A Simple Baseline for Bayesian Uncertainty in Deep Learning

NeurIPS 2019

Wesley Maddox Timur Garipov Pavel Izmailov
Dmitry Vetrov Andrew Gordon Wilson

Bayes rule:

$$p(w \mid Data) = \frac{p(Data \mid w)p(w)}{\int p(Data \mid w')p(w')dw'}$$

Bayesian Model Averaging:

$$p_{BMA}(y \mid x) = \int p(y \mid w, x)p(w \mid Data)dw \approx \sum_i p(y \mid w_i, x), \quad w_i \sim p(w \mid Data)$$

Bayes rule:

$$p(w \mid Data) = \frac{p(Data \mid w)p(w)}{\int p(Data \mid w')p(w')dw'}$$

Intractable

Bayesian Model Averaging:

$$p_{BMA}(y \mid x) = \int p(y \mid w, x)p(w \mid Data)dw \approx \sum_i p(y \mid w_i, x), \quad w_i \sim p(w \mid Data)$$

Bayes rule: $$p(w \mid Data) = \frac{p(Data \mid w)p(w)}{\int p(Data \mid w')p(w')dw'}$$

Intractable

Bayesian Model Averaging:

$$p_{BMA}(y \mid x) = \int p(y \mid w, x)p(w \mid Data)dw \approx \sum_i p(y \mid w_i, x), \quad w_i \sim p(w \mid Data)$$

- Approximate the posterior
- Variational Inference, MCMC, Laplace, ...
- Challenge: capture the *geometry* of the posterior

$posterior(w) \propto \exp(-Loss(w))$

Example:

- 2D quadratic loss
- Constant LR SGD
- Isotropic Gaussian noise

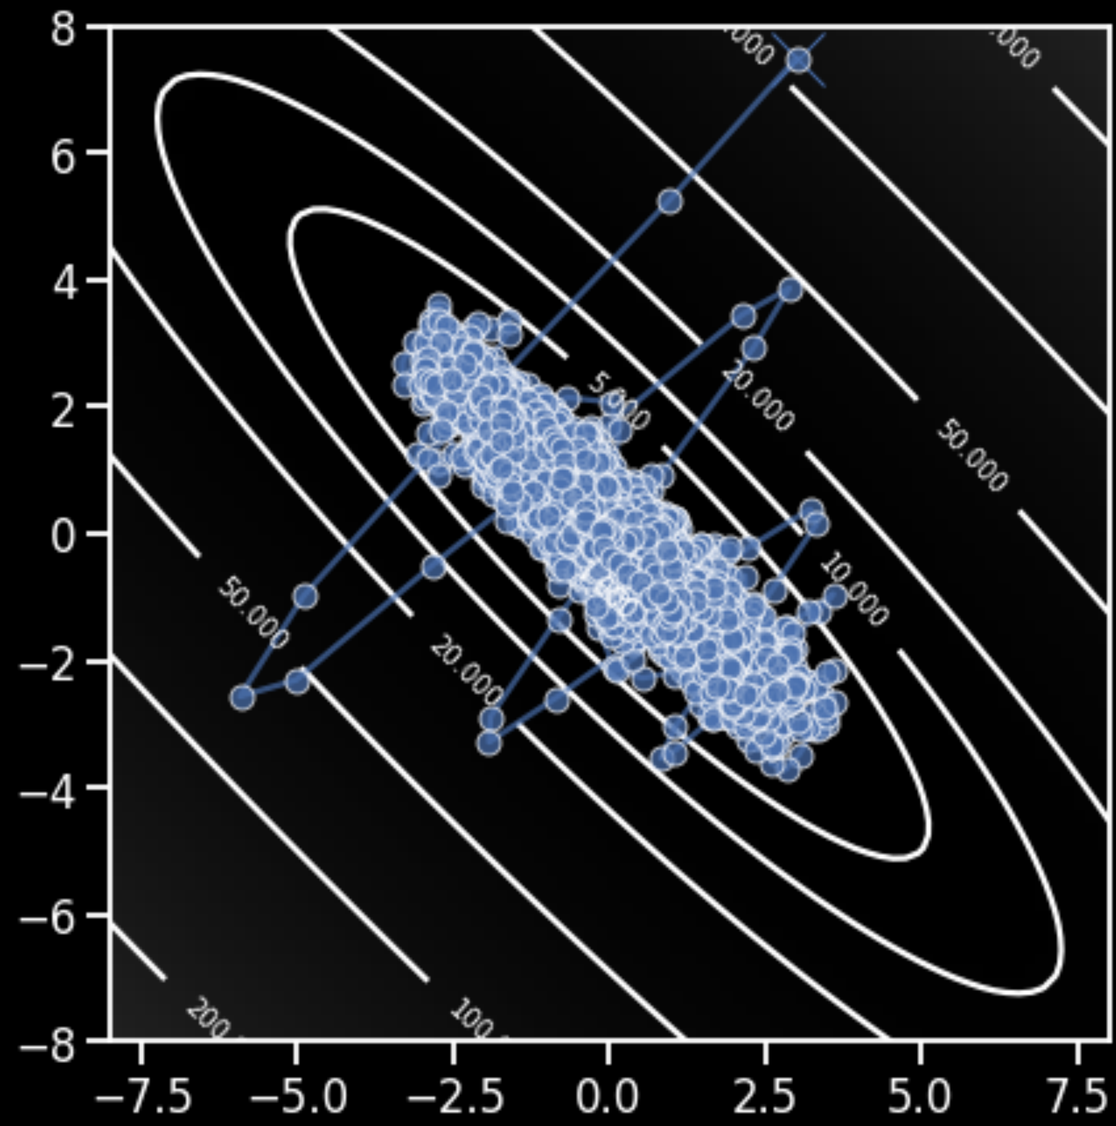

SGD with isotropic noise follows the shape of the posterior!*

We can use the same idea to approximate the posterior in DNNs

*see "*Stochastic Gradient Descent as Approximate Bayesian Inference*" by Mandt et al. JMLR 2017

SWA-Gaussian (SWAG)

