
LEARNING PEPTIDE PROPERTIES WITH POSITIVE EXAMPLES ONLY

Mehrad Ansari and Andrew D. White*

Department of Chemical Engineering, University of Rochester, Rochester, NY, 14627, USA.

ABSTRACT

Deep learning can create accurate predictive models by exploiting existing large-scale experimental data, and guide the design of molecules. However, a major barrier is the requirement of both positive and negative examples in the classical supervised learning frameworks. Notably, most peptide databases come with missing information and low number of observations on negative examples, as such sequences are hard to obtain using high-throughput screening methods. To address this challenge, we solely exploit the limited known positive examples in a semi-supervised setting, and discover peptide sequences that are likely to map to certain antimicrobial properties via positive-unlabeled learning (PU). In particular, we use the two learning strategies of adapting base classifier and reliable negative identification to build deep learning models for inferring solubility, hemolysis, binding against SHP-2, and non-fouling activity of peptides, given their sequence. We evaluate the predictive performance of our PU learning method and show that by only using the positive data, it can achieve competitive performance when compared with the classical positive-negative (PN) classification approach, where there is access to both positive and negative examples.

1 Introduction

As short-chain amino acids, peptides have attracted growing attention in pharmaceuticals [1–3], therapeutics [4–6], immunology [7–9], and biomaterials design [10–12]. However, the development of novel peptides remains a challenge due to poor pharmacokinetic properties that restrict the design space and necessitate unnatural amino acids or cyclization, increasing the complexity of their design.[13] Computational design and data-driven discovery strategies have arisen as promising low-cost techniques in the pre-experiment phase to expedite the process of generating accurate predictions of peptide properties, and shortlist promising candidates for follow-up experimental validation. Some examples of these successful applications include single nucleotide polymorphisms (SNP) and small-indel calling [14], estimating the impact of non-coding variants on DNA-methylation [15], as well as for the prediction of protein function [16], structure [17, 18], and protein-protein interactions [19]. Sequence-based learning strategies aim at mapping peptide’s natural biological function to its sequence. In a supervised learning setting, this is done by training on sequence-function examples. This means that sequence-function relationships are learned by iteratively training on samples of different classes (i.e. positive and negative examples in binary classification). The performance of the classifier is highly dependent on the quality of the training samples and the ratio of the positive and negative samples [20, 21]. In bioinformatics, a variety of supervised-learning algorithms, such as support vector machines [22], random forest [23], logistic regression [24], and naive Bayesian classifier [25], have been successfully applied to develop classification models.

However, lack of negative examples in numerous biological applications [26–29] limits the feasibility of constructing such reliable classifiers. As an example, medical information records typically contain the positively diagnosed diseases of a patient, and the absence of a diagnostic record does not necessarily rule out a disease for him/her. Most high-throughput screening methods solely focus on identifying the positive examples, thus, it is much more straightforward to confirm a property than to ascertain that it does not hold. As an example, a potential binding site is confirmed if a

*andrew.white@rochester.edu

protein binds to a target, but failure to bind only means that the binding conditions were not satisfied under a given experimental setting. With the technological advances, identifying specific properties can be improved, and biological samples formerly not known to have a property can now be classified with confidence. As an example, Li et al. [30] demonstrated on the changes in protein glycosylation site labeling throughout four time points over 10 years. Another example is protein-protein interaction (PPI) [31, 32], where experimentally validated interacting and non-interacting protein pairs are used as positive and negative examples, respectively. However, the selection of non-interacting protein pairs can be challenging for two reasons: 1. As more novel PPIs are constantly being discovered over time, some non-interacting protein pairs (i.e. negative examples) might be mislabeled. 2. The positive examples are significantly outnumbered by a large number of protein pairs for which no interactions have been identified. Similar situations can be found in drug–drug interaction identification [33], small non-coding RNA detection [34], gene function [35, 36] and phage-bacteria interaction [37] prediction, and biological sequence classification [38, 39].

To address the challenges above, we demonstrate on a positive-unlabeled (PU) learning framework to infer peptide sequence-function relationships, by solely exploiting the limited known positive examples in a semi-supervised setting. Semi-supervised learning techniques are a special instance of weak supervision [40, 41], where the training is based on partially labeled training data (i.e. labeled data can be either positive or both positive and negative samples). PU learning builds classification models by primarily leveraging a small number of labeled positive samples and a huge volume of unlabeled samples (i.e. a mixture of both *positive* (P) and *negative* (N) samples) [42]. Depending on how the *unlabeled* (U) data is handled, existing PU learning strategies are divided into two categories. 1. *Reliable negative identification*: this category identifies *reliable negatives* (RN) within U, and then performs ordinary supervised (PN) learning [43, 44]; 2. *Adapting the base classifier*: this treats the U samples as N with smaller weights (biased learning) and adapts the conventional classifiers to directly learn from P and U samples [45, 46]. The former *reliable negative identification* strategies rely on heuristics to identify the RN, and they have been widely used in non-coding RNA identification [34], non-coding RNA-disease association [47], gene function prediction [35, 48], disease gene identification [26, 49, 50], and single-cell RNA sequencing quality control [51]. On the other hand, *adapting the base classifier* algorithms are Bayesian-based approaches that focus on estimating the ratio of positive and negative samples in U (class prior), which then can be applied for classification using the Bayes' rule. One major limitation is that their performance largely depends on good choices of weights of U samples, which are computationally expensive to tune [52]. Thus, compared to the first strategy, there has been a fewer use cases of them in the literature [53–55]. An excellent overview of PU learning strategies can be found in [42]. Li et al. [20] also systematically reviewed the implementation of 29 PU learning methods in a wide range of biological topics.

