 | MeanAbsoluteE

MeanAbsoluteError: 0.0767666995525360 | Loss: 0.0081599744568978 | Epoch: 58 | MeanAbsoluteE

MeanAbsoluteError: 0.0704204365611076 | Loss: 0.0058203193558646 | Epoch: 60 | MeanAbsoluteE

MeanAbsoluteError: 0.1394743323326111 | Loss: 0.0206028744578362 | Epoch: 68 | MeanAbsoluteE

MeanAbsoluteError: 0.0397067368030548 | Loss: 0.0021038191709002 | Epoch: 70 | MeanAbsoluteE

MeanAbsoluteError: 0.0537646487355232 | Loss: 0.0036111596772181 | Epoch: 72 | MeanAbsoluteE

MeanAbsoluteError: 0.0320966355502605 | Loss: 0.0016343049480513 | Epoch: 80 | MeanAbsoluteE

MeanAbsoluteError: 0.0689565166831017 | Loss: 0.0061769180798105 | Epoch: 82 | MeanAbsoluteE

MeanAbsoluteError: 0.0466229729354382 | Loss: 0.0033827678167394 | Epoch: 84 | MeanAbsoluteE

MeanAbsoluteError: 0.0465262047946453 | Loss: 0.0033745758534808 | Epoch: 92 | MeanAbsoluteE

MeanAbsoluteError: 0.0667791366577148 | Loss: 0.0059603773988783 | Epoch: 94 | MeanAbsoluteE

MeanAbsoluteError: 0.0571188181638718 | Loss: 0.0051670750336988 | Epoch: 96 | MeanAbsoluteE

MeanAbsoluteError: 0.0317107960581779 | Loss: 0.0027027608427618 | Epoch: 104 | MeanAbsoluteE

MeanAbsoluteError: 0.0347635187208652 | Loss: 0.0020779435290024 | Epoch: 106 | MeanAbsoluteE

MeanAbsoluteError: 0.0503709279000759 | Loss: 0.0032750917931220 | Epoch: 108 | MeanAbsoluteE

MeanAbsoluteError: 0.0508828833699226 | Loss: 0.0033752926059866 | Epoch: 116 | MeanAbsoluteE

MeanAbsoluteError: 0.0345439016819000 | Loss: 0.0019689274985077 | Epoch: 118 | MeanAbsoluteE

MeanAbsoluteError: 0.0860725715756416 | Loss: 0.0104001417223896 | Epoch: 120 | MeanAbsoluteE

MeanAbsoluteError: 0.0690166875720024 | Loss: 0.0071214307099581 | Epoch: 59 | MeanAbsoluteE

MeanAbsoluteError: 0.0411369353532791 | Loss: 0.0031270795568292 | Epoch: 61 | MeanAbsoluteE

MeanAbsoluteError: 0.0516226813197136 | Loss: 0.0051458943302610 | Epoch: 69 | MeanAbsoluteE

MeanAbsoluteError: 0.0294028241187334 | Loss: 0.0014659352933190 | Epoch: 71 | MeanAbsoluteE

MeanAbsoluteError: 0.0610737465322018 | Loss: 0.0053998500620115 | Epoch: 73 | MeanAbsoluteE

MeanAbsoluteError: 0.0842074230313301 | Loss: 0.0095718503663582 | Epoch: 81 | MeanAbsoluteE

MeanAbsoluteError: 0.0511040054261684 | Loss: 0.0035406228354467 | Epoch: 83 | MeanAbsoluteE

MeanAbsoluteError: 0.0933254882693291 | Loss: 0.0120574829966894 | Epoch: 85 | MeanAbsoluteE

MeanAbsoluteError: 0.0678454786539078 | Loss: 0.0066296332515776 | Epoch: 93 | MeanAbsoluteE

MeanAbsoluteError: 0.0902836546301842 | Loss: 0.0094881671081696 | Epoch: 95 | MeanAbsoluteE

MeanAbsoluteError: 0.0424184240400791 | Loss: 0.0026035617338493 | Epoch: 97 | MeanAbsoluteE

MeanAbsoluteError: 0.0281500890851021 | Loss: 0.0020000811782666 | Epoch: 105 | MeanAbsolutel

MeanAbsoluteError: 0.0628974288702011 | Loss: 0.0056362181369747 | Epoch: 107 | MeanAbsolutel

MeanAbsoluteError: 0.0326974205672741 | Loss: 0.0022886837416861 | Epoch: 109 | MeanAbsolutel

MeanAbsoluteError: 0.0559377782046795 | Loss: 0.0044457019373242 | Epoch: 117 | MeanAbsolutel

MeanAbsoluteError: 0.0769031122326851 | Loss: 0.0100162563446377 | Epoch: 119 | MeanAbsolutel

MeanAbsoluteError: 0.0841882005333900 | Loss: 0.0088208856593285 | Epoch: 121 | MeanAbsolutel

MeanAbsoluteError: 0.0594875141978264 | Loss: 0.0048883236345968 | Epoch: 129 | MeanAbsolutel

MeanAbsoluteError: 0.0620157159864902 | Loss: 0.0051237943116575 | Epoch: 131 | MeanAbsolutel

MeanAbsoluteError: 0.1364165395498276 | Loss: 0.0202788484415838 | Epoch: 67 | MeanAbsoluteE

MeanAbsoluteError: 0.0703717917203903 | Loss: 0.0059199528768659 | Epoch: 75 | MeanAbsoluteE

MeanAbsoluteError: 0.0518786236643791 | Loss: 0.0032977144832590 | Epoch: 77 | MeanAbsoluteE

MeanAbsoluteError: 0.0350730717182159 | Loss: 0.0019058479145834 | Epoch: 79 | MeanAbsoluteE

MeanAbsoluteError: 0.0621104389429092 | Loss: 0.0048273009514170 | Epoch: 87 | MeanAbsoluteE

MeanAbsoluteError: 0.0394702777266502 | Loss: 0.0020904428924301 | Epoch: 89 | MeanAbsoluteE

MeanAbsoluteError: 0.0333004631102085 | Loss: 0.0018569822090545 | Epoch: 91 | MeanAbsoluteE

MeanAbsoluteError: 0.0443581156432629 | Loss: 0.0025392353002514 | Epoch: 99 | MeanAbsoluteE

MeanAbsoluteError: 0.0754874646663666 | Loss: 0.0071894304294671 | Epoch: 101 | MeanAbsoluteE

MeanAbsoluteError: 0.0630412325263023 | Loss: 0.0054307399384145 | Early stopping at epoch 10
Fold: 8
Epoch: 1 | MeanAbsoluteError: 0.1156925112009048 | Loss: 0.0179215836396907 | Epoch: 2 | Mean

MeanAbsoluteError: 0.0813768282532692 | Loss: 0.0098236128687859 | Epoch: 9 | MeanAbsoluteE

MeanAbsoluteError: 0.0755551755428314 | Loss: 0.0088962934179498 | Epoch: 11 | MeanAbsoluteE

MeanAbsoluteError: 0.0575885027647018 | Loss: 0.0055108191445470 | Epoch: 13 | MeanAbsoluteE

MeanAbsoluteError: 0.1062692478299141 | Loss: 0.0163583468113627 | Epoch: 21 | MeanAbsoluteE

MeanAbsoluteError: 0.1012537702918053 | Loss: 0.0134234594048134 | Epoch: 23 | MeanAbsoluteE

MeanAbsoluteError: 0.0791864097118378 | Loss: 0.0082532927127821 | Epoch: 25 | MeanAbsoluteE

MeanAbsoluteError: 0.0822725072503090 | Loss: 0.0087978024967015 | Epoch: 33 | MeanAbsoluteE

MeanAbsoluteError: 0.0484629012644291 | Loss: 0.0041886536803629 | Epoch: 35 | MeanAbsoluteE

MeanAbsoluteError: 0.0359791368246078 | Loss: 0.0024832545812907 | Epoch: 37 | MeanAbsoluteE

MeanAbsoluteError: 0.0768731981515884 | Loss: 0.0075377783505246 | Epoch: 45 | MeanAbsoluteE

MeanAbsoluteError: 0.1013431772589684 | Loss: 0.0154116731137037 | Epoch: 47 | MeanAbsoluteE

MeanAbsoluteError: 0.0352274924516678 | Loss: 0.0021427150474795 | Epoch: 49 | MeanAbsoluteE

MeanAbsoluteError: 0.0375023409724236 | Loss: 0.0027384013270161 | Epoch: 57 | MeanAbsoluteE

MeanAbsoluteError: 0.0589179769158363 | Loss: 0.0046783224679530 | Epoch: 59 | MeanAbsoluteE

MeanAbsoluteError: 0.0746714621782303 | Loss: 0.0067318119441292 | Epoch: 61 | MeanAbsoluteE

MeanAbsoluteError: 0.0765606537461281 | Loss: 0.0071533454715141 | Epoch: 69 | MeanAbsoluteE

MeanAbsoluteError: 0.0550800897181034 | Loss: 0.0044759520928242 | Epoch: 71 | MeanAbsoluteE

MeanAbsoluteError: 0.0353555269539356 | Loss: 0.0019493311909693 | Epoch: 73 | MeanAbsoluteE

MeanAbsoluteError: 0.0814784616231918 | Loss: 0.0109383894928864 | Epoch: 81 | MeanAbsoluteE

MeanAbsoluteError: 0.0487115867435932 | Loss: 0.0037295889508511 | Epoch: 83 | MeanAbsoluteE

MeanAbsoluteError: 0.0370934158563614 | Loss: 0.0033076892557022 | Epoch: 85 | MeanAbsoluteE

MeanAbsoluteError: 0.0440841279923916 | Loss: 0.0033477660202022 | Epoch: 93 | MeanAbsoluteE

MeanAbsoluteError: 0.0456829071044922 | Loss: 0.0033408167738734 | Epoch: 95 | MeanAbsoluteE

MeanAbsoluteError: 0.0472746230661869 | Loss: 0.0032127811407138 | Epoch: 97 | MeanAbsoluteE

MeanAbsoluteError: 0.0784432366490364 | Loss: 0.0079786722947444 | Epoch: 105 | MeanAbsolutel

MeanAbsoluteError: 0.0534877106547356 | Loss: 0.0046265995728650 | Epoch: 107 | MeanAbsolutel

MeanAbsoluteError: 0.0656893104314804 | Loss: 0.0067169601264011 | Epoch: 109 | MeanAbsolutel

