

# Class Balancing for Efficient Active Learning in Imbalanced Datasets

**Yaron Fairstein**

Amazon / yyfairstein@gmail.com

**Zohar Karnin**

Amazon / zkarnin@amazon.com

**Alexander Libov**

Amazon / alibov@amazon.com

**Oren Kalinsky**

Amazon / orenkalinsky@gmail.com

**Guy Kushilevitz**

Amazon / guyk@amazon.com

**Sofia Tolmach**

Amazon / sofiato@amazon.com

## Abstract

Recent developments in active learning algorithms for NLP tasks show promising results in terms of reducing labelling complexity. In this paper we extend this effort to imbalanced datasets; we bridge between the active learning approach of obtaining diverse and informative examples, and the heuristic of class balancing used in imbalanced datasets. We develop a novel tune-free weighting technique that can be applied to various existing active learning algorithms, adding a component of class balancing. We compare several active learning algorithms to their modified version on multiple public datasets and show that when the classes are imbalanced, with manual annotation effort remaining equal the modified version significantly outperforms the original both in terms of the test metric and the number of obtained minority examples. Moreover, when the imbalance is mild or non-existent (classes are completely balanced), our technique does not harm the base algorithms.

## 1 Introduction

Pre-trained Language Models (PLMs) and Masked Language Models (Peters et al., 2018; Devlin et al., 2019; Liu et al., 2019) have revolutionized NLP by supplying meaningful contextual embeddings for tokens and sentences. These models pushed SoTA on many tasks, proving especially effective for text classification tasks (Sun et al., 2019). For many of these tasks, given enough labeled data for fine-tuning a PLM, satisfactory results can be achieved. On the other hand, due to the rapid growth in textual datasets, for many domains and tasks it is often the case that while a vast amount of textual data is available, only a small portion of the text-instances belong to a specific class of interest (Li et al., 2018; Padurariu and Breaban, 2019; Shaikh et al., 2021).

This sparsity makes annotation challenging because naive data sampling methods do not produce enough examples belonging to a class of interest.

It is often possible to use some filtering heuristic before performing manual annotations, so that the percentage of data belonging to the class in question will rise. This approach however is not scalable, since it requires hand-crafting for each specific task or even class. In addition, when filtering is imperfect, the PLM may never observe an important cluster of examples, hurting its generalization capabilities. Therefore, It is essential to minimize the amount of human annotations needed in order to acquire labeled data that will be sufficient for the model to reach a satisfying performance. This calls for the use of Active Learning (AL) techniques. There are many works applying AL techniques for PLM fine-tuning (Gissin and Shalev-Shwartz, 2019; Ash et al., 2020; Dor et al., 2020) but most do not specifically consider the imbalanced dataset case.

Standard AL objectives include obtaining samples which the model is less certain about and increasing the diversity of the chosen sample set. Existing AL methods (C Lin, 2018; Kim and Yoo, 2022) aimed to handle imbalanced datasets suggest doing so by alternating between some approach designed to obtain positive examples<sup>1</sup>, and another approach with standard AL objectives. The shortcoming of this alternating approach is that the exact balance of objectives, i.e., how many samples should be chosen aimed to be positive vs diverse, etc., is dependent on the imbalance of the dataset. As a result, existing methods either require a balancing hyperparameter that cannot be tuned in a practical<sup>2</sup> setting, or need to learn the correct balancing ratio over time, making the methods efficient only after a large amount of labelling.

In this work we take a different approach that modifies an AL algorithm to favor positive exam-

<sup>1</sup>Throughout this paper we assume w.l.o.g. that the positive class is the minority class.

<sup>2</sup>AL is most needed when labeled data is missing, therefore assuming a dev set is impractical.

ples in a way that is self-tuned to obtain an equal amount of positive and negative samples. The core idea of our approach is to debias the inherent skew in the imbalanced dataset by utilizing the inverse-propensity score of the probability of a sample to be positive, according to the model obtained thus far. The simplest form of our algorithm is to sample according to these scores. Intuitively, once the model is reasonable, and this happens quickly in the era of pre-trained models, the amount of positive samples is in the same order of magnitude as the negative samples. We show that these scores can be used not only as a tool to modify the uniform sampling approach, but can be combined with other algorithms providing their own weights to the data points (Yuan et al., 2020), or clustering approaches aimed to maximize diversity (Gissin and Shalev-Shwartz, 2019; Ash et al., 2020).