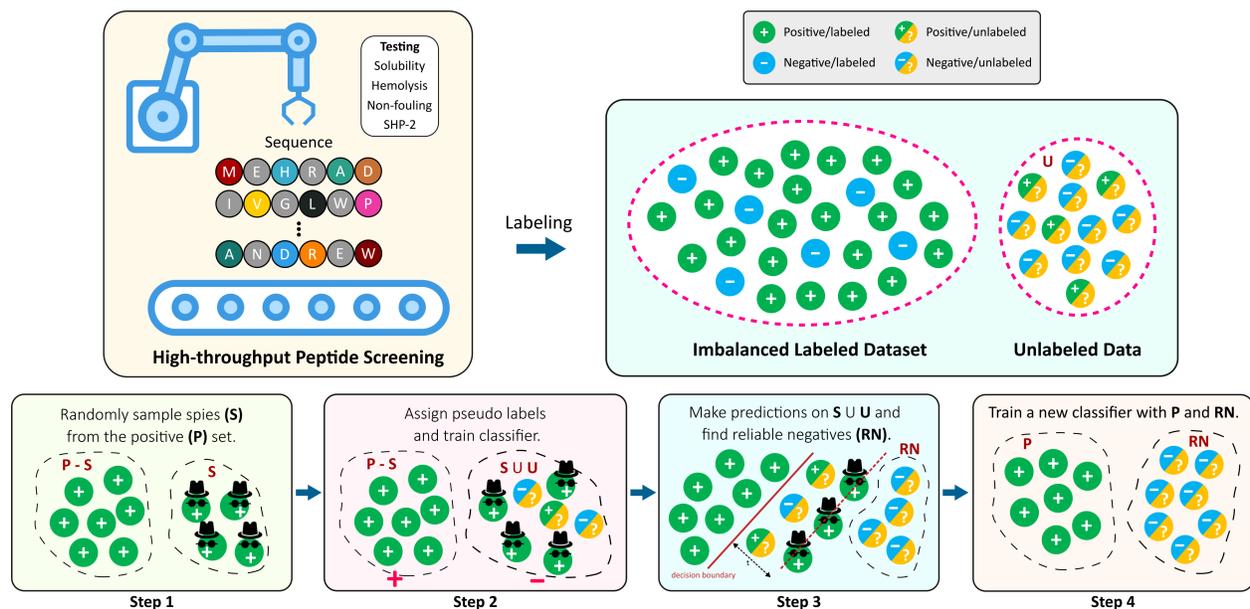


Figure 1: Overview of this work. High-throughput screening methods are commonly good at identifying positive examples, leaving imbalanced datasets (skewed towards the positive class) that are not suitable for supervised learning algorithms. In this work, we use the positive examples only to distinguish between the positive and negative samples using SPY technique.

In this work, we take advantage of the flexibility of reliable negative identification PU strategy, and discover peptide sequences that are likely to map to certain properties. Specifically, we demonstrate on a two-step technique, where Step 1 handles the deficiency of negative training examples by extracting a subset of the U samples that can be confidently labeled as N (i.e. RN). Subsequently, Step 2 involves training a deep neural network classifier using the P and the extracted RN, and applying it to the remaining pool of U. Reliable negative identification in Step 1, is an adaption of the *Spy* technique formerly employed in handling unlabeled text data [43]. In this approach, some randomly selected positive samples are defined as spies, and are intentionally mislabeled as negatives. The reliable negative examples are iteratively found within the unlabeled samples for which the posterior probability is lower than the posterior probability of the spies. We use our approach to predict different peptide properties, such as hemolysis, resistance to non-specific interactions (non-fouling), and solubility.

This manuscript is organized as follows: in Section 2, we describe the datasets, architecture of the deep learning models, and our choices for the hyperparameters. This is followed by evaluating the model in a comparative setting with the classical PN classifier in Section 3. Finally, we conclude the paper in Section. 4, with a discussion of the implications of our findings.

2 Materials and Methods

2.1 Datasets

Hemolysis

Hemolysis is referred to the disruption of erythrocyte membranes that decrease the life span of red blood cells and causes the release of Hemoglobin. It is critical to identify non-hemolytic antimicrobial peptides as a non-toxic and safe measure against bacterial infections. However, distinguishing between hemolytic and non-hemolytic peptides is a challenge, since they primarily exert their activity at the charged surface of the bacterial plasma membrane. In this work, the hemolysis classifier is trained using data from the Database of Antimicrobial Activity and Structure of Peptides (DBAASP v3 [56]). Hemolytic activity is defined by extrapolating a measurement assuming dose response curves to the point at which 50% of red blood cells are lysed. Activities below $100 \frac{\mu\text{g}}{\text{ml}}$, are considered hemolytic. The data contains 9,316 sequences (19.6% positives and 80.4% negatives) of only L- and canonical amino acids. Each measurement is treated independently, so sequences can appear multiple times. This experimental dataset contains noise, and in some observations (~40%), an identical sequence appears in both negative and positive class. As an example, sequence “RVKRVWPLVIRTVIAGYNLYRAIKKK” is found to be both hemolytic and non-hemolytic in two different lab experiments (i.e. two different training examples).

Solubility

This data contains 18,453 sequences (47.6% positives and 52.4% negatives) based on PROSO II [57], where solubility was estimated by retrospective analysis of electronic laboratory notebooks. The notebooks were part of a large effort called the Protein Structure Initiative and consider sequences linearly through the following stages: Selected, Cloned, Expressed, Soluble, Purified, Crystallized, HSQC (heteronuclear single quantum coherence), Structure, and deposited in PDB [58]. The peptides were identified as soluble or insoluble by “Comparing the experimental status at two time points, September 2009 and May 2010, we were able to derive a set of insoluble proteins defined as those which were not soluble in September 2009 and still remained in that state 8 months later.” [57]

Non-fouling

Non-fouling is defined as resistance to non-specific interactions, and this data is obtained from [59]. A non-fouling peptide (positive example) is defined using the mechanism proposed in [60]. Briefly, White et al. [60], showed that the exterior surfaces of proteins have a significantly different frequency of amino acids, and this increases in aggregation prone environments, like the cytoplasm. Synthesizing self-assembling peptides that follow this amino acid distribution and coating surfaces with the peptides creates non-fouling surfaces. This pattern was also found inside chaperone proteins, another area where resistance to non-specific interactions is important[61]. Positive data contains 3,600 sequences. Negative examples are based on 13,585 sequences (79.1% of dataset are negatives) coming from insoluble and hemolytic peptides, as well as, the scrambled positives. The scrambled negatives are generated with lengths sampled from the same length range as their respective positive set, and residues sampled from the frequency distribution of the soluble dataset. Samples are weighted to account for the class imbalance caused by the negative examples dataset size. This dataset is gathered based on the mechanism proposed in [60].