MeanAbsoluteError: 0.0285384468734264 | Loss: 0.0015082774417741 | Epoch: 66 | MeanAbsoluteE
MeanAbsoluteError: 0.0687597841024399 | Loss: 0.0066980745510331 | Epoch: 68 | MeanAbsoluteE
MeanAbsoluteError: 0.0706854835152626 | Loss: 0.0068893689396126 | Epoch: 70 | MeanAbsoluteE
MeanAbsoluteError: 0.0938337594270706 | Loss: 0.0121193771649684 | Epoch: 78 | MeanAbsoluteE
MeanAbsoluteError: 0.0341943055391312 | Loss: 0.0020393597660586 | Epoch: 80 | MeanAbsoluteE
MeanAbsoluteError: 0.0933109447360039 | Loss: 0.0129981432110071 | Epoch: 82 | MeanAbsoluteE
MeanAbsoluteError: 0.0552825257182121 | Loss: 0.0039910430620824 | Epoch: 90 | MeanAbsoluteE
MeanAbsoluteError: 0.0472887232899666 | Loss: 0.0034648361511894 | Epoch: 92 | MeanAbsoluteE
MeanAbsoluteError: 0.0608188360929489 | Loss: 0.0055414781173957 | Epoch: 94 | MeanAbsoluteE
MeanAbsoluteError: 0.0519650690257549 | Loss: 0.0040370161046407 | Epoch: 102 | MeanAbsoluteE
MeanAbsoluteError: 0.0224034544080496 | Loss: 0.0009628019179218 | Epoch: 104 | MeanAbsoluteE
MeanAbsoluteError: 0.0482478886842728 | Loss: 0.0041446289978921 | Epoch: 106 | MeanAbsoluteE
MeanAbsoluteError: 0.0872446447610855 | Loss: 0.0096452876127192 | Epoch: 114 | MeanAbsoluteE
MeanAbsoluteError: 0.0839137732982635 | Loss: 0.0099963785282203 | Epoch: 116 | MeanAbsoluteE
MeanAbsoluteError: 0.0356399789452553 | Loss: 0.0019720518695457 | Epoch: 118 | MeanAbsoluteE
MeanAbsoluteError: 0.0637345165014267 | Loss: 0.0061902959631490 | Epoch: 126 | MeanAbsoluteE
MeanAbsoluteError: 0.0816340520977974 | Loss: 0.0085473517870664 | Epoch: 128 | MeanAbsoluteE
MeanAbsoluteError: 0.0555757433176041 | Loss: 0.0072591837628611 | Epoch: 130 | MeanAbsoluteE
Fold: 10
Epoch: 1 | MeanAbsoluteError: 0.2425923943519592 | Loss: 0.0862785130739212 | Epoch: 2 |

MeanAbsoluteError: 0.0331846438348293 | Loss: 0.0024428384245506 | Epoch: 77 | MeanAbsoluteE
MeanAbsoluteError: 0.0703106746077538 | Loss: 0.0066594999682690 | Epoch: 79 | MeanAbsoluteE
MeanAbsoluteError: 0.0261448640376329 | Loss: 0.0012291404834416 | Epoch: 87 | MeanAbsoluteE
MeanAbsoluteError: 0.0571299977600574 | Loss: 0.0044869496154466 | Epoch: 89 | MeanAbsoluteE
MeanAbsoluteError: 0.0373268276453018 | Loss: 0.0019920719246979 | Epoch: 91 | MeanAbsoluteE
MeanAbsoluteError: 0.0308549497276545 | Loss: 0.0028105406839001 | Epoch: 99 | MeanAbsoluteE
MeanAbsoluteError: 0.0313827544450760 | Loss: 0.0021055158078005 | Epoch: 101 | MeanAbsoluteE
MeanAbsoluteError: 0.0933097377419472 | Loss: 0.0116164674982429 | Epoch: 103 | MeanAbsoluteE
MeanAbsoluteError: 0.0642709657549858 | Loss: 0.0066170622594655 | Epoch: 111 | MeanAbsoluteE
MeanAbsoluteError: 0.0863091275095940 | Loss: 0.0098497354691582 | Epoch: 113 | MeanAbsoluteE
MeanAbsoluteError: 0.0325586535036564 | Loss: 0.0021816564673957 | Epoch: 115 | MeanAbsoluteE

```
metric_name = type(fun_control["metric_torch"]).__name__  
print(f"loss: {df_eval}, Cross-validated {metric_name}: {df_metrics}")
```

loss: 0.005902799277100713, Cross-validated MeanAbsoluteError: 0.06115343049168587

19.10.4 Detailed Hyperparameter Plots

```
filename = "./figures/" + experiment_name  
spot_tuner.plot_important_hyperparameter_contour(filename=filename)
```

```
l1: 0.3845829890028423  
lr_mult: 0.6700872031230476  
batch_size: 34.88391531114552  
epochs: 2.81430669035911  
optimizer: 99.99999999999999  
sgd_momentum: 7.9175041709670655
```





























Figure 19.3: Contour plots.















19.10.5 Parallel Coordinates Plot

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Parallel coordinates plots

Unable to display output for mime type(s): text/html

19.11 Summary and Outlook

This tutorial presents the hyperparameter tuning open source software `spotPython` for PyTorch. Some of the advantages of `spotPython` are:

- Numerical and categorical hyperparameters.
- Powerful surrogate models.
- Flexible approach and easy to use.
- Simple JSON files for the specification of the hyperparameters.
- Extension of default and user specified network classes.
- Noise handling techniques.
- Online visualization of the hyperparameter tuning process with `tensorboard`.

Currently, only rudimentary parallel and distributed neural network training is possible, but these capabilities will be extended in the future. The next version of `spotPython` will also include a more detailed documentation and more examples.

! Important

Important: This tutorial does not present a complete benchmarking study (Bartz-Beielstein et al. 2020). The results are only preliminary and highly dependent on the local configuration (hard- and software). Our goal is to provide a first impression of the performance of the hyperparameter tuning package `spotPython`. The results should be interpreted with care.

20 HPT: PyTorch With VBDP

In this tutorial, we will show how `spotPython` can be integrated into the PyTorch training workflow for a classification task.

 Caution: Data must be downloaded manually

- Ensure that the corresponding data is available as `./data/VBDP/train.csv`.

This document refers to the following software versions:

- python: 3.10.10
- torch: 2.0.1
- torchvision: 0.15.0

```
pip list | grep "spot[RiverPython]"
```

```
spotPython          0.2.38
spotRiver           0.0.93
```

Note: you may need to restart the kernel to use updated packages.

`spotPython` can be installed via `pip`. Alternatively, the source code can be downloaded from `gitHub`: <https://github.com/sequential-parameter-optimization/spotPython>.

```
!pip install spotPython
```

- Uncomment the following lines if you want to for (re-)installation the latest version of `spotPython` from `gitHub`.

```
# import sys
# !{sys.executable} -m pip install --upgrade build
# !{sys.executable} -m pip install --upgrade --force-reinstall spotPython
```

20.1 Step 1: Setup

Before we consider the detailed experimental setup, we select the parameters that affect run time, initial design size and the device that is used.

 Caution: Run time and initial design size should be increased for real experiments

- `MAX_TIME` is set to one minute for demonstration purposes. For real experiments, this should be increased to at least 1 hour.
- `INIT_SIZE` is set to 5 for demonstration purposes. For real experiments, this should be increased to at least 10.

 Note: Device selection

- The device can be selected by setting the variable `DEVICE`.
- Since we are using a simple neural net, the setting "cpu" is preferred (on Mac).
- If you have a GPU, you can use "cuda:0" instead.
- If `DEVICE` is set to `None`, `spotPython` will automatically select the device.
 - This might result in "mps" on Macs, which is not the best choice for simple neural nets.

```
MAX_TIME = 1
INIT_SIZE = 5
DEVICE = None # "cpu" # "cuda:0"
```

```
from spotPython.utils.device import getDevice
DEVICE = getDevice(DEVICE)
print(DEVICE)
```

mps

```
import os
import copy
import socket
from datetime import datetime
from dateutil.tz import tzlocal
start_time = datetime.now(tzlocal())
HOSTNAME = socket.gethostname().split(".")[0]
experiment_name = '25-torch' + "_" + HOSTNAME + "_" + str(MAX_TIME) + "min_" + str(INIT_SIZE)
experiment_name = experiment_name.replace(':', '-')
```

```
print(experiment_name)
if not os.path.exists('./figures'):
    os.makedirs('./figures')
```

25-torch_bartz09_1min_5init_2023-06-19_04-32-01

20.2 Step 2: Initialization of the fun_control Dictionary

 Caution: Tensorboard does not work under Windows

- Since tensorboard does not work under Windows, we recommend setting the parameter `tensorboard_path` to `None` if you are working under Windows.

spotPython uses a Python dictionary for storing the information required for the hyperparameter tuning process, which was described in Section 14.2, see [Initialization of the fun_control Dictionary](#) in the documentation.

```
from spotPython.utils.init import fun_control_init
fun_control = fun_control_init(task="classification",
                               tensorboard_path="runs/25_spot_torch_vbdp",
                               device=DEVICE)
```

20.3 Step 3: PyTorch Data Loading

20.3.1 1. Load VBDP Data

```
import pandas as pd
from sklearn.preprocessing import OrdinalEncoder
train_df = pd.read_csv('./data/VBDP/train.csv')
# remove the id column
train_df = train_df.drop(columns=['id'])
n_samples = train_df.shape[0]
n_features = train_df.shape[1] - 1
target_column = "prognosis"
# # Encoder our prognosis labels as integers for easier decoding later
enc = OrdinalEncoder()
train_df[target_column] = enc.fit_transform(train_df[[target_column]])
train_df.head()
```

```
# convert all entries to int for faster processing
train_df = train_df.astype(int)
```

- Add logical combinations (AND, OR, XOR) of the features to the data set:

```
from spotPython.utils.convert import add_logical_columns
df_new = train_df.copy()
# save the target column using "target_column" as the column name
target = train_df[target_column]
# remove the target column
df_new = df_new.drop(columns=[target_column])
train_df = add_logical_columns(df_new)
# add the target column back
train_df[target_column] = target
train_df = train_df.astype(int)
train_df.head()
```

	sudden_fever	headache	mouth_bleed	nose_bleed	muscle_pain	joint_pain	vomiting	rash	diarrhea
0	1	1	0	1	1	1	1	0	1
1	0	0	0	0	0	0	1	0	1
2	0	1	1	1	0	1	1	1	1
3	0	0	1	1	1	1	0	1	0
4	0	0	0	0	0	0	0	0	1

```
from sklearn.model_selection import train_test_split
import numpy as np

n_samples = train_df.shape[0]
n_features = train_df.shape[1] - 1
train_df.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
train_df.head()
```

	x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	...	x6104	x6105	x6106	x6107	x6108	x6109	x6110
0	1	1	0	1	1	1	1	0	1	1	...	0	0	0	0	0	0	0
1	0	0	0	0	0	0	1	0	1	0	...	0	0	0	0	0	0	0
2	0	1	1	1	0	1	1	1	1	1	...	1	1	0	1	1	0	1
3	0	0	1	1	1	1	0	1	0	1	...	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	1	0	...	0	1	1	0	1	1	0