We demonstrate that for the imbalanced setting, our adaptation of SoTA AL algorithms outperform their original counterparts in multiple datasets. We show that this behavior is consistent across various levels of data imbalance and that in the balanced setting, our adaptation is comparable to the original AL algorithms. The latter property is crucial from a practical perspective as the imbalance ratio of a dataset is rarely known in advance and can be difficult to estimate.

Concluding, we showcase a weakness of recent AL approaches when dealing with imbalanced datasets and devise a novel self-tuned re-weighting solution that complements existing (including SoTA) AL algorithms, improving their performance on imbalanced datasets. We experiment on four imbalanced datasets from different tasks and domains. We publish our code and imbalanced datasets for reproducibility and to encourage future research in this area<sup>3</sup>.

## 2 Related Work

**Conventional active learning.** There is a myriad of works on active learning in different settings (Ren et al., 2021; Fu et al., 2013). Herein, we consider a pool-based active learning scenario where the algorithm has access to a large unlabeled dataset  $\mathcal{U}$  that can be labeled through human annotations. As previously mentioned, our solution modifies existing active learning algorithms in order to better deal with imbalanced datasets. We identify three types of active learning algorithms:

random, embedding-based, and score-based algorithms. Random sampling, the simplest approach, generates a distribution over  $\mathcal{U}$  from which  $k$  distinct samples are randomly selected (e.g., uniform sampling).

Embedding-based algorithms embed unlabeled samples in a high dimensional space. A subset of samples is then selected for annotation using a clustering algorithm based on their embedding, to increase the sample diversity. BADGE (Ash et al., 2020) is an embedding-based algorithm where the sample gradient from the last layer of the model, taken from the log loss of the predicted label, is used as an embedding. ALPS (Gissin and Shalev-Shwartz, 2019) adopts a similar approach, but uses an embedding that captures the language model uncertainty of the different tokens in the sentence.

Score-based algorithms attach a model score to each sample. The most common score-based approach is Least Confidence (Lewis and Gale, 1994) which selects the top-K samples with the highest model uncertainty. A more recent work called DAL (Gissin and Shalev-Shwartz, 2019) trains a classifier to detect new samples dissimilar from the already labeled data.

**Active learning in skewed datasets.** Compared to active learning, relatively only a few studies have specifically addressed the issue of class imbalance. A common approach in handling class imbalance is to skew the active learning algorithm towards ranking positive samples higher. ODAL (Barata et al., 2021) is an extension of DAL that targets the cold-start setting using an outlier detection algorithm until a single positive sample is found. HAL (Kazerouni et al., 2020) suggests a Hybrid Active Learning algorithm that switches between selecting ‘exploitation’ that chooses points where the model is uncertain and ‘exploration’ that either chooses samples uniformly at random or according to how far they are from already chosen samples, yet this approach requires a hyperparameter determining how to balance exploration and exploitation, that is crucial to the performance, and setting it is said to be an open problem.

Recent works try to combine conventional active learning algorithms with positive-skewed sampling. C Lin (2018) explore a method that combines multiple strategies for selecting new examples such as standard active learning, sampling in a skewed way towards positive examples, and generating examples of the minority type (via manual annotators).

<sup>3</sup><https://github.com/balancingAL/ImbAL>

They apply a Multi-armed bandit (MAB) algorithm (UCB) to balance between the strategies. Similarly, BMP (Kim and Yoo, 2022) defines two types of policies, one samples from the positive class and the other uses a standard AL algorithm such as random or VE (Beluch et al., 2018), and uses a MAB algorithm to dynamically allocate batches to the different policies. It is folklore knowledge that MAB algorithms require a large number of rounds (at least 10s) before they can learn anything meaningful<sup>4</sup>, and indeed the number of rounds in the mentioned papers are 55 or more. Having a large number of rounds can be computationally expensive as SoTA active learning algorithms require applying a large ML model on all examples, as well as time consuming in a human annotation setting. We thus focus on a small number of rounds (say, 5), making these works less applicable.