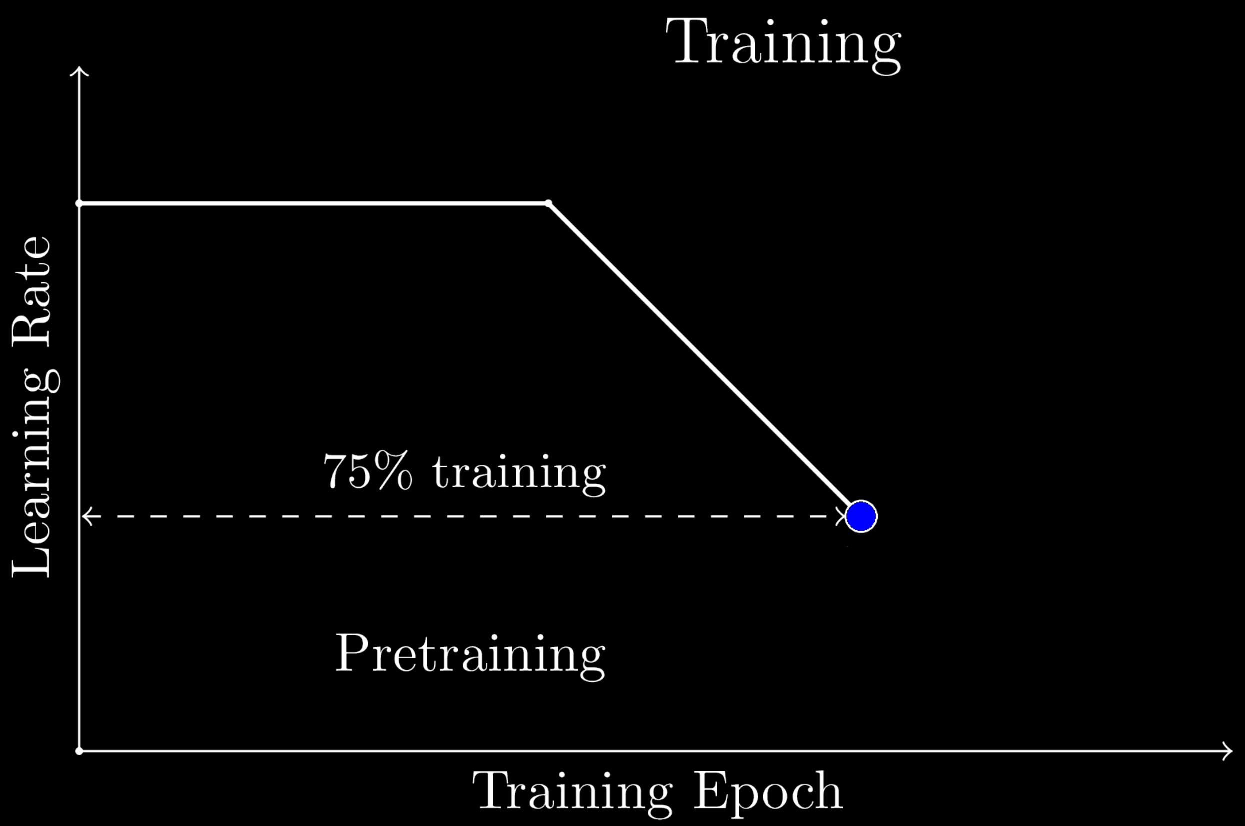

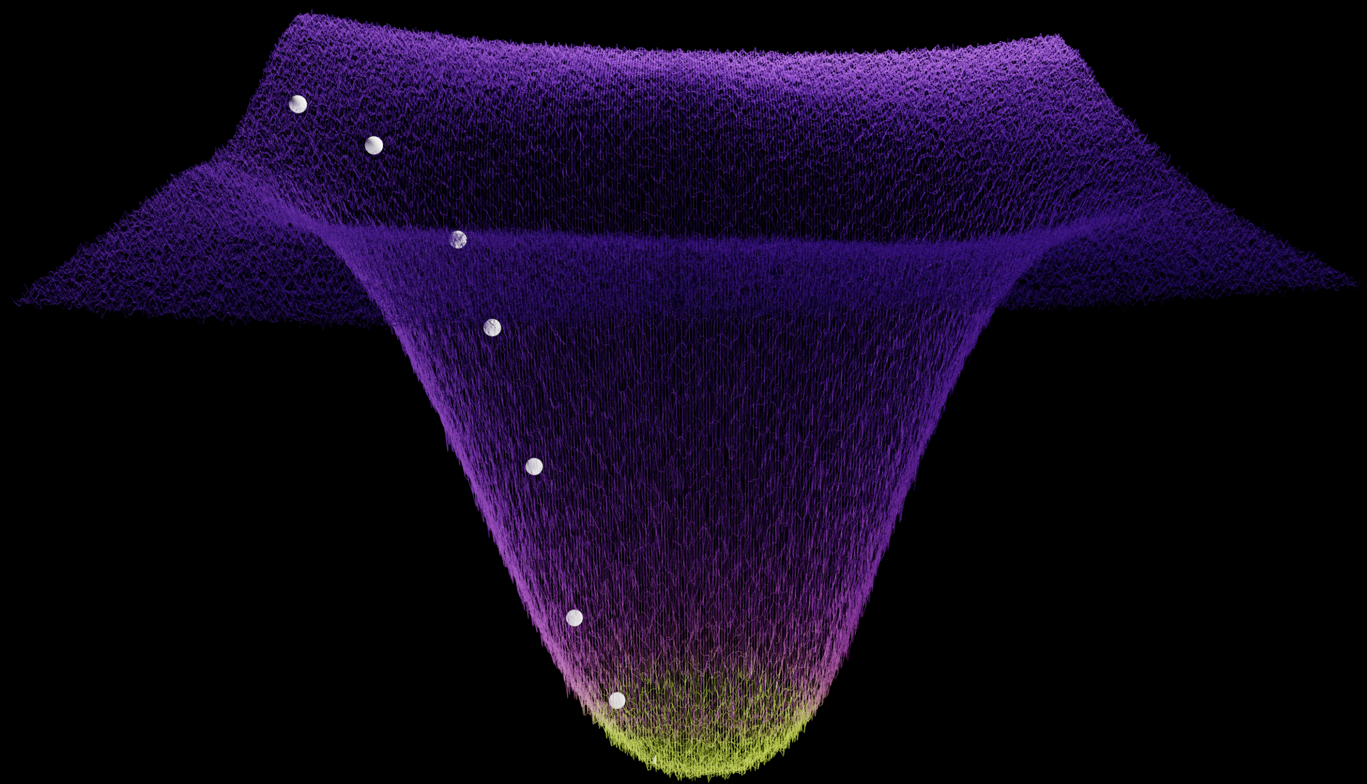

Visualization created in collaboration with Javier Ideami (https://losslandscape.com)