20.3.2 Check content of the target column

```
train_df[target_column].head()
```

```
0    3
1    7
2    3
3   10
4    6
Name: prognosis, dtype: int64
```

```
X_train, X_test, y_train, y_test = train_test_split(train_df.drop(target_column, axis=1),
                                                    random_state=42,
                                                    test_size=0.25,
                                                    stratify=train_df[target_column])
trainset = pd.DataFrame(np.hstack((X_train, np.array(y_train).reshape(-1, 1))))
testset = pd.DataFrame(np.hstack((X_test, np.array(y_test).reshape(-1, 1))))
trainset.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
testset.columns = [f"x{i}" for i in range(1, n_features+1)] + [target_column]
print(train_df.shape)
print(trainset.shape)
print(testset.shape)
```

```
(707, 6113)
```

```
(530, 6113)
```

```
(177, 6113)
```

```
import torch
from sklearn.model_selection import train_test_split
from spotPython.torch.dataframedataset import DataFrameDataset
dtype_x = torch.float32
dtype_y = torch.long
train_df = DataFrameDataset(train_df, target_column=target_column, dtype_x=dtype_x, dtype_y=dtype_y)
train = DataFrameDataset(trainset, target_column=target_column, dtype_x=dtype_x, dtype_y=dtype_y)
test = DataFrameDataset(testset, target_column=target_column, dtype_x=dtype_x, dtype_y=dtype_y)
n_samples = len(train)
```

```
# add the dataset to the fun_control
fun_control.update({"data": train_df, # full dataset,
                  "train": train,
                  "test": test,
                  "n_samples": n_samples,
                  "target_column": target_column})
```

20.4 Step 4: Specification of the Preprocessing Model

After the training and test data are specified and added to the `fun_control` dictionary, `spotPython` allows the specification of a data preprocessing pipeline, e.g., for the scaling of the data or for the one-hot encoding of categorical variables, see Section 14.4. This feature is not used here, so we do not change the default value (which is `None`).

20.5 Step 5: Select algorithm and `core_model_hyper_dict`

20.5.1 Implementing a Configurable Neural Network With `spotPython`

`spotPython` includes the `Net_vbdp` class which is implemented in the file `netvbdp.py`. The class is imported here.

This class inherits from the class `Net_Core` which is implemented in the file `netcore.py`, see Section 14.5.1.

20.5.2 Add the NN Model to the `fun_control` Dictionary

```
from spotPython.torch.netvbdp import Net_vbdp
from spotPython.data.torch_hyper_dict import TorchHyperDict
from spotPython.hyperparameters.values import add_core_model_to_fun_control
fun_control = add_core_model_to_fun_control(core_model=Net_vbdp,
                                          fun_control=fun_control,
                                          hyper_dict=TorchHyperDict)
```

The corresponding entries for the `core_model` class are shown below.

```
fun_control['core_model_hyper_dict']
```

```
{'_L0': {'type': 'int',
        'default': 64,
        'transform': 'None',
        'lower': 64,
        'upper': 64},
 'l1': {'type': 'int',
        'default': 8,
        'transform': 'transform_power_2_int',
        'lower': 8,
        'upper': 16},
 'dropout_prob': {'type': 'float',
                  'default': 0.01,
                  'transform': 'None',
                  'lower': 0.0,
                  'upper': 0.9},
 'lr_mult': {'type': 'float',
             'default': 1.0,
             'transform': 'None',
             'lower': 0.1,
             'upper': 10.0},
 'batch_size': {'type': 'int',
                'default': 4,
                'transform': 'transform_power_2_int',
                'lower': 1,
                'upper': 4},
 'epochs': {'type': 'int',
            'default': 4,
            'transform': 'transform_power_2_int',
            'lower': 4,
            'upper': 9},
 'k_folds': {'type': 'int',
             'default': 1,
             'transform': 'None',
             'lower': 1,
             'upper': 1},
 'patience': {'type': 'int',
              'default': 2,
              'transform': 'transform_power_2_int',
              'lower': 1,
              'upper': 5},
 'optimizer': {'levels': ['Adadelata',
                          'Adagrad',
                          'Adam']},
```

```

'AdamW',
'SparseAdam',
'Adamax',
'ASGD',
'NAdam',
'RAdam',
'RMSprop',
'Rprop',
'SGD'],
'type': 'factor',
'default': 'SGD',
'transform': 'None',
'class_name': 'torch.optim',
'core_model_parameter_type': 'str',
'lower': 0,
'upper': 12},
'sgd_momentum': {'type': 'float',
'default': 0.0,
'transform': 'None',
'lower': 0.0,
'upper': 1.0}}

```

20.6 Step 6: Modify hyper_dict Hyperparameters for the Selected Algorithm aka core_model

spotPython provides functions for modifying the hyperparameters, their bounds and factors as well as for activating and de-activating hyperparameters without re-compilation of the Python source code. These functions were described in Section [14.6](#).

 Caution: Small number of epochs for demonstration purposes

- epochs and patience are set to small values for demonstration purposes. These values are too small for a real application.
- More reasonable values are, e.g.:
 - fun_control = modify_hyper_parameter_bounds(fun_control, "epochs", bounds=[7, 9]) and
 - fun_control = modify_hyper_parameter_bounds(fun_control, "patience", bounds=[2, 7])

```

from spotPython.hyperparameters.values import modify_hyper_parameter_bounds

fun_control = modify_hyper_parameter_bounds(fun_control, "_L0", bounds=[n_features, n_feat
fun_control = modify_hyper_parameter_bounds(fun_control, "l1", bounds=[6, 13])
fun_control = modify_hyper_parameter_bounds(fun_control, "epochs", bounds=[2, 3])
fun_control = modify_hyper_parameter_bounds(fun_control, "patience", bounds=[2, 2])

from spotPython.hyperparameters.values import modify_hyper_parameter_levels
fun_control = modify_hyper_parameter_levels(fun_control, "optimizer", ["Adam", "AdamW", "Ad
# fun_control = modify_hyper_parameter_levels(fun_control, "optimizer", ["Adam"])
# fun_control["core_model_hyper_dict"]

```

20.6.1 Optimizers

Optimizers are described in Section [14.6.1](#).

```

fun_control = modify_hyper_parameter_bounds(fun_control,
    "lr_mult", bounds=[1e-3, 1e-3])
fun_control = modify_hyper_parameter_bounds(fun_control,
    "sgd_momentum", bounds=[0.9, 0.9])

```

20.7 Step 7: Selection of the Objective (Loss) Function

20.7.1 Evaluation

The evaluation procedure requires the specification of two elements:

1. the way how the data is split into a train and a test set (see Section [14.7.1](#))
2. the loss function (and a metric).

20.7.2 Loss Functions and Metrics

The loss function is specified by the key "loss_function". We will use CrossEntropy loss for the multiclass-classification task.

```

from torch.nn import CrossEntropyLoss
loss_function = CrossEntropyLoss()
fun_control.update({"loss_function": loss_function})

```

20.7.3 Metric

- We will use the MAP@k metric for the evaluation of the model. Here is an example how this metric is calculated.

```
from spotPython.torch.mapk import MAPK
import torch
mapk = MAPK(k=2)
target = torch.tensor([0, 1, 2, 2])
preds = torch.tensor(
    [
        [0.5, 0.2, 0.2], # 0 is in top 2
        [0.3, 0.4, 0.2], # 1 is in top 2
        [0.2, 0.4, 0.3], # 2 is in top 2
        [0.7, 0.2, 0.1], # 2 isn't in top 2
    ]
)
mapk.update(preds, target)
print(mapk.compute()) # tensor(0.6250)
```

tensor(0.6250)

```
from spotPython.torch.mapk import MAPK
import torchmetrics
metric_torch = MAPK(k=3)
fun_control.update({"metric_torch": metric_torch})
```

20.8 Step 8: Calling the SPOT Function

20.8.1 Preparing the SPOT Call

The following code passes the information about the parameter ranges and bounds to `spot`.

```
# extract the variable types, names, and bounds
from spotPython.hyperparameters.values import (get_bound_values,
        get_var_name,
        get_var_type,)
var_type = get_var_type(fun_control)
var_name = get_var_name(fun_control)
```

```

fun_control.update({"var_type": var_type,
                  "var_name": var_name})
lower = get_bound_values(fun_control, "lower")
upper = get_bound_values(fun_control, "upper")

```

Now, the dictionary `fun_control` contains all information needed for the hyperparameter tuning. Before the hyperparameter tuning is started, it is recommended to take a look at the experimental design. The method `gen_design_table` generates a design table as follows:

```

from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control))

```

name	type	default	lower	upper	transform
_L0	int	64	6112	6112	None
l1	int	8	6	13	transform_power_2_int
dropout_prob	float	0.01	0	0.9	None
lr_mult	float	1.0	0.001	0.001	None
batch_size	int	4	1	4	transform_power_2_int
epochs	int	4	2	3	transform_power_2_int
k_folds	int	1	1	1	None
patience	int	2	2	2	transform_power_2_int
optimizer	factor	SGD	0	3	None
sgd_momentum	float	0.0	0.9	0.9	None

This allows to check if all information is available and if the information is correct.