### 3 Setting

We consider the pool-based active learning setting, where we are given a pool of unlabeled samples  $\mathcal{U}$ . At each iteration  $t$ , some subset  $S_t \subset \mathcal{U}$  of size  $k$  is selected for annotation. This subset, sometimes together with  $S_1, \dots, S_{t-1}$ , is used to train a model  $M_t$ . The subset  $S_t$  is selected by a sampling algorithm  $\mathcal{A}$ , which is usually dependent on the model  $M_{t-1}$ . A formal description of this process is given in Algorithm 1. The goal in this setting is to select a subset  $\mathcal{D} \subset \mathcal{U}$  to be labeled, such that a model trained on  $\mathcal{D}$  optimizes some metric on a separate test set. We focus on the binary classification problem, as it includes multiple common and important imbalanced scenarios such as phishing attempts, fraud, etc. (Kazerouni et al., 2020; Barata et al., 2021; C Lin, 2018).

### 4 Our approach

Consider some iteration  $t$  where we have at our disposal a model  $M_{t-1}$  trained on samples labelled in previous iterations.<sup>5</sup> For a sample  $x \in \mathcal{U}$ , we consider  $M_{t-1}(x)$  as an approximation of the probability of  $x$  being a positive example. If we uti-

<sup>4</sup>This follows from a simple statistics exercise showing that with a handful of rounds, the posterior distributions of the reward of each option is very close to the prior, or that the confidence intervals remain large.

<sup>5</sup>We assume that the process begins with a model trained on a handful of examples. This is a realistic setting, as it is rarely the case that a practitioner will aim to solve a classification problem without having a single labelled example. Alternatively, we can begin with uniform weights. We discuss this issue in Section 5.9.

---

#### Algorithm 1 Pool-based Active Learning

---

**Input:** Unlabeled data pool  $\mathcal{U}$ , number of samples per iteration  $k$ , number of iterations  $T$ , sampling algorithm  $\mathcal{A}$ , classifier model training algorithm  $\mathcal{M}$  and seed samples  $\mathcal{S}$ .

```

 $\mathcal{D} \leftarrow \mathcal{S}$ 
 $M_0 = \mathcal{M}(\mathcal{D})$  ▷ Initiate classifier model
for  $t \leftarrow 1 : T$  do
   $S_t = \mathcal{A}(M_{t-1}, \mathcal{U}, k)$  ▷ Sampling step
  Annotate  $S_t$ 
   $\mathcal{U} \leftarrow \mathcal{U} \setminus S_t$ 
   $\mathcal{D} \leftarrow \mathcal{D} \cup S_t$ 
   $M_t = \mathcal{M}(\mathcal{D})$  ▷ Learning step
Return  $M_T$ 

```


---

lize the naive random algorithm to select samples  $x \in \mathcal{U}$ , then  $M_{t-1}(x)$  is a random variable in  $[0, 1]$ . Our goal is to construct a distribution  $p$  over  $\mathcal{U}$  such that the resulting random variable  $M_{t-1}(x)$  is uniformly distributed in  $[0, 1]$ . To motivate this idea, consider a case where the model is perfectly calibrated, meaning the probability it provides is correct. Here, if sampled according to  $p$ , half of the examples, in expectation, will be positive (in Observation 1 we provide an additional bound under a weaker assumption). Another advantage of such a distribution appears even in the case of balanced datasets. Consider a case where many of the examples are easy to classify. It is likely that in an early stage easy samples will be concentrated near 0 and 1 in the distribution of  $M_{t-1}(x)$ . This means that our sample strategy will upweight exactly the problematic points where the model is uncertain, such that we obtain the property of biasing towards informative examples.


The distribution  $p$  described above can be achieved by weighting each sample inversely to its propensity results in a uniform sample. Thus, we call this method **InvProp**. In terms of implementation, we discretize the interval  $[0, 1]$  into  $m$  bins uniformly, meaning  $[0, \frac{1}{m}), [\frac{1}{m}, \frac{2}{m}), \dots, [\frac{m-1}{m}, 1]$ , as can be seen in Figure 1(a). For bin  $i$  that contains  $b_i$  samples we give each sample a probability proportional to  $\frac{1}{b_i}$  (see Figure 1(b)).

The following lemma demonstrates the robustness of InvProp-weights to the exact, and unknown, degree of imbalance in the data.