SWAG

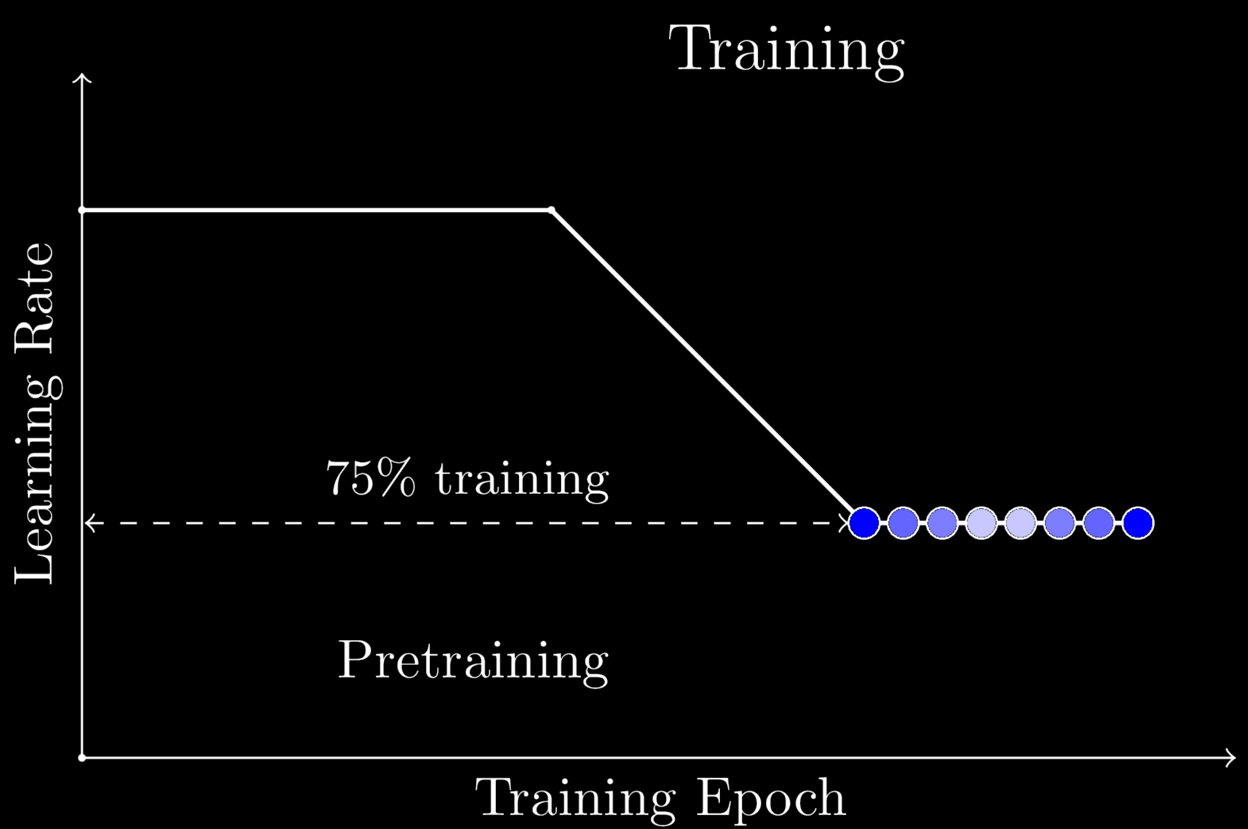

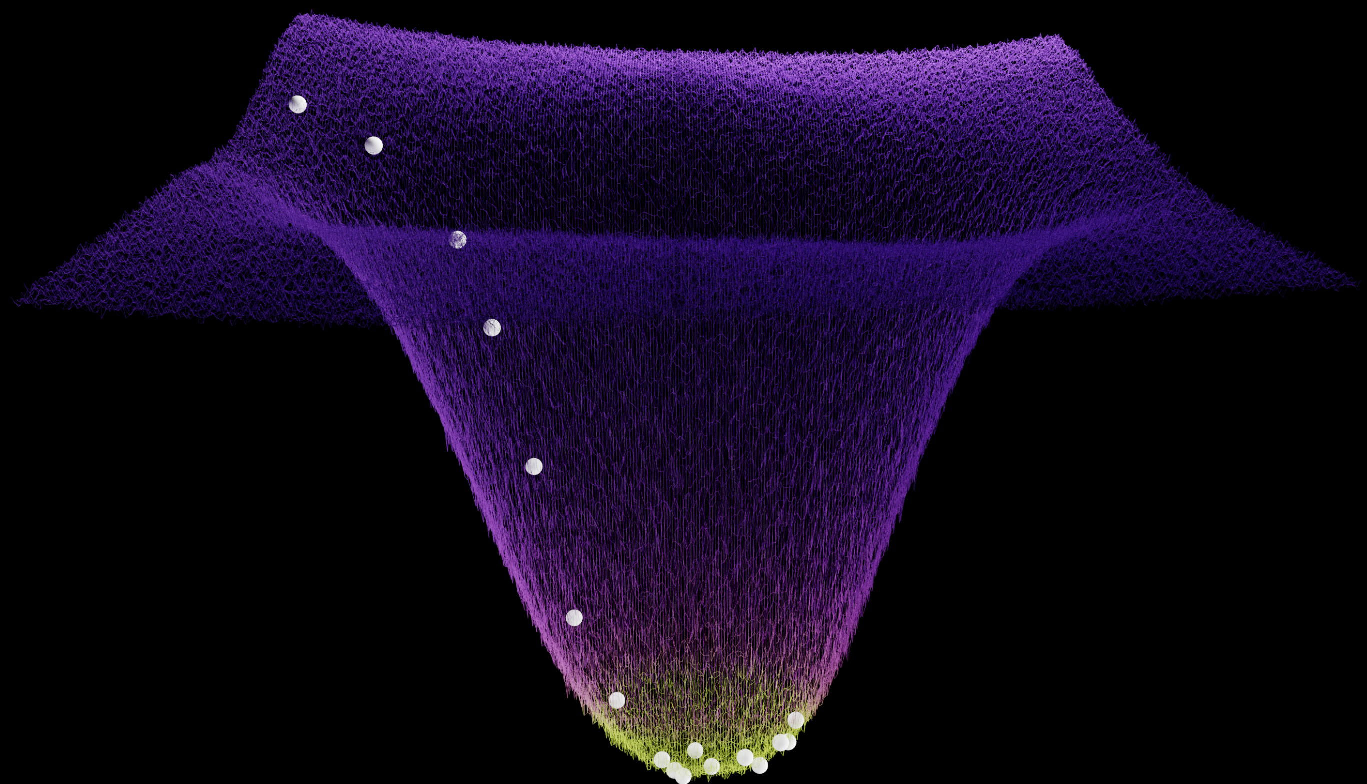

Visualization created in collaboration with Javier Ideami (https://losslandscape.com)