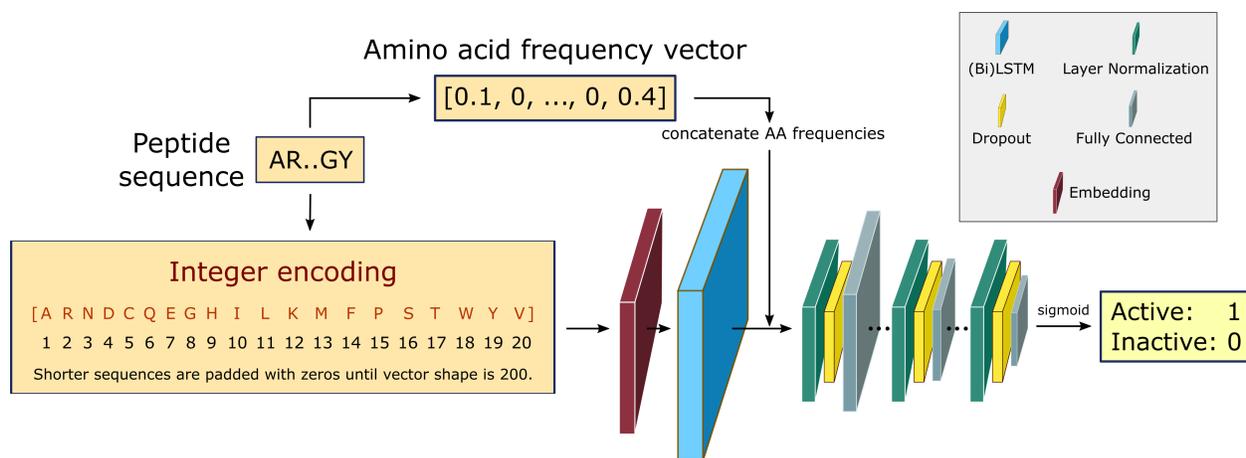


Figure 2: RNN architecture [64]. Padded integer encoded sequences are first fed to a trainable embedding layer, yielding a semantically more compact representation of the input essential amino acids. The use of bidirectional LSTMs and direct inputs of amino acid frequencies prior to the fully connected layers, improves the learning of bidirectional dependency between distant residues within a sequence. The fully connected layers are down-sized in three consecutive steps via layer normalization and dropout regularization. The final layer outputs the probability of being active for the desired training task using a sigmoid activation function.

SHP-2

SHP-2 is a ubiquitous protein tyrosine phosphatase, whose activity is regulated by phosphotyrosine (pY)-containing peptides generated in response to extracellular stimuli. SHP-2 is involved in processes such as cell growth, differentiation, migration, and immune response. [62] The SHP-2 dataset contains fixed-length peptides (5 AA residues) optimized for binding to N-SH2 domain, obtained from [63]. Total dataset size is 300, with 50% positive examples.

Table 1: Summary of used datasets. For more details, refer to [64].

	Hemolysis	Solubility	Non-fouling	SHP-2
Definition	Hemolysis is the process by which red blood cells (RBCs) rupture and release their contents, mainly Hemoglobin, into the surrounding plasma or extracellular fluid. Based on DBAASP v3 [56].	Solubility was defined in PROSO II [57] as a sequence that was transfectable, expressible, secretable, separable, and soluble in E. coli system.	Resistance to non-specific interactions. Gathered using the mechanism proposed in [60].	SHP-2 is a protein encoded by the PTPN11 gene in humans. It is a non-receptor protein tyrosine phosphatase that plays a critical role in various cellular signaling pathways. [62].
Total size	9,316	18,453	17,185	300
Positive examples	19.6%	47.6%	20.9%	50.0%
Length range	1-190 AA residues	19-198 AA residues	5-198 AA residues	5 AA residues

2.2 Model Architecture

We build a recurrent neural network (RNN) to identify the position-invariant patterns in the peptide sequences, using a sequential model from Keras framework [65] and the TensorFlow deep learning library back-end [66]. In specific, the RNN employs bidirectional Long Short Term Memory (LSTM) networks to capture long-range correlations between the amino acid residues. Compared to the conventional RNNs, LSTM networks with gate control units can learn dependency information between distant residues within peptide sequences more effectively [67–69]. An overview of the RNN architecture is shown in Figure 2. This architecture is identical to the one used in our recent work in edge-computing cheminformatics [64].

The input peptide sequences are integer encoded as vectors of shape 200, where the integer at each position in the vector corresponds to the index of the amino acid from the alphabet of the 20 essential amino acids: [A, R, N, D, C, Q, E, G, H, I, L, K, M, F, P, S, T, W, Y, V]. For implementation purposes during the training step, the maximum length of the vector is fixed at 200, padding zeros to shorter length sequences. For those sequences with shorter lengths, zeros are padded to the integer encoding representation to keep the shape fixed at 200 for all examples, to allow input sequences with flexible lengths. Every integer encoded sequence is first fed to an embedding layer with trainable weights, where the indices of discrete symbols (i.e. essential amino acids), into a representation of a fixed-length vector of defined size.

The embedding layer output either goes to a double stacked bi-LSTM layer (for solubility and hemolysis) or a single LSTM layer (for solubility and non-fouling), to identify patterns along a sequence that can be separated by large gaps. The output from the LSTM layer is then concatenated with the relative frequency of each amino acid in the input sequences. This choice is partially based on our earlier work [63], and helps with improving model performance. The concatenated output is then normalized and fed to a dropout layer with a rate of 10%, followed by a dense neural network with ReLU activation function. This is repeated three times, and the final single-node dense layer uses a sigmoid activation function to predict the peptide biological activity as the probability of the label being positive.

The hyperparameters are chosen based on a random search that resulted the best model performance in terms of the area under the receiver operating characteristic curve (AUROC) and accuracy. Readers are encouraged to refer to [64] for more details on the model architecture and its hyperparameters. We compile our Keras model using Adam optimizer [70] with a binary cross-entropy loss function, which is defined as

$$-\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)], \quad (1)$$

where y_i is the true value of the i th example, \hat{y}_i is the corresponding prediction, and N is the size of the dataset.

2.3 Positive-Unlabeled Learning

Let \vec{x} be an example, and $y \in \{0, 1\}$ the true binary label for the instance \vec{x} . If \vec{x} is a positive example, $y = 1$, otherwise $y = 0$. Let $s = 1$, if example \vec{x} is labeled, and $s = 0$, if \vec{x} is unlabeled. Only positive examples are labeled (i.e. $p(s = 1 | \vec{x}, y = 0) = 0$). In other words, the probability that a negative example appears in the labeled set is zero. On the other hand, the unlabeled set ($s = 0 | \vec{x}, y$) can contain positive ($y = 1 | \vec{x}, s = 0$) or negative ($y = 0 | \vec{x}, s = 0$) examples. The goal is to learn a probabilistic binary classifier as a function $f(\vec{x})$, such that $f(\vec{x}) = p(y = 1 | \vec{x})$, i.e. the conditioned probability of being positive given a feature vector \vec{x} .

In this work, we focus on two PU learning strategies; Adapting Base Classifier and Reliable Negative Identification.