**Lemma 1.** *Let  $\mathcal{A}$  be the random algorithm with InvProp weights defined over 2 bins, and let  $M$  be a model with precision  $\alpha$ . Also, let  $S = \mathcal{A}(M, \mathcal{U}, k)$*



(a) Histogram of confidence scores  $M_t(x)$ .



(b) Inverse propensity distribution  $p$ .

Figure 1: Confidence scores histogram and the probability fitting the InvProp-weights.

be the subset of size  $k$  sampled by  $\mathcal{A}$ . If the number of samples in bin  $[1/2, 1]$  is at least  $\beta \cdot k$  for  $\beta > 0$ , then the expected number of positive samples in  $S$  is  $\mathbb{E}[\#\{s \in S \mid y_s = 1\}] \geq \frac{\alpha \cdot \gamma \cdot k}{2}$ , where  $\gamma = \beta/2$  if  $\beta < 1$  and  $\gamma = 1 - \frac{1}{2\beta}$  if  $\beta \geq 1$ .

The proof for the lemma can be found in Appendix A. The lemma does not directly translate to the setting with  $m > 2$  bins, but by altering the precision assumption with that of the model being approximately calibrated, we keep the guarantee of obtaining a constant fraction of positive examples<sup>6</sup>. In our ablation studies (see section 5.8) we see that the precise number of bins has a small effect on the result, resulting in our decision to use an arbitrary default value (of  $m = 10$ ) in our experiments.

#### 4.1 Extending to SoTA algorithms

The random algorithm combined with the InvProp distribution, as described above, enjoys being skewed towards the minority class and informative samples. Nevertheless, recent AL algorithms have been able to provide additional benefit over

<sup>6</sup>The precise assumption is that for a bin  $[\rho, \rho + 1/m]$ , at least  $\alpha\rho$  of the samples have a positive label. This property measures how calibrated the model is; indeed a perfectly calibrated model guarantee this with  $\alpha = 1$ . In this case we are guaranteed that in expectation,  $\sum_{i=0}^{m-1} \frac{i\alpha\gamma k}{2m^2} = \frac{(m-1)}{4m} \alpha\gamma k \geq \alpha\gamma k/6$  samples are positive.

the naive random approach. This leads us to incorporate our weighting scheme in SoTA AL algorithms. To this end, we consider the underlying distribution  $p$  of the InvProp method as a weighting function, assigning each point  $x$  a weight denoted by  $w(x)$  (in the section above this is the probability of choosing  $x$ ). For score based algorithms that choose samples based on a score  $s(x)$ , we modify the score to become  $s_w(x) = w(x) \cdot s(x)$  (see Algorithm 2). For embedding-based algorithms that choose samples based on clustering, we modify the clustering procedure to be weighted according to  $w(x)$  (see Algorithm 3). With this strategy, we are able to modify SoTA algorithms such as BADGE (Ash et al., 2020), ALPS (Yuan et al., 2020) and DAL (Gissin and Shalev-Shwartz, 2019) to be better suited to imbalanced datasets.

---

#### Algorithm 2 Weighted Score-based AL

---

**Input:** Unlabeled data pool  $\mathcal{U}$ , number of samples to select  $k$ , scoring function  $s$ , weight function  $w$ .

$$s_w(x) = \sum_{x \in S} w(x) \cdot s(x)$$

Return  $\operatorname{argmax}_{S \subseteq \mathcal{U}} \{s_w(x) \mid |S| = k\}$

---

---

#### Algorithm 3 Weighted Embedding-based AL

---

**Input:** Unlabeled data pool  $\mathcal{U}$ , number of samples to select  $k$ , embeddings  $e_x$ , weight function  $w$ .

Cluster  $\{e_x\}_{x \in \mathcal{U}}$  with weights  $w(x)$  into  $k$  clusters, pick  $S$  to be the samples nearest to the cluster centers.