SWAG

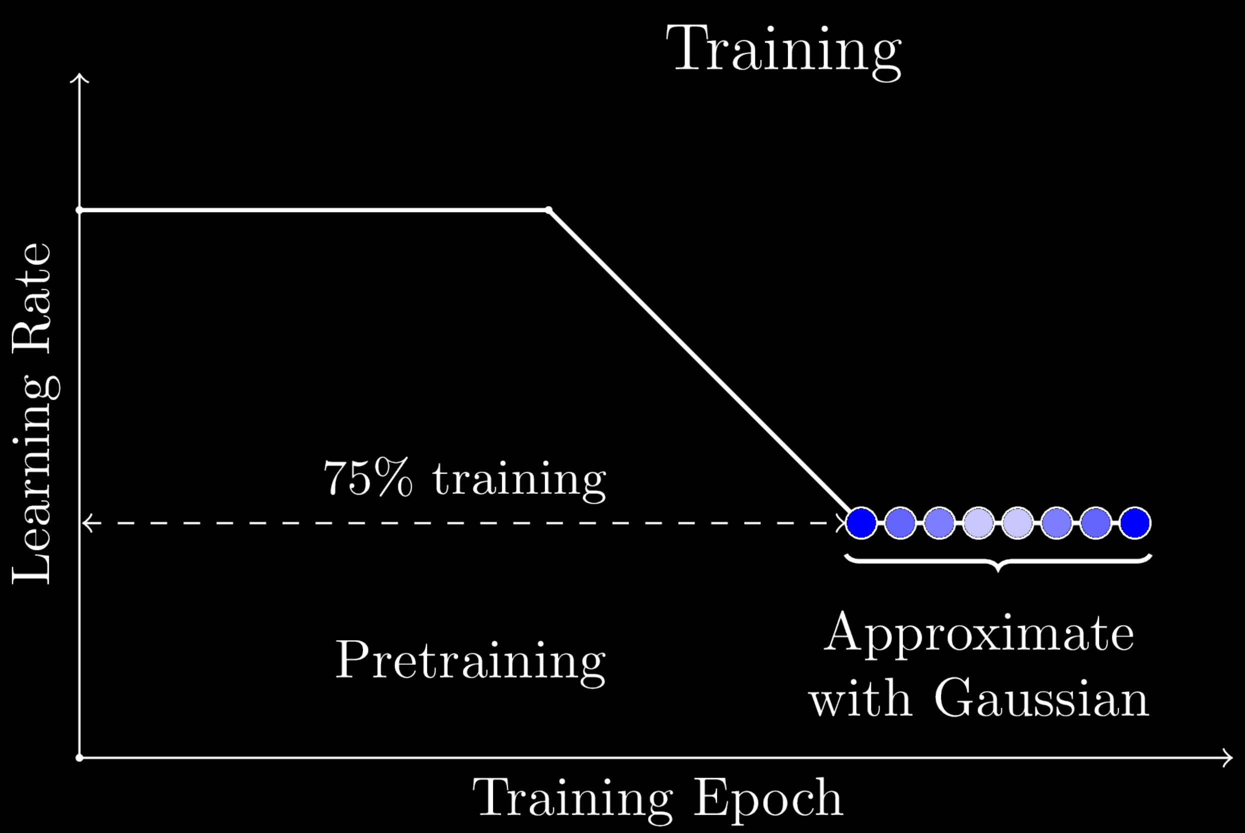

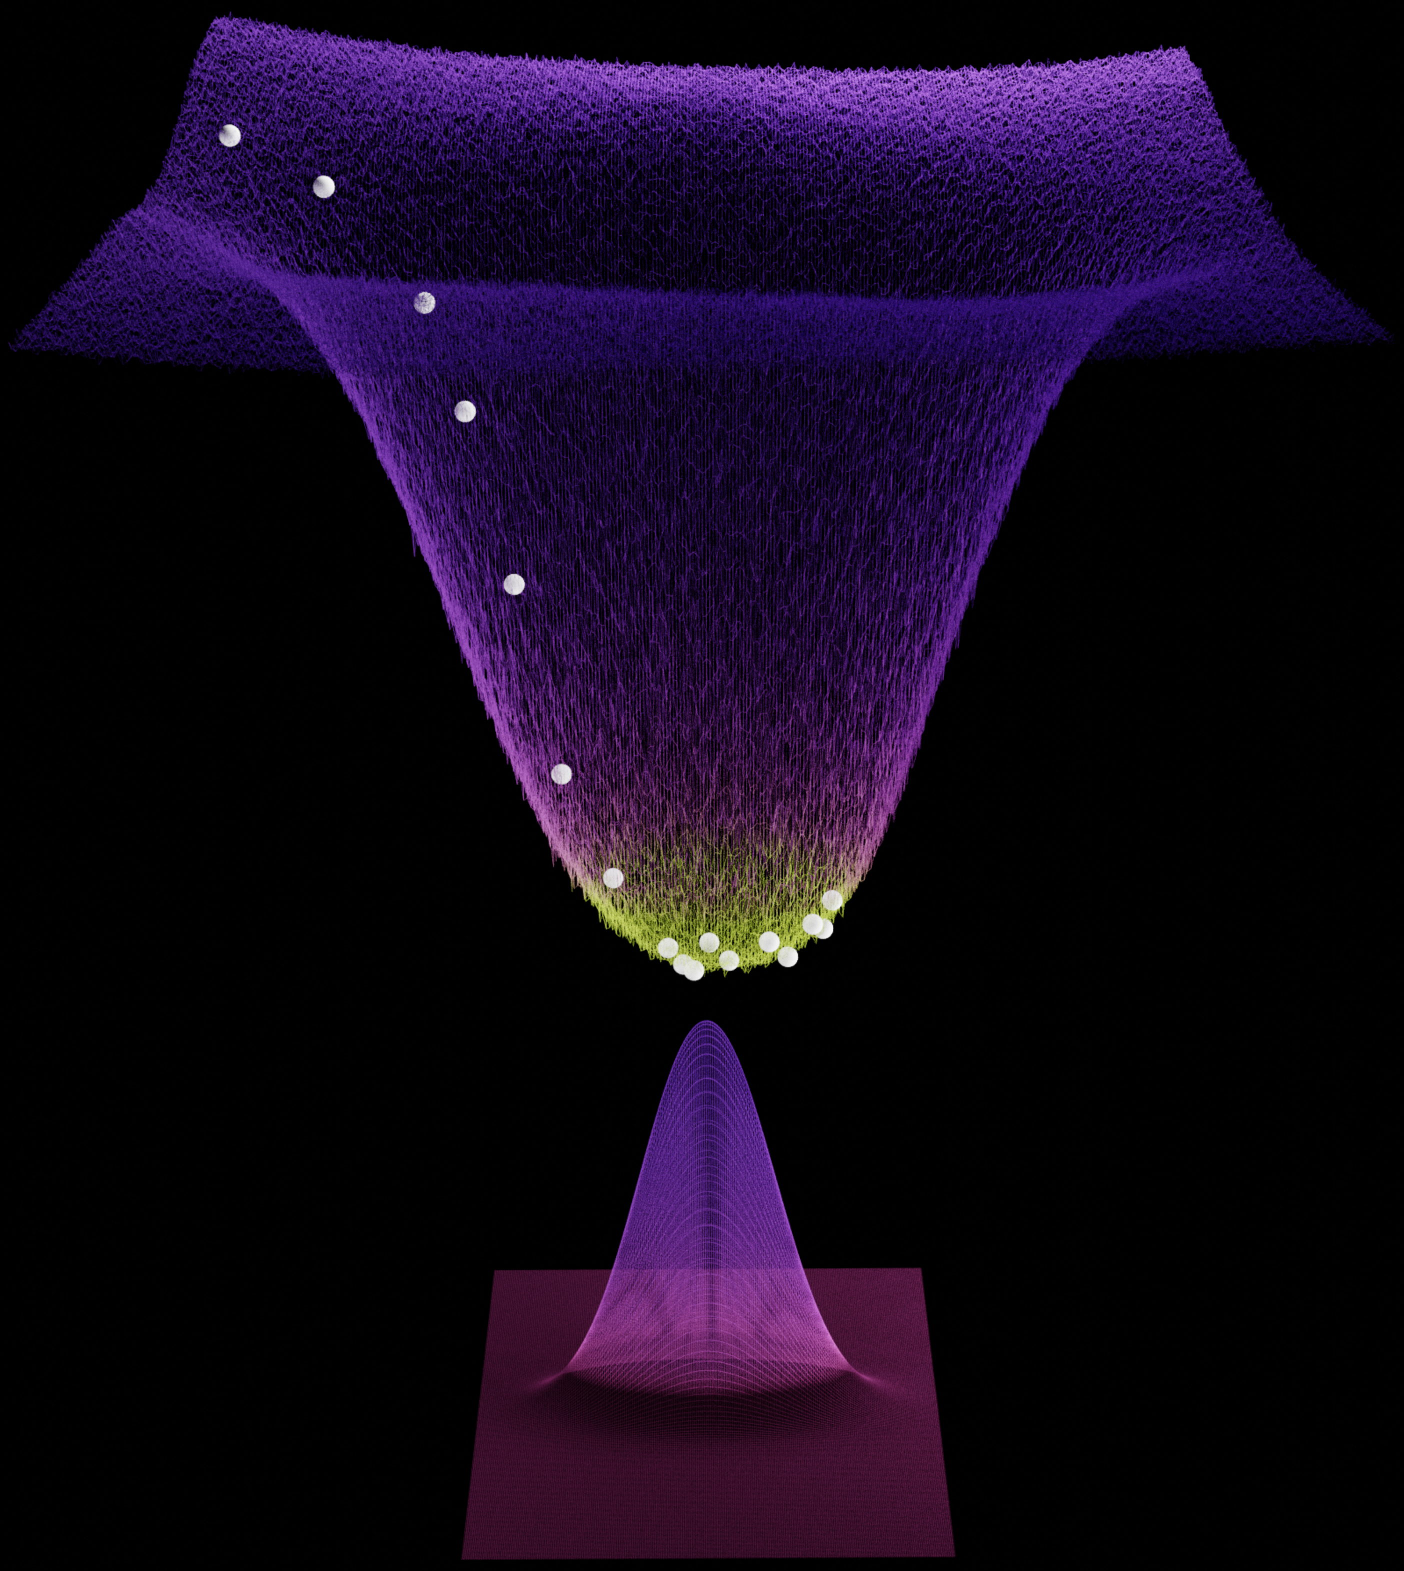

Visualization created in collaboration with Javier Ideami (https://losslandscape.com)