2.3.1 Adapting Base Classifier

Adapting base classifier, also known as class prior estimation, are Bayesian-based methods that adapt the base classifier (i.e. SVM) to estimate the expected ratio of positive or negative examples in the unlabeled set. Note that in this work, we use an RNN as our base classifier. This approach simply tries to adjust the probability of being positive estimated by a traditional classifier trained with positive and unlabeled examples, where the unlabeled is treated as the negative class. The positive likelihood score $p(y = 1 | \vec{x})$ is estimated by Elkan and Noto [71] as

$$f(\vec{x}) = p(y = 1 | \vec{x}) = \frac{p(s = 1 | \vec{x})}{p(s = 1 | y = 1, \vec{x})}, \quad (2)$$

where $p(s = 1 | \vec{x})$ is the likelihood of the example \vec{x} being labeled (thus, being positive), learned from the labeled and unlabeled data. $p(s = 1 | y = 1, \vec{x})$ denotes the posterior probability of the example \vec{x} , i.e. positive sample being labeled as positive in the training data. Assuming that the labeled positive samples are chosen completely randomly from all positive examples, $p(s = 1 | y = 1, \vec{x})$ is treated as a constant factor (c) for all the samples, that can be obtained through a validation (held-out) set [53]. This “*selected completely at random*” assumption can be also written as $c = p(s = 1 | y = 1, \vec{x}) = p(s = 1 | y = 1)$, where c is a constant probability that a positive sample is labeled. This assumption is analogous to the “*missing completely at random*” assumption that is made when learning data with missing values [72–74]. Among the empirical estimators for c proposed in [71], we use the following average:

$$c = p(s = 1 | y = 1) = \frac{\sum_{x \in R} p(s = 1 | \vec{x})}{\sum_{x \in V} p(s = 1 | \vec{x})}, \quad (3)$$

where V is the validation set, drawn in the same manner as the training set, and $R \subseteq V$ is a set of positive examples in V . A threshold is adjusted within range $(0 - 1/c)$ to discriminate if the sample belongs to the positive or negative class,

by maximizing Cohen’s kappa coefficient [75]. It is important to note that the Elkan and Noto [71] algorithm was not developed to handle noisy labeled data. In addition, the theory behind its estimator limits its use to classify conditional distributions with non-overlapping support [76].

2.3.2 Reliable Negative Identification

Reliable negative identification adopts two independent algorithms: 1) identify the reliable negatives (RN) within the unlabeled set given the likelihood and 2) train a binary classifier to distinguish the labeled positive examples from the identified RN set. This approach is based on two assumptions of smoothness and separability, which simply means that all the positive examples are similar to the labeled examples, and that the negative examples are very different from them, respectively [42]. Several techniques have been proposed to extract the reliable negatives or positives from the unlabeled set, such as Spy [43], Cosine-Rocchio [77], Rocchio [44], 1DNF [78], PNLH [79], and Augmented Negatives [80], and DILCA [81].

In this work, we use Spy to find the reliable negatives. First, a small randomly selected group of positive examples (S) are removed and put in the unlabeled data as spies. This allows us to define new datasets P_s and U_s , respectively. The percentage of positive instances used as spies is defined by *spy-rate* (in this work, we use 0.2). Then, a classifier f_1 is trained based on P_s and U_s . Next, the boundary of RN under the rule that most of the spies are classified as positives is found, based on spy-tolerance (ϵ). ϵ determines what percentage of spies can remain in the unlabeled set when the decision boundary threshold (t_s) is calculated (in this work, we use 0.05). In other words, t_s is the posterior likelihood such that all added spies during training f_1 are classified as positives. All samples in U_s , whose posterior likelihood is smaller than t_s are considered *RN*. Finally, we train a new classifier f_2 given original positive samples (P) and the found *RN*.

Algorithm 1 Reliable Negative Identification with Spy

```
1: Initialize  $RN = \{\}$ 
2: Randomly sample spy-rate% from  $P$ 
3:  $P_s = P - S, U_s = U \cup S$ 
4: Assign  $P_s$  to  $y = 1$  class, and  $U_s$  to  $y = 0$  class
5: Train classifier  $f_1$  on  $P_s$  and  $U_s$ 
6: Use  $f_1$  to classify  $U_s$  and find  $p(y = 1|U_s)$ 
7: Adjust  $t_s$  such that ratio% of the positively classed  $U_s$  is less than  $\epsilon$ 
8: for  $x \in U_s$  do
9:   if  $p(y = 1|x) < t_s$  then
10:      $RN = RN \cup x$ 
11:      $U_s = U_s - x$ 
12:   end if
13: end for
14: return  $RN, U_s$ 
15: Train classifier  $f_2$  on  $P$  and  $RN$ 
```

3 Results and Discussion

In this section, we evaluate the estimated generalization error of our PU approach, and compare it with the classical PN classification, where both positive and negative examples are available for training. Note that the test data contains unobserved real positive and negative examples. We take two approaches to generate the unlabeled data: 1) Unlabeled Data Generated from Positive and Negatives Samples. In this setting, the unlabeled data is generated from a mixture of known positive and negative examples for each task. 2) Unlabeled Data Generated from Mutated Positive Samples. Given a distribution of positive examples, we generate unlabeled examples by randomly breaking the positive examples into sub-sequences, and filling up a similar-length sequence, with these sub-sequences. Duplicate sequence are removed after the generation step. This allows us to generate the unlabeled data, by creating mutations of the positive examples *without* any knowledge on what the true negative examples are.

3.1 Unlabeled Data Generated from Positive and Negatives Samples

Performance comparison between our PU learning methods and classical PN learning for different prediction tasks are presented in Table 2. Results for all the PN models are based on our earlier work in [64]. For every task, we make comparisons of the model accuracy (*ACC*%), and the area under the receiver operating characteristic curve (*AUROC*),

using the two the Adapting Base classifier, and the Reliable Negative Identification PU methods. Across all prediction tasks, with one exception of Hemolysis and Solubility with the Adapting Base Classifier method, the accuracy of our PU methods are considerably higher than the PN classification. Comparing the two PU methods, it is observed that Reliable Negative Identification outperforms Adapting Base Classifier method for all prediction tasks. Surprisingly, for the non-fouling and SHP-2 predictions, both PU methods outperform the PN classifier.

Table 2: Performance comparison between PU learning and classical PN learning for different prediction tasks, with the unlabeled data generated from positive and negatives samples. PN models are trained by having access to both positive and negative data, based on our earlier work in [64].