Return  $S$

---

## 5 Experiments

### 5.1 Datasets & Measure

We follow recent AL work (Ein-Dor et al., 2020; Yuan et al., 2020; Wang et al., 2021) and use 4 popular text classification datasets. **DBpedia-14** (Zhang et al., 2015) contains text snippets from 14 different DBpedia (Lehmann et al., 2015) classes. Each snippet’s label is the ontology class it was taken from. **SST-2** (Socher et al., 2013) is a binary classification dataset containing sentences from movie reviews labeled by their sentiment. **PubMed-20K-RCT** (Dernoncourt and Lee, 2017) contains sentences from PubMed abstracts labeled by their role in the abstract. **AG\_news** (Zhang et al., 2015) contains news snippets classified into 4 categories.

**Binary-Imbalanced data.** We generate several Binary-Imbalanced versions from each of the above

mentioned datasets. This is done by first selecting a class from the original dataset to serve as the positive (minority) class in the generated dataset. Samples from all other classes are regarded as negative. Both train and test sets are updated this way. For DBpedia-14, due to its larger number of classes and small variance we observe on this dataset, we randomly select 3 of the classes to serve as the minority class. For the rest of the datasets we use each of the original classes as a minority class. After selecting a positive class and updating the train and test sets, we sample positive examples from the training set so that the fraction of positive samples in the resulting dataset will be equal to the desired value. We construct datasets where the fraction of positives is  $1/x$  for  $x \in \{2, 10, 20, 50, 100\}$ , i.e., 4 different imbalance ratios and a balanced dataset. See Table 1 for further details about the datasets.

**Measure.** As customary in studies related to imbalanced classes, we measure Balanced Accuracy and ROC-AUC, metrics that are insensitive to change in class distribution. Thus, the test set is not down-sampled, as this would only add noise to the evaluation. In addition, we compare the Positives Ratio, the fraction of minority samples found by each algorithm. This will allow us to evaluate the skew produced by our technique. Due to aggregations over different datasets, we found the standard deviations tend to be large, making confidence intervals uninformative. We thus measure statistical significance via p-values<sup>7</sup> and mark in **bold** results whose p-value is smaller than 5% for all compared results.

## 5.2 Active Learning Algorithms

We test our proposed weighting method by incorporating it into well established AL algorithms.

- **BADGE** (Ash et al., 2020) is an embedding-based algorithm. It generates a weak label according to the prediction of the current model. The weak label is used to calculate the gradient of the last layer of the model. The gradients are used as the embeddings.
- **ALPS**. (Yuan et al., 2020) Is an embedding-based algorithm. It generates an embedding according to the MLM objective. Some tokens are masked and predicted by the model. The embeddings are defined as the cross entropy distance between the prediction of the model

<sup>7</sup>The p-value is calculated using a relative t-test of the elements whose mean is calculated.

and the actual masked tokens.

- **DAL**. (Gissin and Shalev-Shwartz, 2019) Is a score-based algorithm. It trains a classification model to predict whether a sample is labelled or unlabelled. The “top” unlabelled samples (the samples with the highest confidence score) are selected.
- **Random**. The naive random baseline which selects samples uniformly at random.

## 5.3 Weighting Methods

Other than our main approach of **InvProp**, we evaluate two additional weighting methods.

- **Uniform** ignores the imbalance issue and assigns the same weight to all samples, effectively running the original AL algorithm.
- **PosProb** assigns a point  $x$  the weight  $M_{t-1}(x)$ , meaning the probability according to the available model that the sample is positive. This option is explored as a strawman that promotes positive samples, but is not adaptive to the imbalance ratio.

We test 3 weighting methods (Uniform, PosProb, InvProp), incorporated into the 4 mentioned AL algorithms. Overall, 12 procedures are tested.

Previous studies came to the conclusion that even though AL algorithms improve over naive random sampling, there is no single algorithm which performs best for all models and tasks (Lowell et al., 2019; Dor et al., 2020). Hence, we do not aim to show that one algorithm outperforms another, we aim to test whether utilizing our weighting method improves the performance of each of the base algorithms when applied on imbalanced datasets.

## 5.4 Experimental Setup

We evaluate each algorithm on all Binary-Imbalanced datasets through 10 runs, each with a different random seed that determines the shuffling and random initialization. The results of these runs were averaged to produce the reported results. We ran the algorithms for 5 AL iterations, selecting 50 samples at each iteration.