20.8.2 The Objective Function `fun_torch`

The objective function `fun_torch` is selected next. It implements an interface from PyTorch's training, validation, and testing methods to `spotPython`.

```

from spotPython.fun.hypertorch import HyperTorch
fun = HyperTorch().fun_torch

```

```

from spotPython.hyperparameters.values import get_default_hyperparameters_as_array
hyper_dict=TorCHHyperDict().load()
X_start = get_default_hyperparameters_as_array(fun_control, hyper_dict)

```

20.8.3 Starting the Hyperparameter Tuning

The spotPython hyperparameter tuning is started by calling the Spot function as described in Section 14.8.4.

```
import numpy as np
from spotPython.spot import spot
from math import inf
spot_tuner = spot.Spot(fun=fun,
                      lower = lower,
                      upper = upper,
                      fun_evals = inf,
                      fun_repeats = 1,
                      max_time = MAX_TIME,
                      noise = False,
                      tolerance_x = np.sqrt(np.spacing(1)),
                      var_type = var_type,
                      var_name = var_name,
                      infill_criterion = "y",
                      n_points = 1,
                      seed=123,
                      log_level = 50,
                      show_models= False,
                      show_progress= True,
                      fun_control = fun_control,
                      design_control={"init_size": INIT_SIZE,
                                     "repeats": 1},
                      surrogate_control={"noise": True,
                                       "cod_type": "norm",
                                       "min_theta": -4,
                                       "max_theta": 3,
                                       "n_theta": len(var_name),
                                       "model_fun_evals": 10_000,
                                       "log_level": 50
                                       })

spot_tuner.run(X_start=X_start)
```

```
config: {'_L0': 6112, 'l1': 2048, 'dropout_prob': 0.17031221661559992, 'lr_mult': 0.001, 'ba
Epoch: 1 |
```

```
MAPK: 0.1778273731470108 | Loss: 2.3978227717535838 | Acc: 0.1132075471698113.
```

Epoch: 2 |

MAPK: 0.1755952388048172 | Loss: 2.3978894608361379 | Acc: 0.1084905660377359.

Epoch: 3 |

MAPK: 0.1845238059759140 | Loss: 2.3977681398391724 | Acc: 0.1273584905660377.

Epoch: 4 |

MAPK: 0.1889880895614624 | Loss: 2.3977671180452620 | Acc: 0.1367924528301887.

Epoch: 5 | MAPK: 0.1956845223903656 | Loss: 2.3977283069065640 | Acc: 0.1650943396226415.

Epoch: 6 |

MAPK: 0.1934523731470108 | Loss: 2.3976699795041765 | Acc: 0.1320754716981132.

Epoch: 7 |

MAPK: 0.1867559403181076 | Loss: 2.3975695712225780 | Acc: 0.1320754716981132.

Epoch: 8 |

MAPK: 0.1912202388048172 | Loss: 2.3975362437111989 | Acc: 0.1273584905660377.

Returned to Spot: Validation loss: 2.397536243711199

config: {'_L0': 6112, 'l1': 256, 'dropout_prob': 0.19379790035512987, 'lr_mult': 0.001, 'bat

Epoch: 1 |

MAPK: 0.1589506417512894 | Loss: 2.3977898138540761 | Acc: 0.0801886792452830.

Epoch: 2 |

MAPK: 0.1635802537202835 | Loss: 2.3977428807152643 | Acc: 0.0801886792452830.

Epoch: 3 |

MAPK: 0.1620370596647263 | Loss: 2.3977394987035683 | Acc: 0.0707547169811321.

Epoch: 4 |

MAPK: 0.1658950746059418 | Loss: 2.3977184383957475 | Acc: 0.0801886792452830.

Returned to Spot: Validation loss: 2.3977184383957475

config: {'_L0': 6112, 'l1': 4096, 'dropout_prob': 0.6759063718076167, 'lr_mult': 0.001, 'bat

Epoch: 1 |

MAPK: 0.1871068924665451 | Loss: 2.3979439443012454 | Acc: 0.1179245283018868.
Epoch: 2 |

MAPK: 0.2421383112668991 | Loss: 2.3975673734017140 | Acc: 0.1415094339622641.
Epoch: 3 |

MAPK: 0.2531446516513824 | Loss: 2.3974374240299441 | Acc: 0.1462264150943396.
Epoch: 4 |

MAPK: 0.2303458899259567 | Loss: 2.3974605888690590 | Acc: 0.1320754716981132.
Epoch: 5 |

MAPK: 0.2287735193967819 | Loss: 2.3970276544678888 | Acc: 0.1226415094339623.
Epoch: 6 |

MAPK: 0.2264150828123093 | Loss: 2.3968495602877633 | Acc: 0.1320754716981132.
Epoch: 7 |

MAPK: 0.2091194838285446 | Loss: 2.3965910470710612 | Acc: 0.1084905660377359.
Epoch: 8 |

MAPK: 0.2224842458963394 | Loss: 2.3959322245615833 | Acc: 0.1320754716981132.
Returned to Spot: Validation loss: 2.3959322245615833

config: {'_L0': 6112, 'l1': 128, 'dropout_prob': 0.37306669346546995, 'lr_mult': 0.001, 'batch_size': 128}
Epoch: 1 |

MAPK: 0.1588050574064255 | Loss: 2.3982574624835320 | Acc: 0.0896226415094340.
Epoch: 2 |

MAPK: 0.1588050574064255 | Loss: 2.3983343322322055 | Acc: 0.0849056603773585.
Epoch: 3 |

MAPK: 0.1588050425052643 | Loss: 2.3981994233041442 | Acc: 0.0849056603773585.
Epoch: 4 |

MAPK: 0.1572327017784119 | Loss: 2.3983834644533553 | Acc: 0.0801886792452830.
Returned to Spot: Validation loss: 2.3983834644533553

config: {'_L0': 6112, 'l1': 1024, 'dropout_prob': 0.870137281216666, 'lr_mult': 0.001, 'batch_size': 128}
Epoch: 1 |

MAPK: 0.1766975373029709 | Loss: 2.3976071499012135 | Acc: 0.0990566037735849.
Epoch: 2 |

MAPK: 0.1929012537002563 | Loss: 2.3974499084331371 | Acc: 0.1084905660377359.
Epoch: 3 |

MAPK: 0.2060185223817825 | Loss: 2.3974460937358715 | Acc: 0.1367924528301887.
Epoch: 4 |

MAPK: 0.2183642238378525 | Loss: 2.3971893434171325 | Acc: 0.1084905660377359.
Epoch: 5 |

MAPK: 0.2052469253540039 | Loss: 2.3975013008824102 | Acc: 0.1084905660377359.
Epoch: 6 |

MAPK: 0.1867283880710602 | Loss: 2.3977151093659579 | Acc: 0.1320754716981132.
Epoch: 7 |

MAPK: 0.2160493582487106 | Loss: 2.3974010591153747 | Acc: 0.1320754716981132.
Epoch: 8 |

MAPK: 0.2083333283662796 | Loss: 2.3972757657368979 | Acc: 0.1226415094339623.
Early stopping at epoch 7
Returned to Spot: Validation loss: 2.397275765736898

config: {'_L0': 6112, 'l1': 4096, 'dropout_prob': 0.4132005099912892, 'lr_mult': 0.001, 'bat
Epoch: 1 |

MAPK: 0.1800314486026764 | Loss: 2.3976828557140424 | Acc: 0.0943396226415094.
Epoch: 2 |

MAPK: 0.2114779800176620 | Loss: 2.3973630531778873 | Acc: 0.1132075471698113.
Epoch: 3 |

MAPK: 0.2311320602893829 | Loss: 2.3970617730662509 | Acc: 0.1320754716981132.
Epoch: 4 |

MAPK: 0.2295597791671753 | Loss: 2.3964778162398428 | Acc: 0.1320754716981132.
Epoch: 5 |

MAPK: 0.2507861554622650 | Loss: 2.3957185700254620 | Acc: 0.1745283018867924.
Epoch: 6 |

MAPK: 0.2421383261680603 | Loss: 2.3945849076756893 | Acc: 0.1792452830188679.
Epoch: 7 |

MAPK: 0.2319181859493256 | Loss: 2.3930445887007803 | Acc: 0.1650943396226415.
Epoch: 8 |

MAPK: 0.2405660003423691 | Loss: 2.3900625323349574 | Acc: 0.1792452830188679.
Returned to Spot: Validation loss: 2.3900625323349574
spotPython tuning: 2.3900625323349574 [####-----] 35.87%

config: {'_L0': 6112, 'l1': 512, 'dropout_prob': 0.41307021636436325, 'lr_mult': 0.001, 'bat
Epoch: 1 |

MAPK: 0.2106918543577194 | Loss: 2.3969543834902205 | Acc: 0.1273584905660377.
Epoch: 2 |

MAPK: 0.2209119200706482 | Loss: 2.3968141033964336 | Acc: 0.1320754716981132.
Epoch: 3 |

MAPK: 0.2264150977134705 | Loss: 2.3968117731922076 | Acc: 0.1462264150943396.
Epoch: 4 |

MAPK: 0.2382074892520905 | Loss: 2.3967072468883588 | Acc: 0.1603773584905660.
Epoch: 5 |

MAPK: 0.2499999403953552 | Loss: 2.3965636019436820 | Acc: 0.1745283018867924.
Epoch: 6 |

MAPK: 0.2460691630840302 | Loss: 2.3965558735829480 | Acc: 0.1698113207547170.
Epoch: 7 |

MAPK: 0.2555030882358551 | Loss: 2.3965241774073185 | Acc: 0.1745283018867924.
Epoch: 8 |

MAPK: 0.2555031180381775 | Loss: 2.3964265607437998 | Acc: 0.1886792452830189.
Returned to Spot: Validation loss: 2.3964265607438
spotPython tuning: 2.3900625323349574 [#####-----] 46.13%

config: {'_L0': 6112, 'l1': 128, 'dropout_prob': 0.41300574671604434, 'lr_mult': 0.001, 'bat
Epoch: 1 |

MAPK: 0.1462264358997345 | Loss: 2.3985383555574238 | Acc: 0.0613207547169811.
Epoch: 2 |

MAPK: 0.1525157392024994 | Loss: 2.3986691416434520 | Acc: 0.0613207547169811.
Epoch: 3 |

MAPK: 0.1477987617254257 | Loss: 2.3986250297078548 | Acc: 0.0613207547169811.
Epoch: 4 |

MAPK: 0.1477987766265869 | Loss: 2.3985712978075133 | Acc: 0.0613207547169811.
Epoch: 5 |

MAPK: 0.1509434133768082 | Loss: 2.3984091821706519 | Acc: 0.0613207547169811.
Epoch: 6 |

MAPK: 0.1525157392024994 | Loss: 2.3983121615535810 | Acc: 0.0613207547169811.
Epoch: 7 |

MAPK: 0.1556603908538818 | Loss: 2.3983128160800575 | Acc: 0.0613207547169811.
Epoch: 8 |

MAPK: 0.1509434282779694 | Loss: 2.3983534754447215 | Acc: 0.0613207547169811.
Returned to Spot: Validation loss: 2.3983534754447215
spotPython tuning: 2.3900625323349574 [#####----] 63.56%

config: {'_L0': 6112, 'l1': 8192, 'dropout_prob': 0.41444870177080323, 'lr_mult': 0.001, 'ba