SWAG

Training

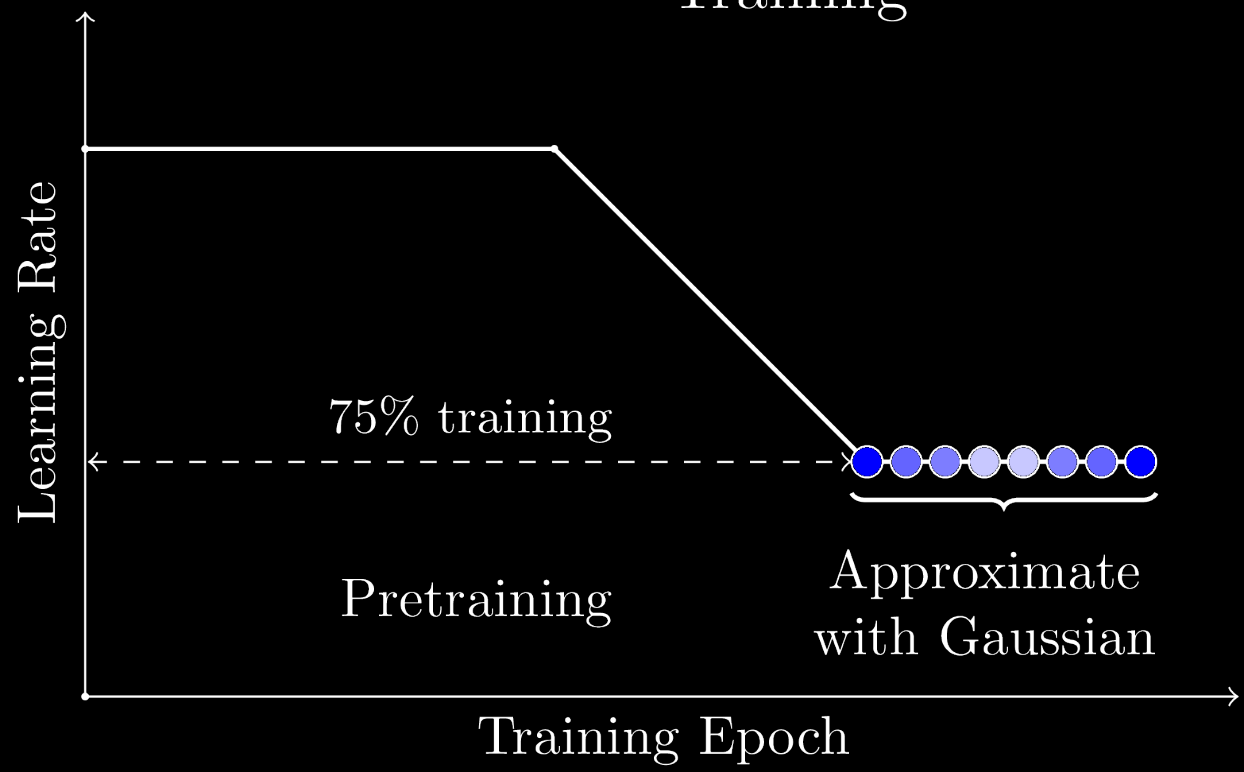

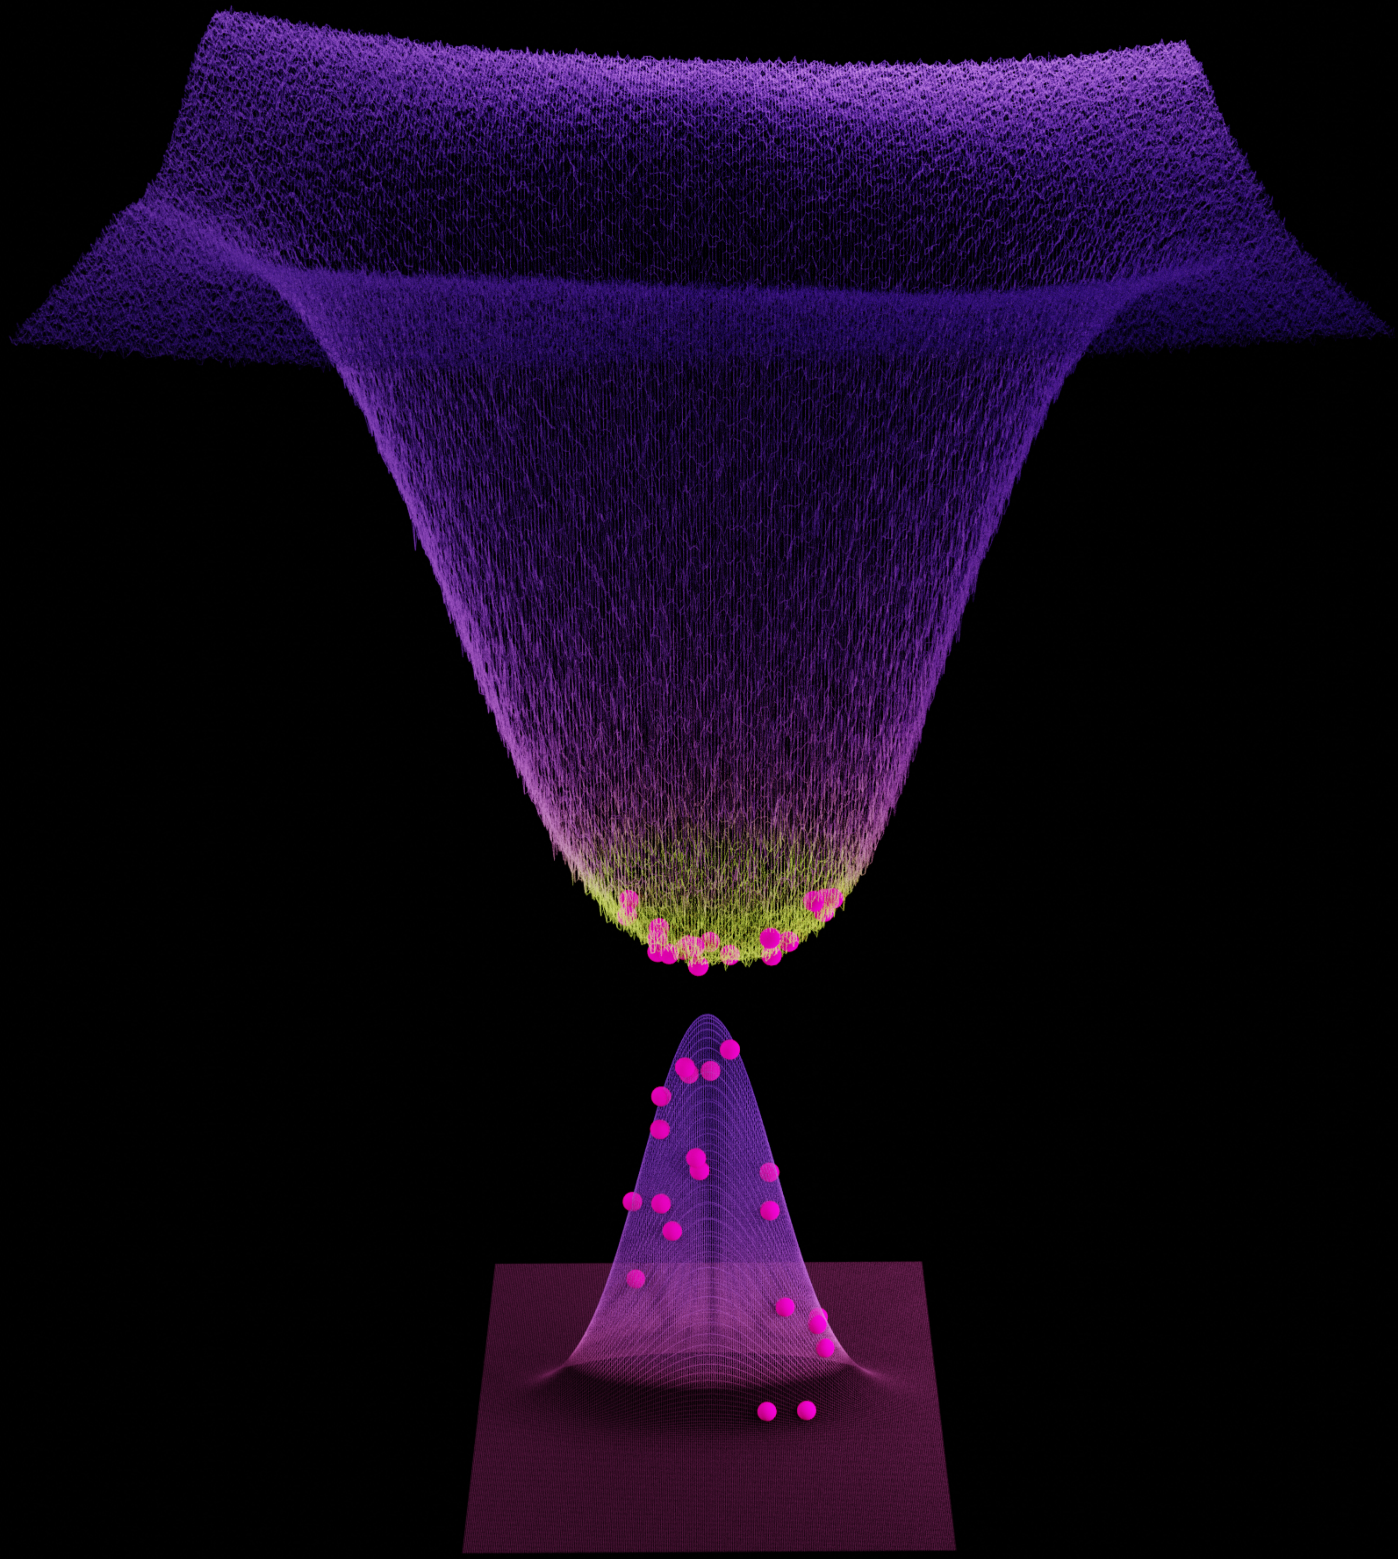

Evaluation

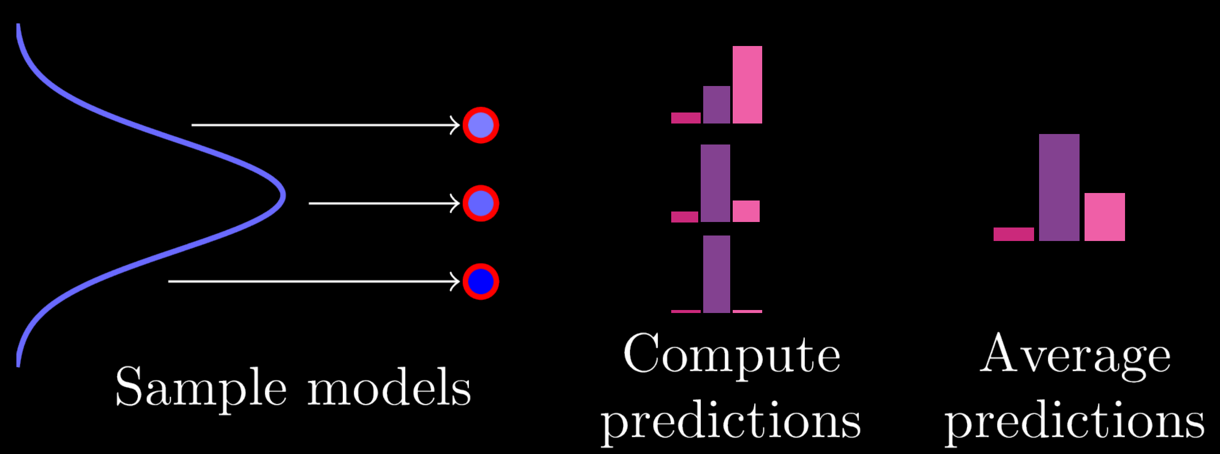

Visualization created in collaboration with Javier Ideami (https://losslandscape.com)