We use the uncased bert-base model with 110M parameters<sup>8</sup> (Devlin et al., 2019) as the PLM, and set the maximum sequence length to 128 and the number of train epochs to 2. To cope with the imbalance of the data that was already labeled, we use oversampling to ensure that the positive samples

<sup>8</sup><https://huggingface.co/transformers/>

Dataset	Classes	Original train/test	Selected minority classes
DBpedia-14	14	560K/70K	Company, Animal, Album
SST-2	2	~ 67K/872	all classes
Pubmed-20K-RCT	5	~ 180K/30K	all classes
AG_news	4	120K/7.6K	all classes

Table 1: **Dataset details.** Number of classes and train/test sizes in the original dataset, and the set of classes from the original dataset used as the minority class in the Binary-Imbalanced datasets we created (see section 5.1).

are at least a quarter of the training samples. This heuristic is known to vastly reduce the negative effects of the imbalance (Estabrooks and Japkowicz, 2001). At each iteration, we fine-tune the original pre-trained model using all examples, rather than fine-tuning the resulting model of a previous iteration, as this is known to provide better performance (Ash and Adams, 2020).

At the beginning of each run, we generate a random labeled subset of 5 negative and 5 positive samples that is given to the algorithm. An intuitive design would be that of a ‘cold start’, beginning with no labeled data. However, we found this setup to have a very high variance in terms of the final performance. In Section 5.9 we elaborate on the conjectured reason for this, provide experimental results to support our claim, and motivate our choice of starting with 5 examples of each class.

## 5.5 Main Results

In order to concisely compare between the weighting methods, we aggregate the results over either algorithms or datasets. Table 2 compares the performance of the InvProp and Uniform weighting method on the 4 datasets. A detailed comparison to PosProb can be found in Section 5.6. For each Binary-Imbalanced dataset and weighting method, we select the algorithm that performed best for the combination, based on its average score over 10 random seeds. We then take the average metric over Binary-Imbalanced variants of a dataset to obtain a single score for each (metric, dataset, weighting procedure) triplet. Here, we consider only Binary-Imbalanced datasets with an imbalance ratio of  $\frac{1}{50}$ . In both AUC and Balanced Accuracy, InvProp provides either superior or comparable results. As a justification, there is an overwhelming advantage to our weighting method in terms of positive ratio, nearly doubling the number of positive examples.

In Table 3 we compare the performance of the weighting methods between different AL algorithms. For each algorithm we average the score obtained by a specific weighting method over all

	Balanced Accuracy		AUC		Positives Ratio	
	InvProp	Uniform	InvProp	Uniform	InvProp	Uniform
AG_news	<b>0.7889</b>	0.7674	<b>0.9475</b>	0.9414	<b>0.3601</b>	0.2030
DBpedia	<b>0.9876</b>	0.9826	0.9983	0.9983	<b>0.4377</b>	0.2287
Pubmed	0.6338	0.6387	<b>0.8813</b>	0.8606	<b>0.2516</b>	0.1217
SST-2	0.6523	0.6299	<b>0.8709</b>	0.8538	<b>0.2045</b>	0.0989

Table 2: Comparison of InvProp with Uniform across multiple datasets (for an imbalance ratio of  $\frac{1}{50}$ ).

	Balanced Accuracy		AUC		Positives Ratio	
	InvProp	Uniform	InvProp	Uniform	InvProp	Uniform
ALPS	<b>0.7002</b>	0.6887	0.8712	0.8625	<b>0.0368</b>	0.0271
BADGE	0.7203	0.7181	0.8914	0.8901	<b>0.2730</b>	0.1646
DAL	<b>0.7036</b>	0.6692	<b>0.8948</b>	0.8560	<b>0.2823</b>	0.0142
Random	<b>0.7248</b>	0.6763	<b>0.8986</b>	0.8681	<b>0.2442</b>	0.0196

Table 3: Comparison of InvProp and Uniform across multiple base algorithms (for an imbalance ratio of  $\frac{1}{50}$ ).

Binary-Imbalanced datasets, averaged on 10 random seeds. Here we also consider an imbalance ratio of  $\frac{1}{50}$ . For both AUC and Balanced Accuracy we see a clear advantage to InvProp-weighting compared to Uniform weighting, in that it is either better or comparable for all algorithms. In terms of the Positives Ratio, the effect on the different algorithms varies, but we see a clear increase when using the InvProp weighting method. The method is able to produce more positives, and as shown in Appendix B is able to do this throughout the different iterations.