Epoch: 1 |

MAPK: 0.1973270177841187 | Loss: 2.3975810734730847 | Acc: 0.1132075471698113.
Epoch: 2 |

MAPK: 0.2287735342979431 | Loss: 2.3971492974263318 | Acc: 0.1415094339622641.
Epoch: 3 |

MAPK: 0.1973270773887634 | Loss: 2.3966462432213551 | Acc: 0.1179245283018868.
Epoch: 4 |

MAPK: 0.1918238997459412 | Loss: 2.3957121237268986 | Acc: 0.1132075471698113.
Returned to Spot: Validation loss: 2.3957121237268986
spotPython tuning: 2.3900625323349574 [#####--] 83.85%

config: {'_L0': 6112, 'l1': 4096, 'dropout_prob': 0.41474501146092996, 'lr_mult': 0.001, 'ba

Epoch: 1 |

MAPK: 0.1604938507080078 | Loss: 2.3979456159803600 | Acc: 0.0801886792452830.
Epoch: 2 |

MAPK: 0.2060185074806213 | Loss: 2.3978245876453541 | Acc: 0.1179245283018868.
Epoch: 3 |

MAPK: 0.2037036716938019 | Loss: 2.3978362613254123 | Acc: 0.1320754716981132.
Epoch: 4 |

MAPK: 0.1844135522842407 | Loss: 2.3978672910619667 | Acc: 0.1084905660377359.
Epoch: 5 |

MAPK: 0.1944444328546524 | Loss: 2.3977656982563160 | Acc: 0.1132075471698113.
Epoch: 6 |

MAPK: 0.2322530597448349 | Loss: 2.3975858246838606 | Acc: 0.1603773584905660.
Epoch: 7 |

MAPK: 0.2075617164373398 | Loss: 2.3976335437209517 | Acc: 0.1179245283018868.
Epoch: 8 |

```
MAPK: 0.2060184776782990 | Loss: 2.3975747161441379 | Acc: 0.1226415094339623.  
Returned to Spot: Validation loss: 2.397574716144138  
spotPython tuning: 2.3900625323349574 [#####-] 92.60%
```

```
config: {'_L0': 6112, 'l1': 4096, 'dropout_prob': 0.41094634811682174, 'lr_mult': 0.001, 'ba  
Epoch: 1 |
```

```
MAPK: 0.2044025063514709 | Loss: 2.3975967506192766 | Acc: 0.1462264150943396.  
Epoch: 2 |
```

```
MAPK: 0.1926100701093674 | Loss: 2.3974419134967730 | Acc: 0.1084905660377359.  
Epoch: 3 |
```

```
MAPK: 0.1957547366619110 | Loss: 2.3972251460237324 | Acc: 0.0943396226415094.  
Epoch: 4 |
```

```
MAPK: 0.2106918096542358 | Loss: 2.3968959259537033 | Acc: 0.0943396226415094.  
Epoch: 5 |
```

```
MAPK: 0.2059748619794846 | Loss: 2.3968038244067498 | Acc: 0.0896226415094340.  
Epoch: 6 |
```

```
MAPK: 0.1933962404727936 | Loss: 2.3964082070116728 | Acc: 0.0896226415094340.  
Epoch: 7 |
```

```
MAPK: 0.1926100403070450 | Loss: 2.3958087822176375 | Acc: 0.0896226415094340.  
Epoch: 8 |
```

```
MAPK: 0.1745283007621765 | Loss: 2.3954800929663316 | Acc: 0.0896226415094340.  
Returned to Spot: Validation loss: 2.3954800929663316  
spotPython tuning: 2.3900625323349574 [#####] 100.00% Done...
```

```
<spotPython.spot.spot.Spot at 0x2c97937f0>
```

20.9 Step 9: Tensorboard

The textual output shown in the console (or code cell) can be visualized with Tensorboard as described in Section 14.9, see also the description in the documentation: [Tensorboard](#).

20.10 Step 10: Results

After the hyperparameter tuning run is finished, the results can be analyzed as described in Section 14.10.

```
spot_tuner.plot_progress(log_y=False,
    filename="./figures/" + experiment_name+"_progress.png")
```




Figure 20.1: Progress plot. *Black* dots denote results from the initial design. *Red* dots illustrate the improvement found by the surrogate model based optimization.

```
from spotPython.utils.eda import gen_design_table
print(gen_design_table(fun_control=fun_control, spot=spot_tuner))
```

name	type	default	lower	upper	tuned	transform
_L0	int	64	6112.0	6112.0	6112.0	None
l1	int	8	6.0	13.0	12.0	transform_pow
dropout_prob	float	0.01	0.0	0.9	0.4132005099912892	None
lr_mult	float	1.0	0.001	0.001	0.001	None
batch_size	int	4	1.0	4.0	1.0	transform_pow
epochs	int	4	2.0	3.0	3.0	transform_pow
k_folds	int	1	1.0	1.0	1.0	None
patience	int	2	2.0	2.0	2.0	transform_pow
optimizer	factor	SGD	0.0	3.0	3.0	None
sgd_momentum	float	0.0	0.9	0.9	0.9	None

```
spot_tuner.plot_importance(threshold=0.025,
                           filename="./figures/" + experiment_name+"_importance.png")
```




Figure 20.2: Variable importance plot, threshold 0.025.

20.10.1 Get the Tuned Architecture

```
from spotPython.hyperparameters.values import get_one_core_model_from_X
X = spot_tuner.to_all_dim(spot_tuner.min_X.reshape(1,-1))
model_spot = get_one_core_model_from_X(X, fun_control)
model_spot
```

```
Net_vbdp(
  (fc1): Linear(in_features=6112, out_features=4096, bias=True)
  (fc2): Linear(in_features=4096, out_features=2048, bias=True)
  (fc3): Linear(in_features=2048, out_features=1024, bias=True)
  (fc4): Linear(in_features=1024, out_features=512, bias=True)
  (fc5): Linear(in_features=512, out_features=11, bias=True)
  (relu): ReLU()
  (softmax): Softmax(dim=1)
  (dropout1): Dropout(p=0.4132005099912892, inplace=False)
  (dropout2): Dropout(p=0.2066002549956446, inplace=False)
)
```

20.10.2 Evaluation of the Tuned Architecture

```
from spotPython.torch.traintest import (  
    train_tuned,  
    test_tuned,  
)  
train_tuned(net=model_spot, train_dataset=train,  
            loss_function=fun_control["loss_function"],  
            metric=fun_control["metric_torch"],  
            shuffle=True,  
            device = fun_control["device"],  
            path=None,  
            task=fun_control["task"],)
```

Epoch: 1 |

MAPK: 0.1753144711256027 | Loss: 2.3976820324951746 | Acc: 0.1084905660377359.

Epoch: 2 |

MAPK: 0.1753144562244415 | Loss: 2.3973140176737084 | Acc: 0.1132075471698113.

Epoch: 3 |

MAPK: 0.1816037595272064 | Loss: 2.3969395947906205 | Acc: 0.1179245283018868.

Epoch: 4 |

MAPK: 0.1816037744283676 | Loss: 2.3961891678144349 | Acc: 0.1179245283018868.

Epoch: 5 |

MAPK: 0.1823899447917938 | Loss: 2.3951856910057789 | Acc: 0.1179245283018868.

Epoch: 6 |

MAPK: 0.1784591376781464 | Loss: 2.3940786213245033 | Acc: 0.1179245283018868.

Epoch: 7 |

MAPK: 0.1800314337015152 | Loss: 2.3925867305611663 | Acc: 0.1179245283018868.

Epoch: 8 |

MAPK: 0.1816037744283676 | Loss: 2.3907402348968216 | Acc: 0.1179245283018868.

Returned to Spot: Validation loss: 2.3907402348968216

If `path` is set to a filename, e.g., `path = "model_spot_trained.pt"`, the weights of the trained model will be loaded from this file.

```
test_tuned(net=model_spot, test_dataset=test,
           shuffle=False,
           loss_function=fun_control["loss_function"],
           metric=fun_control["metric_torch"],
           device = fun_control["device"],
           task=fun_control["task"],)
```

```
MAPK: 0.1910112500190735 | Loss: 2.3872228097379877 | Acc: 0.1186440677966102.
Final evaluation: Validation loss: 2.3872228097379877
Final evaluation: Validation metric: 0.1910112500190735
-----
```

```
(2.3872228097379877, nan, tensor(0.1910))
```

20.10.3 Cross-validated Evaluations

- This is the evaluation that will be used in the comparison.

Caution: Cross-validated Evaluations

- The number of folds is set to 1 by default.
- Here it was changed to 3 for demonstration purposes.
- Set the number of folds to a reasonable value, e.g., 10.
- This can be done by setting the `k_folds` attribute of the model as follows:
- `setattr(model_spot, "k_folds", 10)`

```
from spotPython.torch.traintest import evaluate_cv
# modify k-folds:
setattr(model_spot, "k_folds", 3)
df_eval, df_preds, df_metrics = evaluate_cv(net=model_spot,
      dataset=fun_control["data"],
      loss_function=fun_control["loss_function"],
      metric=fun_control["metric_torch"],
      task=fun_control["task"],
      writer=fun_control["writer"],
      writerId="model_spot_cv",
      device = fun_control["device"])
```

Fold: 1
Epoch: 1 |

MAPK: 0.1744349896907806 | Loss: 2.3973758847026501 | Acc: 0.0805084745762712.
Epoch: 2 |

MAPK: 0.2146892547607422 | Loss: 2.3967134689880631 | Acc: 0.1483050847457627.
Epoch: 3 |

MAPK: 0.2161016911268234 | Loss: 2.3958126387353671 | Acc: 0.1271186440677966.
Epoch: 4 |

MAPK: 0.2266948670148849 | Loss: 2.3945004879418068 | Acc: 0.1271186440677966.
Epoch: 5 |

MAPK: 0.2457626760005951 | Loss: 2.3925800707380649 | Acc: 0.1271186440677966.
Epoch: 6 |

MAPK: 0.2598869800567627 | Loss: 2.3892519857923862 | Acc: 0.1271186440677966.
Epoch: 7 |

MAPK: 0.2853107452392578 | Loss: 2.3845437946966137 | Acc: 0.1313559322033898.
Epoch: 8 |

MAPK: 0.3227401077747345 | Loss: 2.3799093315156838 | Acc: 0.1949152542372881.
Fold: 2
Epoch: 1 |