Task	PU Method	PU		PN	
		ACC(%)	AUROC	ACC(%)	AUROC
Hemolysis	Adapting Base Classifier	83.1	0.78	84.0	0.84
Hemolysis	Reliable Negative Identification	84.1	0.80		
Non-fouling	Adapting Base Classifier	93.8	0.93	82.0	0.93
Non-fouling	Reliable Negative Identification	95.0	0.93		
Solubility	Adapting Base Classifier	53.0	0.59	70.0	0.76
Solubility	Reliable Negative Identification	86.7	0.68		
SHP-2	Adapting Base Classifier	84.1	0.87	83.3	0.82
SHP-2	Reliable Negative Identification	90.2	0.93		

3.2 Unlabeled Data Generated from Mutated Positive Samples

Table 3 shows performance comparison between our PU learning method and classical PN learning for different prediction tasks. Considering the much better performance of Reliable Negative Identification compared to the Adapting Base Classifier observed in Table 2, we only consider the Reliable Negative Identification PU method for this unlabeled data generation scenario. Note that the solubility model in this setting showed a poor performance, and was excluded in our comparison. Considering the ACC and AUROC reported in Table 3, our PU method is able to reasonably discriminate between the positive and the reliable negatives identified from the generated unlabeled examples.

Table 3: Performance comparison between PU learning and classical PN learning for different prediction tasks, with the unlabeled data generated from mutated positive samples. Generated unlabeled is 8 times larger than the positive size. PN models are trained by having access to both positive and negative data, based on our earlier work in [64].

Task	PU Method	PU		PN	
		ACC(%)	AUROC	ACC(%)	AUROC
Hemolysis	Reliable Negative Identification	76.8	0.75	84.0	0.84
Non-fouling	Reliable Negative Identification	94.1	0.87	82.0	0.93
SHP-2	Reliable Negative Identification	84.8	0.91	83.3	0.82

It is important to note that with the unlabeled data generation, we can control how large the size of the generated unlabeled examples are. The generated unlabeled:labeled ratio reported in Table 3 is fixed at 8.0. Next, we investigate the effect of unlabeled:labeled ratio on the performance of Reliable Negative Identification strategy across all prediction tasks in Figure 3. Each point represents the average value of AUROC and ACC% (left and right panel, respectively) over 6 models trained with a different choice of randomly selected spy positives, and error bars show the magnitude of the standard deviation. Horizontal dashed lines show the performance of the PN classifier for each task represented as a baseline for performance comparison. With very small generated unlabeled samples (i.e. unlabeled:labeled ratio ≈ 2.0), the exploration of new examples that can qualify as reliable negatives will be largely limited. Thus, the trained f_2 classifier has a significantly lower performance compared to the baseline PN classifier and to the other PU models trained with higher generated unlabeled:labeled ratios. With larger unlabeled:labeled ratios (i.e. > 10.0), we see a better prediction performance across all the tasks. There are two significant observations; 1. With more unlabeled sequences generated, the trained PU models have a competitive performance with the PN models. In specific, for binding against SHP2, we observe that the PU model beats the PN classifier in both AUROC and ACC%. 2. Surprisingly, the PU models become more confident in their predictions with the increase of the unlabeled:labeled ratio (compare magnitude of error bars in Figure 3). This can bring a major advantage in implementing our approach in a generative setting, where we can

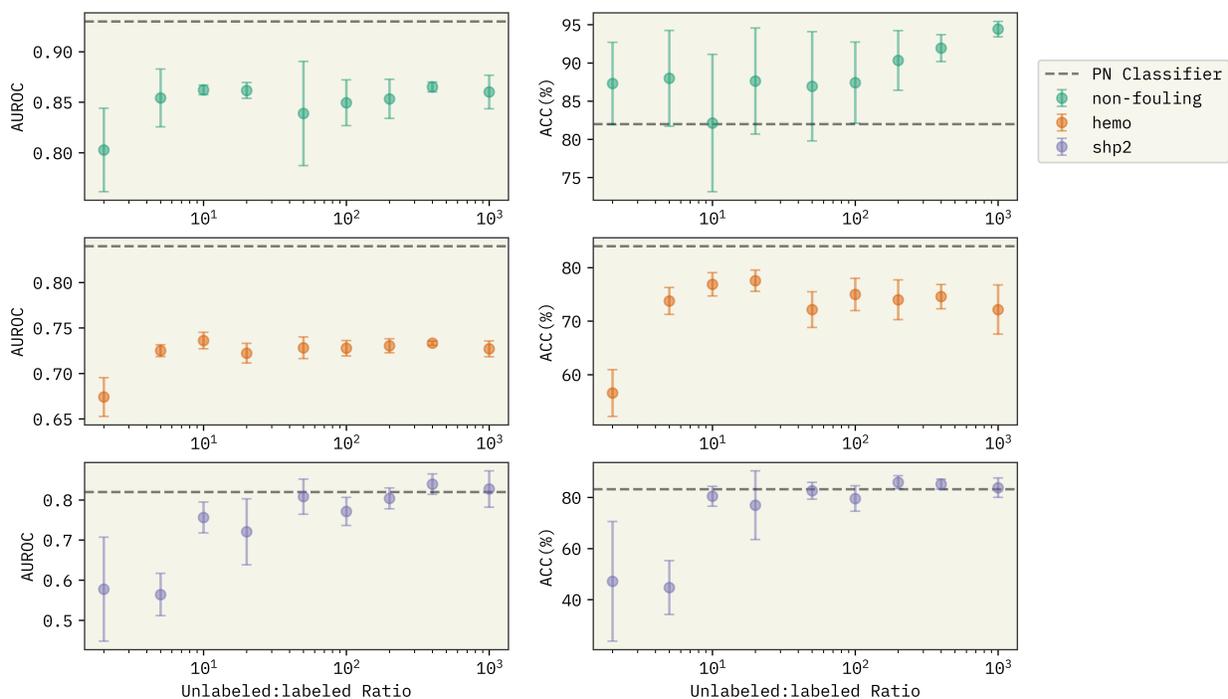


Figure 3: Effect of generated unlabeled:labeled ratio on the performance of the Reliable Negative Identification strategy for the three prediction tasks. Horizontal dashed lines show the performance of the PN classifier from Table 3 used as a baseline for comparison. At the low ratio regime, the pool of unlabeled data is not big enough to obtain promising candidates as reliable negatives. With larger unlabeled:labeled ratios, the PU model gets to identify a better choice of sequences as reliable negatives, despite the major existing class imbalance in the training data.

predict the properties of new peptide sequences without having to worry much about the class imbalance between the positive and the negative examples, which can majorly reduce model performance, if the learning is supervised.

Comparing AUROC and ACC in Tables 2 and 3, we observe that Reliable Negative Identification with mutated positive samples has a relative lower performance compared to the other scenario, where the unlabeled data is generated from a distribution of positive and negative examples. Despite this minor lower performance, using the new unlabeled sequence generation, one can explore the newly unlabeled samples, and make predictions on peptide properties by only having access to the examples from one class (i.e. positive). The sequence-based peptide property prediction in this work is limited to four different tasks. However, with the positive data available, this work can be further extended to developing predictive models for inferring other peptide properties.