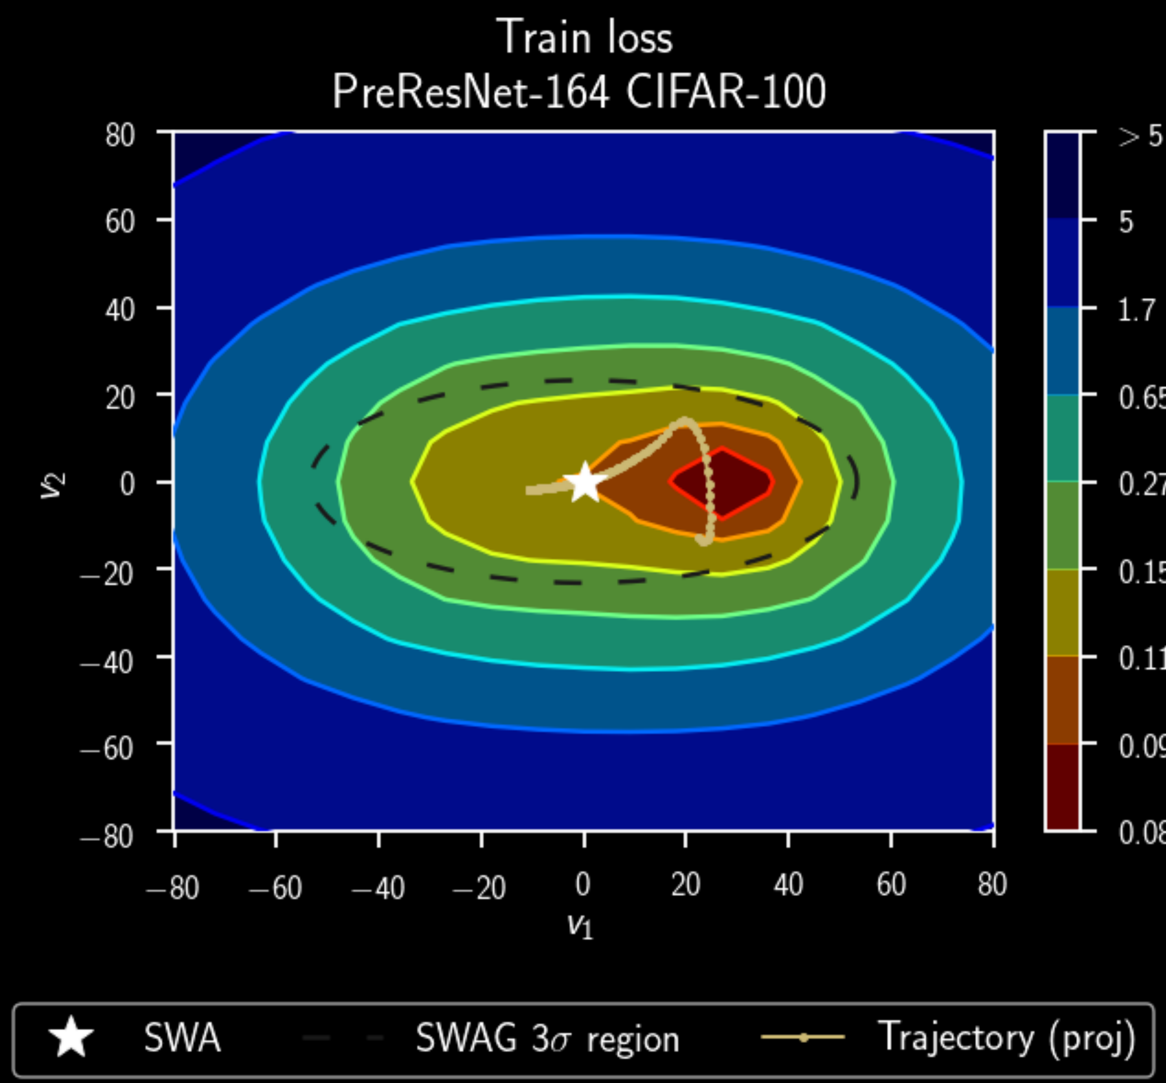

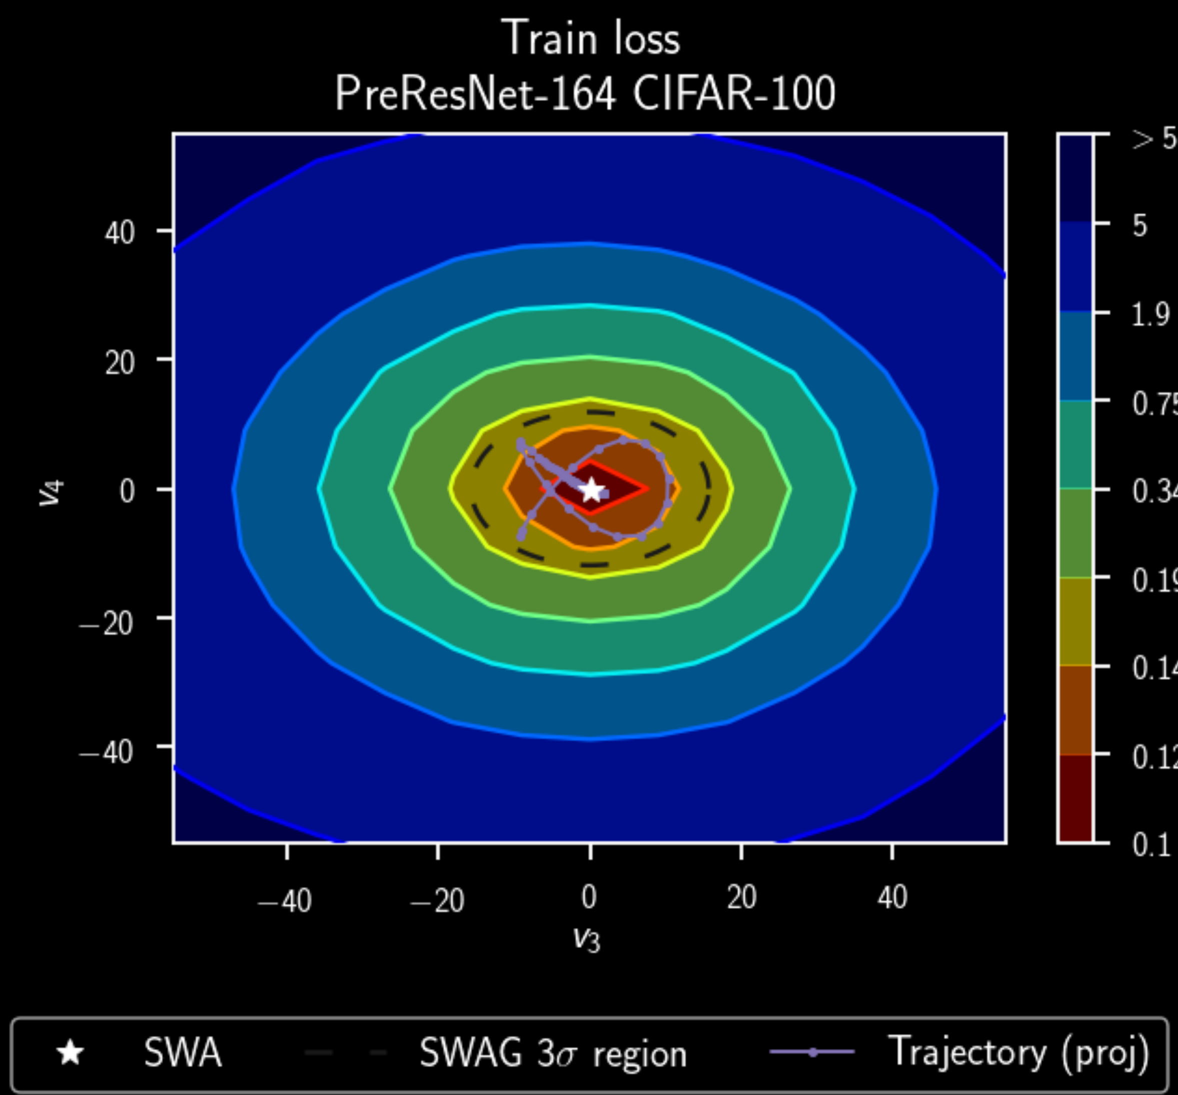

SWAG captures local geometry of the posterior

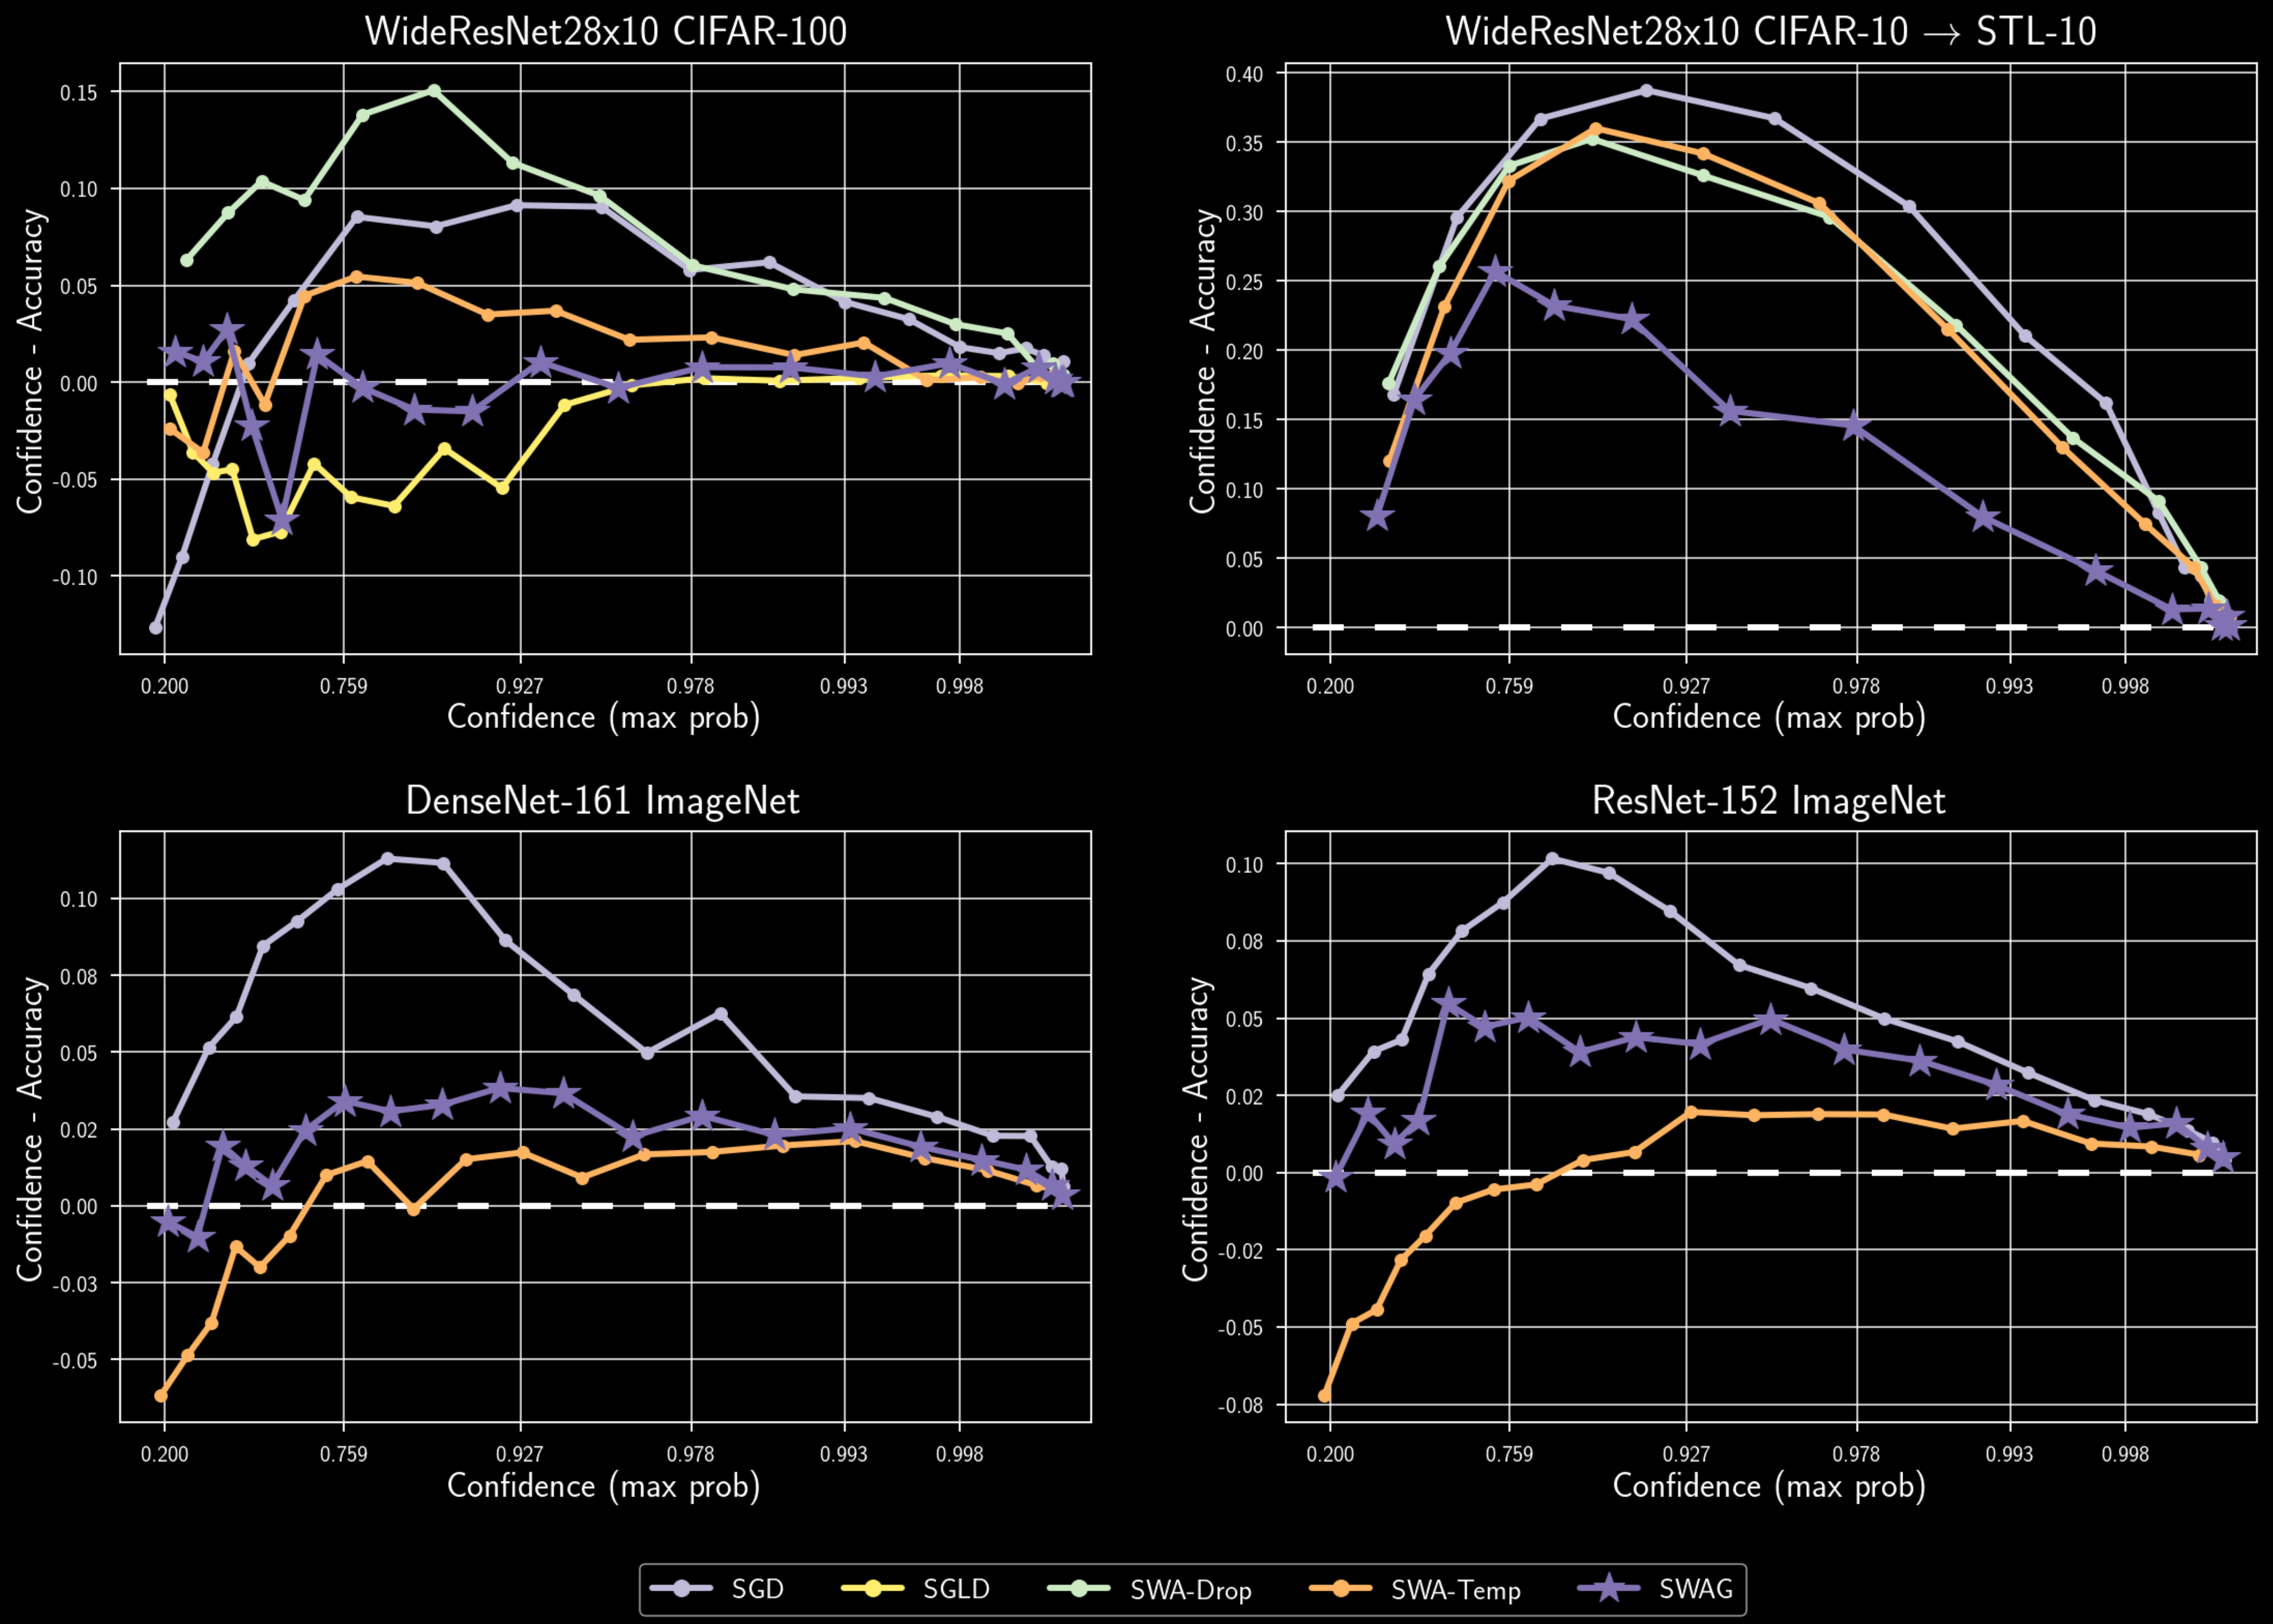

SWAG provides better-calibrated uncertainties!

Summary

- SWAG is simple, easy to use and consistently improves accuracy and uncertainty
- Applicable to ImageNet-scale problems
- Captures local geometry of the posterior

PyTorch code

github.com/wjmaddox/swa_gaussian

Come to our poster poster **#146**

Tue Dec 10th at 10:45 AM

East Exhibition Hall B + C