MAPK: 0.1871468871831894 | Loss: 2.3975351159855469 | Acc: 0.0889830508474576.
Epoch: 2 |

MAPK: 0.2281073182821274 | Loss: 2.3968891127634855 | Acc: 0.1186440677966102.
Epoch: 3 |

MAPK: 0.2358756810426712 | Loss: 2.3959750260336925 | Acc: 0.1186440677966102.
Epoch: 4 |

MAPK: 0.2337570190429688 | Loss: 2.3945896524493979 | Acc: 0.1186440677966102.
Epoch: 5 |

MAPK: 0.2337570339441299 | Loss: 2.3920788482084112 | Acc: 0.1228813559322034.
Epoch: 6 |

MAPK: 0.2302259504795074 | Loss: 2.3886764635473994 | Acc: 0.1186440677966102.
Epoch: 7 |

MAPK: 0.2422316074371338 | Loss: 2.3847005751173374 | Acc: 0.1228813559322034.
Epoch: 8 |

MAPK: 0.2824858725070953 | Loss: 2.3816417132393788 | Acc: 0.1737288135593220.
Fold: 3
Epoch: 1 |

MAPK: 0.1829096078872681 | Loss: 2.3976690526735984 | Acc: 0.0978723404255319.
Epoch: 2 |

MAPK: 0.1956214606761932 | Loss: 2.3971529390852329 | Acc: 0.1234042553191489.
Epoch: 3 |

MAPK: 0.2026835680007935 | Loss: 2.3964017508393627 | Acc: 0.1191489361702128.
Epoch: 4 |

MAPK: 0.1991525441408157 | Loss: 2.3954021526595293 | Acc: 0.1148936170212766.
Epoch: 5 |

MAPK: 0.2083333134651184 | Loss: 2.3936718278012035 | Acc: 0.1148936170212766.
Epoch: 6 |

MAPK: 0.2111581861972809 | Loss: 2.3913376068664811 | Acc: 0.1148936170212766.
Epoch: 7 |

MAPK: 0.2139830440282822 | Loss: 2.3877145233800854 | Acc: 0.1148936170212766.
Epoch: 8 |

MAPK: 0.2351694703102112 | Loss: 2.3849314491627580 | Acc: 0.1148936170212766.

```
metric_name = type(fun_control["metric_torch"]).__name__  
print(f"loss: {df_eval}, Cross-validated {metric_name}: {df_metrics}")
```

loss: 2.3821608313059404, Cross-validated MAPK: 0.28013181686401367

20.10.4 Detailed Hyperparameter Plots

```
filename = "./figures/" + experiment_name  
spot_tuner.plot_important_hyperparameter_contour(filename=filename)
```

```
l1: 21.13433159251595  
dropout_prob: 100.0  
batch_size: 87.40863288608274  
optimizer: 81.7067553304582
```













Figure 20.3: Contour plots.







20.10.5 Parallel Coordinates Plot

```
spot_tuner.parallel_plot()
```

Unable to display output for mime type(s): text/html

Parallel coordinates plots

Unable to display output for mime type(s): text/html

```
# close tensorboard writer
if fun_control["writer"] is not None:
    fun_control["writer"].close()
```

20.10.6 Plot all Combinations of Hyperparameters

- Warning: this may take a while.

```
PLOT_ALL = False
if PLOT_ALL:
    n = spot_tuner.k
    for i in range(n-1):
        for j in range(i+1, n):
            spot_tuner.plot_contour(i=i, j=j, min_z=min_z, max_z = max_z)
```

21 Documentation of the Sequential Parameter Optimization

This document describes the Spot features.

21.1 Example: spot

```
import numpy as np
from math import inf
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
from scipy.optimize import shgo
from scipy.optimize import direct
from scipy.optimize import differential_evolution
import matplotlib.pyplot as plt
```


21.1.1 The Objective Function

The spotPython package provides several classes of objective functions. We will use an analytical objective function, i.e., a function that can be described by a (closed) formula:

$$f(x) = x^2$$

```
fun = analytical().fun_sphere

x = np.linspace(-1,1,100).reshape(-1,1)
y = fun(x)
plt.figure()
plt.plot(x,y, "k")
plt.show()
```



```
spot_1 = spot.Spot(fun=fun,
                  lower = np.array([-10]),
                  upper = np.array([100]),
                  fun_evals = 7,
                  fun_repeats = 1,
                  max_time = inf,
                  noise = False,
                  tolerance_x = np.sqrt(np.spacing(1)),
                  var_type=["num"],
                  infill_criterion = "y",
                  n_points = 1,
                  seed=123,
                  log_level = 50,
                  show_models=True,
                  fun_control = {},
                  design_control={"init_size": 5,
                                "repeats": 1},
                  surrogate_control={"noise": False,
                                   "cod_type": "norm",
                                   "min_theta": -4,
                                   "max_theta": 3,
                                   "n_theta": 1,
                                   "model_optimizer": differential_evolution,
                                   "model_fun_evals": 1000,
```

})

spot's `__init__` method sets the control parameters. There are two parameter groups:

1. external parameters can be specified by the user
2. internal parameters, which are handled by `spot`.

21.1.2 External Parameters

external parameter	type	description	default	mandatory
<code>fun</code>	object	objective function		yes
<code>lower</code>	array	lower bound		yes
<code>upper</code>	array	upper bound		yes
<code>fun_evals</code>	int	number of function evaluations	15	no
<code>fun_evals</code>	int	number of function evaluations	15	no
<code>fun_control</code>	dict	noise etc.	{}	n
<code>max_time</code>	int	max run time budget	<code>inf</code>	no
<code>noise</code>	bool	if repeated evaluations of <code>fun</code> results in different values, then <code>noise</code> should be set to <code>True</code> .	<code>False</code>	no

external parameter	type	description	default	mandatory
<code>tolerance_x</code>	float	tolerance for new x solutions. Minimum distance of new solutions, generated by <code>suggest_new_X</code> , to already existing solutions. If zero (which is the default), every new solution is accepted.	0	no
<code>var_type</code>	list	list of type information, can be either "num" or "factor"	["num"]	no
<code>infill_criterion</code>	string	Can be "y", "s", "ei" (negative expected improvement), or "all"	"y"	no
<code>n_points</code>	int	number of infill points	1	no
<code>seed</code>	int	initial seed. If <code>Spot.run()</code> is called twice, different results will be generated. To reproduce results, the <code>seed</code> can be used.	123	no

external parameter	type	description	default	mandatory
log_level	int	log level with the following settings: NOTSET (0), DEBUG (10: Detailed information, typically of interest only when diagnosing problems.), INFO (20: Confirmation that things are working as expected.), WARNING (30: An indication that something unexpected happened, or indicative of some problem in the near future (e.g. 'disk space low'). The software is still working as expected.), ERROR (40: Due to a more serious problem, the software has not been able to perform some function.), and CRITICAL (50: A serious error, indicating that the program itself may be unable to continue running.)	50	no

external parameter	type	description	default	mandatory
<code>show_models</code>	bool	Plot model. Currently only 1-dim functions are supported	<code>False</code>	no
<code>design</code>	object	experimental design	<code>None</code>	no
<code>design_control</code>	dict	control parameters	see below	no
<code>surrogate</code>		surrogate model	<code>kriging</code>	no
<code>surrogate_control</code>	dict	control parameters	see below	no
<code>optimizer</code>	object	optimizer	see below	no
<code>optimizer_control</code>	dict	control parameters	see below	no

- Besides these single parameters, the following parameter dictionaries can be specified by the user:

- `fun_control`
- `design_control`
- `surrogate_control`
- `optimizer_control`

21.2 The `fun_control` Dictionary

external parameter	type	description	default	mandatory
<code>sigma</code>	float	noise: standard deviation	<code>0</code>	yes
<code>seed</code>	int	seed for rng	<code>124</code>	yes

21.3 The `design_control` Dictionary

external parameter	type	description	default	mandatory
<code>init_size</code>	int	initial sample size	<code>10</code>	yes

external parameter	type	description	default	mandatory
repeats	int	number of repeats of the initial sammples	1	yes

21.4 The surrogate_control Dictionary


external parameter	type	description	default	mandatory
noise				
model_optimizer	object	optimizer	differential_evolution	
model_fun_evals				
min_theta			-3.	
max_theta			3.	
n_theta			1	
n_p			1	
optim_p			False	
cod_type			"norm"	
var_type				
use_cod_y	bool		False	



21.5 The optimizer_control Dictionary

external parameter	type	description	default	mandatory
max_iter	int	max number of iterations. Note: these are the cheap evaluations on the surrogate.	1000	no

21.6 Run

```
spot_1.run()
```





<spotPython.spot.spot.Spot at 0x17ed8a9e0>

21.7 Print the Results

```
spot_1.print_results()
```


```
min y: 0.30656551286610595
```

```
x0: -0.5536835855126157
```

```
[['x0', -0.5536835855126157]]
```

21.8 Show the Progress

```
spot_1.plot_progress()
```




21.9 Visualize the Surrogate

- The plot method of the `kriging` surrogate is used.
- Note: the plot uses the interval defined by the ranges of the natural variables.

```
spot_1.surrogate.plot()
```

<Figure size 2700x1800 with 0 Axes>



21.10 Init: Build Initial Design

```

from spotPython.design.spacefilling import spacefilling
from spotPython.build.kriging import Kriging
from spotPython.fun.objectivefunctions import analytical
gen = spacefilling(2)
rng = np.random.RandomState(1)
lower = np.array([-5,-0])
upper = np.array([10,15])
fun = analytical().fun_branin
fun_control = {"sigma": 0,
               "seed": 123}

X = gen.scipy_lhd(10, lower=lower, upper = upper)
print(X)
y = fun(X, fun_control=fun_control)
print(y)

```

```

[[ 8.97647221 13.41926847]
 [ 0.66946019  1.22344228]
 [ 5.23614115 13.78185824]
 [ 5.6149825  11.5851384 ]

```

```

[-1.72963184  1.66516096]
[-4.26945568  7.1325531 ]
[ 1.26363761 10.17935555]
[ 2.88779942  8.05508969]
[-3.39111089  4.15213772]
[ 7.30131231  5.22275244]]
[128.95676449  31.73474356 172.89678121 126.71295908  64.34349975
 70.16178611  48.71407916  31.77322887  76.91788181  30.69410529]