4 Conclusions

We've showed a semi-supervised learning framework to infer the mapping from peptides' sequence to function for properties such as hemolysis, solubility, non-fouling, and binding against SHP-2. Our positive unlabeled learning method aims at identifying likely negative candidates (reliable negatives), from the generated unlabeled sequences, given random permutations of subsequences within the available positive samples. The reliable negative identification strategy is agnostic with respect to the model architecture used, giving generality. Our method will be most beneficial in biology screening experiments, where most high-throughput screening methods solely focus on identifying the positive example. All PU models showed a comparative predictive ability and robustness across the different prediction tasks, when compared to training with both positive and negative examples. This learning strategy can provide a robust feasible path towards estimating how amino acids positional substitutions can affect peptide's functional response for unknown sequences, and accelerate the design and discovery of novel therapeutics.

Acknowledgements

Research reported in this work was supported by the National Institute of General Medical Sciences of the National Institutes of Health under award number R35GM137966. We thank the Center for Integrated Research Computing (CIRC) at University of Rochester for providing computational resources and technical support.

Data and Code Availability

All data and code used to produce results in this study are publically available in the following GitHub repository: <https://github.com/ur-whitelab/pu-peptides>.

References

- [1] Jeffrey B Sperry, Christopher J Minter, JingYa Tao, Rebecca Johnson, Remzi Duzguner, Michael Hawksworth, Samantha Oke, Paul F Richardson, Richard Barnhart, David R Bill, et al. Thermal stability assessment of peptide coupling reagents commonly used in pharmaceutical manufacturing. *Organic Process Research & Development*, 22(9):1262–1275, 2018.
- [2] Lucia Ferrazzano, Dario Corbisiero, Giulia Martelli, Alessandra Tolomelli, Angelo Viola, Antonio Ricci, and Walter Cabri. Green solvent mixtures for solid-phase peptide synthesis: A dimethylformamide-free highly efficient synthesis of pharmaceutical-grade peptides. *ACS Sustainable Chemistry & Engineering*, 7(15):12867–12877, 2019.
- [3] Markus Muttenthaler, Glenn F King, David J Adams, and Paul F Alewood. Trends in peptide drug discovery. *Nature Reviews Drug Discovery*, 20(4):309–325, 2021.
- [4] Daniel J Drucker. Advances in oral peptide therapeutics. *Nature reviews Drug discovery*, 19(4):277–289, 2020.
- [5] Kohei Sato, Mark P Hendricks, Liam C Palmer, and Samuel I Stupp. Peptide supramolecular materials for therapeutics. *Chemical Society Reviews*, 47(20):7539–7551, 2018.
- [6] Fatemeh Araste, Khalil Abnous, Maryam Hashemi, Seyed Mohammad Taghdisi, Mohammad Ramezani, and Mona Alibolandi. Peptide-based targeted therapeutics: Focus on cancer treatment. *Journal of controlled release*, 292:141–162, 2018.
- [7] Brian P Lazzaro, Michael Zasloff, and Jens Rolff. Antimicrobial peptides: Application informed by evolution. *Science*, 368(6490):eaau5480, 2020.
- [8] Annika Nelde, Tatjana Bilich, Jonas S Heitmann, Yacine Maringer, Helmut R Salih, Malte Roerden, Maren Lübke, Jens Bauer, Jonas Rieth, Marcel Wacker, et al. Sars-cov-2-derived peptides define heterologous and covid-19-induced t cell recognition. *Nature immunology*, 22(1):74–85, 2021.
- [9] Natsuki Furukawa and Aleksander S Popel. Peptides that immunoactivate the tumor microenvironment. *Biochimica et Biophysica Acta (BBA)-Reviews on Cancer*, 1875(1):188486, 2021.
- [10] L Zhang, JR Lu, and TA Waigh. Electronics of peptide-and protein-based biomaterials. *Advances in Colloid and Interface Science*, 287:102319, 2021.
- [11] Janna N Sloand, Michael A Miller, and Scott H Medina. Fluorinated peptide biomaterials. *Peptide Science*, 113(2):e24184, 2021.
- [12] Christina Karavasili and Dimitrios G Fatouros. Self-assembling peptides as vectors for local drug delivery and tissue engineering applications. *Advanced drug delivery reviews*, 174:387–405, 2021.
- [13] Andy Chi-Lung Lee, Janelle Louise Harris, Kum Kum Khanna, and Ji-Hong Hong. A comprehensive review on current advances in peptide drug development and design. *International journal of molecular sciences*, 20(10):2383, 2019.
- [14] Ryan Poplin, Pi-Chuan Chang, David Alexander, Scott Schwartz, Thomas Colthurst, Alexander Ku, Dan Newburger, Jojo Dijamco, Nam Nguyen, Pegah T Afshar, et al. A universal snp and small-indel variant caller using deep neural networks. *Nature biotechnology*, 36(10):983–987, 2018.
- [15] Haoyang Zeng and David K. Gifford. Predicting the impact of non-coding variants on DNA methylation. *Nucleic Acids Research*, 45(11):e99–e99, 03 2017. ISSN 0305-1048. doi: 10.1093/nar/gkx177. URL <https://doi.org/10.1093/nar/gkx177>.
- [16] Maxat Kulmanov, Mohammed Asif Khan, and Robert Hoehndorf. Deepgo: predicting protein functions from sequence and interactions using a deep ontology-aware classifier. *Bioinformatics*, 34(4):660–668, 2018.