```

21.11 Replicability

Seed

```

gen = spacefilling(2, seed=123)
X0 = gen.scipy_lhd(3)
gen = spacefilling(2, seed=345)
X1 = gen.scipy_lhd(3)
X2 = gen.scipy_lhd(3)
gen = spacefilling(2, seed=123)
X3 = gen.scipy_lhd(3)
X0, X1, X2, X3

```

```

(array([[0.77254938, 0.31539299],
        [0.59321338, 0.93854273],
        [0.27469803, 0.3959685 ]]),
 array([[0.78373509, 0.86811887],
        [0.06692621, 0.6058029 ],
        [0.41374778, 0.00525456]]),
 array([[0.121357  , 0.69043832],
        [0.41906219, 0.32838498],
        [0.86742658, 0.52910374]]),
 array([[0.77254938, 0.31539299],
        [0.59321338, 0.93854273],
        [0.27469803, 0.3959685 ]]))

```

21.12 Surrogates

21.12.1 A Simple Predictor

The code below shows how to use a simple model for prediction. Assume that only two (very costly) measurements are available:

1. $f(0) = 0.5$
2. $f(2) = 2.5$

We are interested in the value at $x_0 = 1$, i.e., $f(x_0 = 1)$, but cannot run an additional, third experiment.

```
from sklearn import linear_model
X = np.array([[0], [2]])
y = np.array([0.5, 2.5])
S_lm = linear_model.LinearRegression()
S_lm = S_lm.fit(X, y)
X0 = np.array([[1]])
y0 = S_lm.predict(X0)
print(y0)
```

[1.5]

Central Idea: Evaluation of the surrogate model S_{lm} is much cheaper (or / and much faster) than running the real-world experiment f .

21.13 Demo/Test: Objective Function Fails

SPOT expects `np.nan` values from failed objective function values. These are handled. Note: SPOT's counter considers only successful executions of the objective function.

```
import numpy as np
from spotPython.fun.objectivefunctions import analytical
from spotPython.spot import spot
import numpy as np
from math import inf
# number of initial points:
ni = 20
# number of points
n = 30
```

```

fun = analytical().fun_random_error
lower = np.array([-1])
upper = np.array([1])
design_control={"init_size": ni}

spot_1 = spot.Spot(fun=fun,
                  lower = lower,
                  upper= upper,
                  fun_evals = n,
                  show_progress=False,
                  design_control=design_control,)
spot_1.run()
# To check whether the run was successfully completed,
# we compare the number of evaluated points to the specified
# number of points.
assert spot_1.y.shape[0] == n

```

```

[ 0.53176481 -0.9053821 -0.02203599 -0.21843718  0.78240941 -0.58120945
 -0.3923345   0.67234256  0.31802454 -0.68898927 -0.75129705  0.97550354
  0.41757584  0.0786237   0.82585329  0.23700598 -0.49274073 -0.82319082
      nan  0.1481835 ]

```

```
[-1.]
```

```
[0.17335968]
```

```
[-0.58552368]
```

```
[-0.20126111]
```

```
[-0.60100809]
```

```
[-0.97897336]
```

```
[-0.2748985]
```

```
[0.8359486]
```

```
[0.99035591]
```

```
[0.01641232]
```

```
[0.5629346]
```

21.14 PyTorch: Detailed Description of the Data Splitting

21.14.1 Description of the "train_hold_out" Setting

The "train_hold_out" setting is used by default. It uses the loss function specified in `fun_control` and the metric specified in `fun_control`.

1. First, the method `HyperTorch().fun_torch` is called.
2. `fun_torc()`, which is implemented in the file `hypertorch.py`, calls `evaluate_hold_out()` as follows:

```
df_eval, _ = evaluate_hold_out(  
    model,  
    train_dataset=fun_control["train"],  
    shuffle=self.fun_control["shuffle"],  
    loss_function=self.fun_control["loss_function"],  
    metric=self.fun_control["metric_torch"],  
    device=self.fun_control["device"],  
    show_batch_interval=self.fun_control["show_batch_interval"],  
    path=self.fun_control["path"],  
    task=self.fun_control["task"],  
    writer=self.fun_control["writer"],  
    writerId=config_id,  
)
```

Note: Only the data set `fun_control["train"]` is used for training and validation. It is used in `evaluate_hold_out` as follows:

```
trainloader, valloader = create_train_val_data_loaders(  
    dataset=train_dataset, batch_size=batch_size_instance, shuffle=shuffle  
)
```

`create_train_val_data_loaders()` splits the `train_dataset` into `trainloader` and `valloader` using `torch.utils.data.random_split()` as follows:

```
def create_train_val_data_loaders(dataset, batch_size, shuffle, num_workers=0):  
    test_abs = int(len(dataset) * 0.6)  
    train_subset, val_subset = random_split(dataset, [test_abs, len(dataset) - test_abs])  
    trainloader = torch.utils.data.DataLoader(  
        train_subset, batch_size=int(batch_size), shuffle=shuffle, num_workers=num_workers  
    )  
    valloader = torch.utils.data.DataLoader(  
        val_subset, batch_size=int(batch_size), shuffle=shuffle, num_workers=num_workers  
    )
```

```

        val_subset, batch_size=int(batch_size), shuffle=shuffle, num_workers=num_workers
    )
    return trainloader, valloader

```

The optimizer is set up as follows:

```

optimizer_instance = net.optimizer
lr_mult_instance = net.lr_mult
sgd_momentum_instance = net.sgd_momentum
optimizer = optimizer_handler(
    optimizer_name=optimizer_instance,
    params=net.parameters(),
    lr_mult=lr_mult_instance,
    sgd_momentum=sgd_momentum_instance,
)

```

3. `evaluate_hold_out()` sets the net attributes such as `epochs`, `batch_size`, `optimizer`, and `patience`. For each epoch, the methods `train_one_epoch()` and `validate_one_epoch()` are called, the former for training and the latter for validation and early stopping. The validation loss from the last epoch (not the best validation loss) is returned from `evaluate_hold_out`.
4. The method `train_one_epoch()` is implemented as follows:

```

def train_one_epoch(
    net,
    trainloader,
    batch_size,
    loss_function,
    optimizer,
    device,
    show_batch_interval=10_000,
    task=None,
):
    running_loss = 0.0
    epoch_steps = 0
    for batch_nr, data in enumerate(trainloader, 0):
        input, target = data
        input, target = input.to(device), target.to(device)
        optimizer.zero_grad()
        output = net(input)
        if task == "regression":

```

```

        target = target.unsqueeze(1)
        if target.shape == output.shape:
            loss = loss_function(output, target)
        else:
            raise ValueError(f"Shapes of target and output do not match:
                               {target.shape} vs {output.shape}")
    elif task == "classification":
        loss = loss_function(output, target)
    else:
        raise ValueError(f"Unknown task: {task}")
    loss.backward()
    torch.nn.utils.clip_grad_norm_(net.parameters(), max_norm=1.0)
    optimizer.step()
    running_loss += loss.item()
    epoch_steps += 1
    if batch_nr % show_batch_interval == (show_batch_interval - 1):
        print(
            "Batch: %5d. Batch Size: %d. Training Loss (running): %.3f"
            % (batch_nr + 1, int(batch_size), running_loss / epoch_steps)
        )
        running_loss = 0.0
return loss.item()

```

5. The method `validate_one_epoch()` is implemented as follows:

```

def validate_one_epoch(net, valloader, loss_function, metric, device, task):
    val_loss = 0.0
    val_steps = 0
    total = 0
    correct = 0
    metric.reset()
    for i, data in enumerate(valloader, 0):
        # get batches
        with torch.no_grad():
            input, target = data
            input, target = input.to(device), target.to(device)
            output = net(input)
            # print(f"target: {target}")
            # print(f"output: {output}")
            if task == "regression":
                target = target.unsqueeze(1)

```

```

        if target.shape == output.shape:
            loss = loss_function(output, target)
        else:
            raise ValueError(f"Shapes of target and output
                do not match: {target.shape} vs {output.shape}")
        metric_value = metric.update(output, target)
    elif task == "classification":
        loss = loss_function(output, target)
        metric_value = metric.update(output, target)
        _, predicted = torch.max(output.data, 1)
        total += target.size(0)
        correct += (predicted == target).sum().item()
    else:
        raise ValueError(f"Unknown task: {task}")
    val_loss += loss.cpu().numpy()
    val_steps += 1
loss = val_loss / val_steps
print(f"Loss on hold-out set: {loss}")
if task == "classification":
    accuracy = correct / total
    print(f"Accuracy on hold-out set: {accuracy}")
# metric on all batches using custom accumulation
metric_value = metric.compute()
metric_name = type(metric).__name__
print(f"{metric_name} value on hold-out data: {metric_value}")
return metric_value, loss

```

21.14.1.1 Description of the "test_hold_out" Setting

It uses the loss function specified in `fun_control` and the metric specified in `fun_control`.

1. First, the method `HyperTorch().fun_torch` is called.
2. `fun_torch()` calls `spotPython.torch.traintest.evaluate_hold_out()` similar to the "train_hold_out" setting with one exception: It passes an additional test data set to `evaluate_hold_out()` as follows:

```
test_dataset=fun_control["test"]
```

`evaluate_hold_out()` calls `create_train_test_data_loaders` instead of `create_train_val_data_loaders`: The two data sets are used in `create_train_test_data_loaders` as follows:

```

def create_train_test_data_loaders(dataset, batch_size, shuffle, test_dataset,
    num_workers=0):
    trainloader = torch.utils.data.DataLoader(
        dataset, batch_size=int(batch_size), shuffle=shuffle,
        num_workers=num_workers
    )
    testloader = torch.utils.data.DataLoader(
        test_dataset, batch_size=int(batch_size), shuffle=shuffle,
        num_workers=num_workers
    )
    return trainloader, testloader

```

3. The following steps are identical to the "train_hold_out" setting. Only a different data loader is used for testing.

21.14.1.2 Detailed Description of the "train_cv" Setting

It uses the loss function specified in `fun_control` and the metric specified in `fun_control`.