- [17] Buzhong Zhang, Jinyan Li, and Qiang Lü. Prediction of 8-state protein secondary structures by a novel deep learning architecture. *BMC bioinformatics*, 19(1):1–13, 2018.
- [18] Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. Broadly applicable and accurate protein design by integrating structure prediction networks and diffusion generative models. *bioRxiv*, pages 2022–12, 2022.
- [19] Somaye Hashemifar, Behnam Neyshabur, Aly A Khan, and Jinbo Xu. Predicting protein–protein interactions through sequence-based deep learning. *Bioinformatics*, 34(17):i802–i810, 2018.
- [20] Fuyi Li, Shuangyu Dong, André Leier, Meiya Han, Xudong Guo, Jing Xu, Xiaoyu Wang, Shirui Pan, Cangzhi Jia, Yang Zhang, et al. Positive-unlabeled learning in bioinformatics and computational biology: a brief review. *Briefings in bioinformatics*, 23(1):bbab461, 2022.
- [21] Katarzyna Sidorczuk, Przemyslaw Gagat, Filip Pietluch, Jakub Kala, Dominik Rafacz, Laura Bakala, Jadwiga Slowik, Rafal Kolenda, Stefan Roediger, Legana CHW Fingerhut, et al. The impact of negative data sampling on antimicrobial peptide prediction. *bioRxiv*, 2022.
- [22] Evgeny Byvatov and Gisbert Schneider. Support vector machine applications in bioinformatics. *Applied bioinformatics*, 2(2):67–77, 2003.
- [23] Anne-Laure Boulesteix, Silke Janitza, Jochen Kruppa, and Inke R König. Overview of random forest methodology and practical guidance with emphasis on computational biology and bioinformatics. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 2(6):493–507, 2012.
- [24] Tong Tong Wu, Yi Fang Chen, Trevor Hastie, Eric Sobel, and Kenneth Lange. Genome-wide association analysis by lasso penalized logistic regression. *Bioinformatics*, 25(6):714–721, 2009.
- [25] Qiong Wang, George M Garrity, James M Tiedje, and James R Cole. Naive bayesian classifier for rapid assignment of rna sequences into the new bacterial taxonomy. *Applied and environmental microbiology*, 73(16):5261–5267, 2007.
- [26] Peng Yang, Xiao-Li Li, Jian-Ping Mei, Chee-Keong Kwoh, and See-Kiong Ng. Positive-unlabeled learning for disease gene identification. *Bioinformatics*, 28(20):2640–2647, 2012.
- [27] Akram Vasighizaker, Alok Sharma, and Abdollah Dehzangi. A novel one-class classification approach to accurately predict disease-gene association in acute myeloid leukemia cancer. *PloS one*, 14(12):e0226115, 2019.
- [28] Zhazhan Cheng, Shuigeng Zhou, and Jihong Guan. Computationally predicting protein-rna interactions using only positive and unlabeled examples. *Journal of bioinformatics and computational biology*, 13(03):1541005, 2015.
- [29] Hyebin Song, Bennett J Bremer, Emily C Hinds, Garvesh Raskutti, and Philip A Romero. Inferring protein sequence-function relationships with large-scale positive-unlabeled learning. *Cell Systems*, 12(1):92–101, 2021.
- [30] Fuyi Li, Yang Zhang, Anthony W Purcell, Geoffrey I Webb, Kuo-Chen Chou, Trevor Lithgow, Chen Li, and Jiangning Song. Positive-unlabelled learning of glycosylation sites in the human proteome. *BMC bioinformatics*, 20(1):1–17, 2019.
- [31] Hongfang Liu, Manabu Torii, Guixian Xu, Zhangzhi Hu, and Johannes Goll. Learning from positive and unlabeled documents for retrieval of bacterial protein-protein interaction literature. In *Linking Literature, Information, and Knowledge for Biology*, pages 62–70. Springer, 2010.
- [32] Cumhur Kılıç and Mehmet Tan. Positive unlabeled learning for deriving protein interaction networks. *Network modeling analysis in health informatics and bioinformatics*, 1(3):87–102, 2012.
- [33] Pathima Nusrath Hameed, Karin Verspoor, Snezana Kusljic, and Saman Halgamuge. Positive-unlabeled learning for inferring drug interactions based on heterogeneous attributes. *BMC bioinformatics*, 18(1):1–15, 2017.
- [34] Chunlin Wang, Chris Ding, Richard F Meraz, and Stephen R Holbrook. Psol: a positive sample only learning algorithm for finding non-coding rna genes. *Bioinformatics*, 22(21):2590–2596, 2006.
- [35] Xing-Ming Zhao, Yong Wang, Luonan Chen, and Kazuyuki Aihara. Gene function prediction using labeled and unlabeled data. *BMC bioinformatics*, 9(1):1–14, 2008.
- [36] Nitin Bhardwaj, Mark Gerstein, and Hui Lu. Genome-wide sequence-based prediction of peripheral proteins using a novel semi-supervised learning technique. *BMC bioinformatics*, 11(1):1–8, 2010.
- [37] Juan Fernando López, Jesús Alfonso López Sotelo, Diogo Leite, and Carlos Peña-Reyes. Applying one-class learning algorithms to predict phage-bacteria interactions. In *2019 IEEE Latin American Conference on Computational Intelligence (LA-CCI)*, pages 1–6. IEEE, 2019.

- [38] Yuanyuan Xiao and Mark R Segal. Biological sequence classification utilizing positive and unlabeled data. *Bioinformatics*, 24(9):1198–1205, 2008.
- [39] Pratiti Bhadra, Jielu Yan, Jinyan Li, Simon Fong, and Shirley WI Siu. Ampep: Sequence-based prediction of antimicrobial peptides using distribution patterns of amino acid properties and random forest. *Scientific reports*, 8(1):1–10, 2018.
- [40] Zhi-Hua Zhou. A brief introduction to weakly supervised learning. *National science review*, 5(1):44–53, 2018.
- [41] Jieyu Zhang, Cheng-Yu Hsieh, Yue Yu, Chao Zhang, and Alexander Ratner. A survey on programmatic weak supervision. *arXiv preprint arXiv:2202.05433*, 2022.
- [42] Jessa Bekker and Jesse Davis. Learning from positive and unlabeled data: A survey. *Machine Learning*, 109(4):719–760, 2020.
- [43] Bing Liu, Wee Sun Lee, Philip S Yu, and Xiaoli Li. Partially supervised classification of text documents. In *ICML*, volume 2, pages 387–394. Sydney, NSW, 2002.
- [44] Xiaoli Li and Bing Liu. Learning to classify texts using positive and unlabeled data. In *IJCAI*, volume 3, pages 587–592, 2003.
- [45] Wee Sun Lee and Bing Liu. Learning with positive and unlabeled examples using weighted logistic regression. In *ICML*, volume 3, pages 448–455, 2003.
- [46] Bing Liu, Yang Dai, Xiaoli Li, Wee Sun Lee, and Philip S Yu. Building text classifiers using positive and unlabeled examples. In *Third IEEE international conference on data mining*, pages 179–186. IEEE, 2003.
- [47] Hang Wei, Yong Xu, and Bin Liu. ipidi-pul: identifying piwi-interacting rna-disease associations based on positive unlabeled learning. *Briefings in Bioinformatics*, 22(3):bbaa058, 2021.
- [48] Yiming Chen, Zhoujun Li, Xiaofeng Wang, Jiali Feng, and Xiaohua Hu. Predicting gene function using few positive examples and unlabeled ones. *BMC genomics*, 11(2):1–9, 2010.
- [49] Peng Yang, Xiaoli Li, Hon-Nian Chua, Chee-Keong Kwoh, and See-Kiong Ng. Ensemble positive unlabeled learning for disease gene identification. *PLoS one*, 9(5):e97079, 2014.
- [50] Gholam-Hossein Jowkar and Eghbal G Mansoori. Perceptron ensemble of graph-based positive-unlabeled learning for disease gene identification. *Computational biology and chemistry*, 64:263–270, 2016.
- [51] Fangfang Yan, Zhongming Zhao, and Lukas M Simon. Emptynn: A neural network based on positive and unlabeled learning to remove cell-free droplets and recover lost cells in scRNA-seq data. *Patterns*, 2(8):100311, 2021.
- [52] Ryuichi Kiryo, Gang Niu, Marthinus C Du Plessis, and Masashi Sugiyama. Positive-unlabeled learning with non-negative risk estimator. *Advances in neural information processing systems*, 30, 2017.
- [53] Luigi Cerulo, Charles Elkan, and Michele Ceccarelli. Learning gene regulatory networks from only positive and unlabeled data. *BMC bioinformatics*, 11(1):1–16, 2010.
- [54] Vikas Pejaver, Jorge Urresti, Jose Lugo-Martinez, Kymberleigh A Pagel, Guan Ning Lin, Hyun-Jun Nam, Matthew Mort, David N Cooper, Jonathan Sebat, Lilia M Iakoucheva, et al. Inferring the molecular and phenotypic impact of amino acid variants with mutpred2. *Nature communications*, 11(1):1–13, 2020.
- [55] Zhenfeng Li, Lun Hu, Zehai Tang, and Cheng Zhao. Predicting hiv-1 protease cleavage sites with positive-unlabeled learning. *Frontiers in Genetics*, 12:658078, 2021.
- [56] Malak Pirtskhalava, Anthony A Armstrong, Maia Grigolava, Mindia Chubinidze, Evgenia Alimbarashvili, Boris Vishnepolsky, Andrei Gabrielian, Alex Rosenthal, Darrell E Hurt, and Michael Tartakovsky. Dbaasp v3: database of antimicrobial/cytotoxic activity and structure of peptides as a resource for development of new therapeutics. *Nucleic acids research*, 49(D1):D288–D297, 2021.
- [57] Pawel Smialowski, Gero Doose, Philipp Torkler, Stefanie Kaufmann, and Dmitriy Frishman. Proso ii—a new method for protein solubility prediction. *The FEBS journal*, 279(12):2192–2200, 2012.
- [58] Helen M Berman, John D Westbrook, Margaret J Gabanyi, Wendy Tao, Raship Shah, Andrei Kouranov, Torsten Schwede, Konstantin Arnold, Florian Kiefer, Lorenza Bordoli, et al. The protein structure initiative structural genomics knowledgebase. *Nucleic acids research*, 37(suppl_1):D365–D368, 2009.
- [59] Rainier Barrett, Shaoyi Jiang, and Andrew D White. Classifying antimicrobial and multifunctional peptides with bayesian network models. *Peptide Science*, 110(4):e24079, 2018.
- [60] Andrew D White, Ann K Nowinski, Wenjun Huang, Andrew J Keefe, Fang Sun, and Shaoyi Jiang. Decoding nonspecific interactions from nature. *Chemical Science*, 3(12):3488–3494, 2012.

- [61] Andrew D White, Wenjun Huang, and Shaoyi Jiang. Role of nonspecific interactions in molecular chaperones through model-based bioinformatics. *Biophysical journal*, 103(12):2484–2491, 2012.
- [62] Michelangelo Marasco, John Kirkpatrick, Vittoria Nanna, Justyna Sikorska, and Teresa Carlomagno. Phosphotyrosine couples peptide binding and shp2 activation via a dynamic allosteric network. *Computational and structural biotechnology journal*, 19:2398–2415, 2021.
- [63] Rainier Barrett and Andrew D White. Investigating active learning and meta-learning for iterative peptide design. *Journal of chemical information and modeling*, 61(1):95–105, 2020.
- [64] Mehrad Ansari and Andrew D White. Serverless prediction of peptide properties with recurrent neural networks. *bioRxiv*, 2022.
- [65] François Chollet. Keras. <https://github.com/fchollet/keras>, 2015.
- [66] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- [67] Ilya Sutskever, James Martens, and Geoffrey E Hinton. Generating text with recurrent neural networks. In *ICML*, 2011.
- [68] Marwin HS Segler, Thierry Kogej, Christian Tyrchan, and Mark P Waller. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS central science*, 4(1):120–131, 2018.
- [69] Yilin Ye, Jian Wang, Yunwan Xu, Yi Wang, Youdong Pan, Qi Song, Xing Liu, and Ji Wan. Mathla: a robust framework for hla-peptide binding prediction integrating bidirectional lstm and multiple head attention mechanism. *BMC bioinformatics*, 22(1):1–12, 2021.
- [70] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [71] Charles Elkan and Keith Noto. Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 213–220, 2008.
- [72] Roderick JA Little and Donald B Rubin. *Statistical analysis with missing data*, volume 793. John Wiley & Sons, 2019.
- [73] Andrew Smith and Charles Elkan. A bayesian network framework for reject inference. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 286–295, 2004.
- [74] Andrew T Smith and Charles Elkan. Making generative classifiers robust to selection bias. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 657–666, 2007.
- [75] Carmen Esposito, Gregory A Landrum, Nadine Schneider, Nikolaus Stiefl, and Sereina Riniker. Ghost: adjusting the decision threshold to handle imbalanced data in machine learning. *Journal of Chemical Information and Modeling*, 61(6):2623–2640, 2021.
- [76] Shantanu Jain, Martha White, and Predrag Radivojac. Estimating the class prior and posterior from noisy positives and unlabeled data. *Advances in neural information processing systems*, 29, 2016.
- [77] Hwanjo Yu, Jiawei Han, and KC-C Chang. Pebl: Web page classification without negative examples. *IEEE Transactions on Knowledge and Data Engineering*, 16(1):70–81, 2004.
- [78] Tao Peng, Wanli Zuo, and Fengling He. Svm based adaptive learning method for text classification from positive and unlabeled documents. *Knowledge and Information Systems*, 16(3):281–301, 2008.
- [79] Gabriel Pui Cheong Fung, Jeffrey Xu Yu, Hongjun Lu, and Philip S Yu. Text classification without negative examples revisit. *IEEE transactions on Knowledge and Data Engineering*, 18(1):6–20, 2005.
- [80] Xiao-Li Li and Bing Liu. Learning from positive and unlabeled examples with different data distributions. In *European conference on machine learning*, pages 218–229. Springer, 2005.
- [81] Dino Ienco, Ruggero G Pensa, and Rosa Meo. From context to distance: Learning dissimilarity for categorical data clustering. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(1):1–25, 2012.