1. First, the method `HyperTorch().fun_torch` is called.
2. `fun_torch()` calls `spotPython.torch.traintest.evaluate_cv()` as follows (Note: Only the data set `fun_control["train"]` is used for CV.):

```

df_eval, _ = evaluate_cv(
    model,
    dataset=fun_control["train"],
    shuffle=self.fun_control["shuffle"],
    device=self.fun_control["device"],
    show_batch_interval=self.fun_control["show_batch_interval"],
    task=self.fun_control["task"],
    writer=self.fun_control["writer"],
    writerId=config_id,
)

```

3. In 'evaluate_cv()', the following steps are performed: The optimizer is set up as follows:

```

optimizer_instance = net.optimizer
lr_instance = net.lr
sgd_momentum_instance = net.sgd_momentum
optimizer = optimizer_handler(optimizer_name=optimizer_instance,
    params=net.parameters(), lr_mult=lr_mult_instance)

```

`evaluate_cv()` sets the `net` attributes such as `epochs`, `batch_size`, `optimizer`, and `patience`. CV is implemented as follows:

```
def evaluate_cv(
    net,
    dataset,
    shuffle=False,
    loss_function=None,
    num_workers=0,
    device=None,
    show_batch_interval=10_000,
    metric=None,
    path=None,
    task=None,
    writer=None,
    writerId=None,
):
    lr_mult_instance = net.lr_mult
    epochs_instance = net.epochs
    batch_size_instance = net.batch_size
    k_folds_instance = net.k_folds
    optimizer_instance = net.optimizer
    patience_instance = net.patience
    sgd_momentum_instance = net.sgd_momentum
    removed_attributes, net = get_removed_attributes_and_base_net(net)
    metric_values = {}
    loss_values = {}
    try:
        device = getDevice(device=device)
        if torch.cuda.is_available():
            device = "cuda:0"
            if torch.cuda.device_count() > 1:
                print("We will use", torch.cuda.device_count(), "GPUs!")
                net = nn.DataParallel(net)
        net.to(device)
        optimizer = optimizer_handler(
            optimizer_name=optimizer_instance,
            params=net.parameters(),
            lr_mult=lr_mult_instance,
            sgd_momentum=sgd_momentum_instance,
        )
        kfold = KFold(n_splits=k_folds_instance, shuffle=shuffle)
```

```

for fold, (train_ids, val_ids) in enumerate(kfold.split(dataset)):
    print(f"Fold: {fold + 1}")
    train_subsampler = torch.utils.data.SubsetRandomSampler(train_ids)
    val_subsampler = torch.utils.data.SubsetRandomSampler(val_ids)
    trainloader = torch.utils.data.DataLoader(
        dataset, batch_size=batch_size_instance,
        sampler=train_subsampler, num_workers=num_workers
    )
    valloader = torch.utils.data.DataLoader(
        dataset, batch_size=batch_size_instance,
        sampler=val_subsampler, num_workers=num_workers
    )
    # each fold starts with new weights:
    reset_weights(net)
    # Early stopping parameters
    best_val_loss = float("inf")
    counter = 0
    for epoch in range(epochs_instance):
        print(f"Epoch: {epoch + 1}")
        # training loss from one epoch:
        training_loss = train_one_epoch(
            net=net,
            trainloader=trainloader,
            batch_size=batch_size_instance,
            loss_function=loss_function,
            optimizer=optimizer,
            device=device,
            show_batch_interval=show_batch_interval,
            task=task,
        )
        # Early stopping check. Calculate validation loss from one epoch:
        metric_values[fold], loss_values[fold] = validate_one_epoch(
            net, valloader=valloader, loss_function=loss_function,
            metric=metric, device=device, task=task
        )
        # Log the running loss averaged per batch
        metric_name = "Metric"
        if metric is None:
            metric_name = type(metric).__name__
            print(f"{metric_name} value on hold-out data:
                {metric_values[fold]}")

```

```

    if writer is not None:
        writer.add_scalars(
            "evaluate_cv fold:" + str(fold + 1) +
            ". Train & Val Loss and Val Metric" + writerId,
            {"Train loss": training_loss, "Val loss":
            loss_values[fold], metric_name: metric_values[fold]},
            epoch + 1,
        )
        writer.flush()
    if loss_values[fold] < best_val_loss:
        best_val_loss = loss_values[fold]
        counter = 0
        # save model:
        if path is not None:
            torch.save(net.state_dict(), path)
    else:
        counter += 1
        if counter >= patience_instance:
            print(f"Early stopping at epoch {epoch}")
            break

    df_eval = sum(loss_values.values()) / len(loss_values.values())
    df_metrics = sum(metric_values.values()) / len(metric_values.values())
    df_preds = np.nan
except Exception as err:
    print(f"Error in Net_Core. Call to evaluate_cv() failed. {err=},
          {type(err)=}")
    df_eval = np.nan
    df_preds = np.nan
add_attributes(net, removed_attributes)
if writer is not None:
    metric_name = "Metric"
    if metric is None:
        metric_name = type(metric).__name__
    writer.add_scalars(
        "CV: Val Loss and Val Metric" + writerId,
        {"CV-loss": df_eval, metric_name: df_metrics},
        epoch + 1,
    )
    writer.flush()
return df_eval, df_preds, df_metrics

```

4. The method `train_fold()` is implemented as shown above.

5. The method `validate_one_epoch()` is implemented as shown above. In contrast to the hold-out setting, it is called for each of the k folds. The results are stored in a dictionaries `metric_values` and `loss_values`. The results are averaged over the k folds and returned as `df_eval`.

21.14.1.3 Detailed Description of the "test_cv" Setting

It uses the loss function specified in `fun_control` and the metric specified in `fun_control`.

1. First, the method `HyperTorch().fun_torch` is called.
2. `fun_torch()` calls `spotPython.torch.traintest.evaluate_cv()` as follows:

```
df_eval, _ = evaluate_cv(  
    model,  
    dataset=fun_control["test"],  
    shuffle=self.fun_control["shuffle"],  
    device=self.fun_control["device"],  
    show_batch_interval=self.fun_control["show_batch_interval"],  
    task=self.fun_control["task"],  
    writer=self.fun_control["writer"],  
    writerId=config_id,  
)
```

Note: The data set `fun_control["test"]` is used for CV. The rest is the same as for the "train_cv" setting.

21.14.1.4 Detailed Description of the Final Model Training and Evaluation

There are two methods that can be used for the final evaluation of a Pytorch model:

1. "train_tuned and
2. "test_tuned".

`train_tuned()` is just a wrapper to `evaluate_hold_out` using the `train` data set. It is implemented as follows:

```
def train_tuned(  
    net,  
    train_dataset,  
    shuffle,  
    loss_function,  
    metric,
```

```

device=None,
show_batch_interval=10_000,
path=None,
task=None,
writer=None,
):
    evaluate_hold_out(
        net=net,
        train_dataset=train_dataset,
        shuffle=shuffle,
        test_dataset=None,
        loss_function=loss_function,
        metric=metric,
        device=device,
        show_batch_interval=show_batch_interval,
        path=path,
        task=task,
        writer=writer,
    )

```

The `test_tuned()` procedure is implemented as follows:

```

def test_tuned(net, shuffle, test_dataset=None, loss_function=None,
metric=None, device=None, path=None, task=None):
    batch_size_instance = net.batch_size
    removed_attributes, net = get_removed_attributes_and_base_net(net)
    if path is not None:
        net.load_state_dict(torch.load(path))
        net.eval()
    try:
        device = getDevice(device=device)
        if torch.cuda.is_available():
            device = "cuda:0"
            if torch.cuda.device_count() > 1:
                print("We will use", torch.cuda.device_count(), "GPUs!")
                net = nn.DataParallel(net)
        net.to(device)
        valloader = torch.utils.data.DataLoader(
            test_dataset, batch_size=int(batch_size_instance),
            shuffle=shuffle,
            num_workers=0
        )

```

```

metric_value, loss = validate_one_epoch(
    net, valloader=valloader, loss_function=loss_function,
    metric=metric, device=device, task=task
)
df_eval = loss
df_metric = metric_value
df_preds = np.nan
except Exception as err:
    print(f"Error in Net_Core. Call to test_tuned() failed. {err=},
          {type(err)=}")
    df_eval = np.nan
    df_metric = np.nan
    df_preds = np.nan
add_attributes(net, removed_attributes)
print(f"Final evaluation: Validation loss: {df_eval}")
print(f"Final evaluation: Validation metric: {df_metric}")
print("-----")
return df_eval, df_preds, df_metric

```

References

- Bartz, Eva, Thomas Bartz-Beielstein, Martin Zaefferer, and Olaf Mersmann, eds. 2022. *Hyperparameter Tuning for Machine and Deep Learning with R - A Practical Guide*. Springer.
- Bartz-Beielstein, Thomas. 2023. “PyTorch Hyperparameter Tuning with SPOT: Comparison with Ray Tuner and Default Hyperparameters on CIFAR10.” https://github.com/sequential-parameter-optimization/spotPython/blob/main/notebooks/14_spot_ray_hpt_torch_cifar10.ipynb.
- Bartz-Beielstein, Thomas, Jürgen Branke, Jörn Mehnen, and Olaf Mersmann. 2014. “Evolutionary Algorithms.” *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 4 (3): 178–95.
- Bartz-Beielstein, Thomas, Carola Doerr, Jakob Bossek, Sowmya Chandrasekaran, Tome Eftimov, Andreas Fischbach, Pascal Kerschke, et al. 2020. “Benchmarking in Optimization: Best Practice and Open Issues.” arXiv. <https://arxiv.org/abs/2007.03488>.
- Bartz-Beielstein, Thomas, Christian Lasarczyk, and Mike Preuss. 2005. “Sequential Parameter Optimization.” In *Proceedings 2005 Congress on Evolutionary Computation (CEC’05), Edinburgh, Scotland*, edited by B McKay et al., 773–80. Piscataway NJ: IEEE Press.
- Lewis, R M, V Torczon, and M W Trosset. 2000. “Direct search methods: Then and now.” *Journal of Computational and Applied Mathematics* 124 (1–2): 191–207.
- Li, Lisha, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. 2016. “Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization.” *arXiv e-Prints*, March, arXiv:1603.06560.
- Meignan, David, Sigrid Knust, Jean-Marc Frayet, Gilles Pesant, and Nicolas Gaud. 2015. “A Review and Taxonomy of Interactive Optimization Methods in Operations Research.” *ACM Transactions on Interactive Intelligent Systems*, September.
- Montiel, Jacob, Max Halford, Saulo Martiello Mastelini, Geoffrey Bolmier, Raphael Sourty, Robin Vaysse, Adil Zouitine, et al. 2021. “River: Machine Learning for Streaming Data in Python.”
- PyTorch. 2023a. “Hyperparameter Tuning with Ray Tune.” https://pytorch.org/tutorials/beginner/hyperparameter_tuning_tutorial.html.
- . 2023b. “Training a Classifier.” https://pytorch.org/tutorials/beginner/blitz/cifar10_tutorial.html.