Two interesting insights that are not directly related to our study are (1) the algorithms differ in their robustness to imbalance, with BADGE proving to be quite robust, especially when Uniform weighting is applied; (2) the Random algorithm combined with our weighting method becomes a strong baseline and is in fact the leader w.r.t both Balanced Accuracy and AUC.

We ran a similar comparison against the PosProb baseline and concluded that for an imbalance ratio of  $\frac{1}{50}$ , the results are comparable. We provide the full experiments comparing all 3 weighting methods on multiple imbalance ratio combinations, on all datasets and base algorithms in Section 5.6. In

	Balanced Accuracy		AUC		Positives Ratio	
	InvProp	Uniform	InvProp	Uniform	InvProp	Uniform
2.0	0.9009	0.9029	0.9518	<b>0.9534</b>	0.4816	0.4855
10.0	<b>0.8253</b>	0.8149	<b>0.9348</b>	0.9306	<b>0.3259</b>	0.1436
20.0	<b>0.7709</b>	0.7519	<b>0.9175</b>	0.9061	<b>0.2779</b>	0.0946
50.0	<b>0.7122</b>	0.6881	<b>0.8890</b>	0.8692	<b>0.2091</b>	0.0564
100.0	<b>0.6839</b>	0.6612	<b>0.8552</b>	0.8454	<b>0.1692</b>	0.0375

Table 4: Comparison of InvProp with Uniform across multiple degrees of imbalance.

addition, a qualitative analysis of the informativeness of selected positive samples can be found in Appendix C.

## 5.6 Sensitivity to Imbalance Ratio

The main approach of techniques addressing dataset imbalance is to skew the sampling process towards positive samples. This leads to the question of when each technique should be applied; optimizing towards positive samples may have a negative effect on the results when the real-world distribution is almost balanced. Therefore, we evaluate the performance of InvProp-weighting on datasets with different imbalance ratios. As described in Section 5.1, we generated datasets with different imbalance ratios. Specifically, we use datasets with an imbalance ratio of 100, 50, 20, 10 and 2 (i.e. balanced). As in earlier experiments, we repeat the experiment on each Binary-Imbalanced dataset with 10 different random seeds. For each imbalance ratio we compare the average performance of each weighting method over all AL algorithms and datasets (the results per dataset and algorithm appear in Appendix D). Results of this experiment are reported in Table 4; it is easy to see that the InvProp-weighting method outperforms the Uniform baseline across all imbalance ratios which represent an imbalanced dataset. In addition, even for the balanced scenario, the baseline slightly outperforms our method in only one metric, AUC, while being comparable in the Balanced Accuracy metric and in Positives Ratio. This allows us to recommend the usage of InvProp-weighting even if the imbalance ratio is unknown.

As mentioned in Section 5.5, the PosProb weighting method performance is comparable to InvProp when considering a dataset with an imbalance ratio of  $\frac{1}{50}$ . In Table 5 we extend this comparison and compare the effect of the imbalance ratio on these two weighting methods. This shows that InvProp either outperforms or is comparable to PosProb in both Balanced Accuracy and AUC.

	Balanced Accuracy		AUC		Positives Ratio	
	InvProp	PosProb	InvProp	PosProb	InvProp	PosProb
2.0	<b>0.9009</b>	0.8313	<b>0.9518</b>	0.9442	0.4816	<b>0.6425</b>
10.0	<b>0.8253</b>	0.8155	<b>0.9348</b>	0.9310	0.3259	<b>0.3702</b>
20.0	0.7709	0.7744	<b>0.9175</b>	0.9140	0.2779	0.2821
50.0	0.7122	0.7128	<b>0.8890</b>	0.8833	<b>0.2091</b>	0.1831
100.0	0.6839	0.6800	0.8552	0.8513	<b>0.1692</b>	0.1273

Table 5: Comparison of InvProp with PosProb across multiple degrees of imbalance.

A deeper dive into the Balanced Accuracy scores shows that PosProb performs best (compared to InvProp) on datasets with an imbalance ratio of 20. Since the weights of PosProb are independent of the imbalance of the dataset, i.e., down-sampling the minority class does not affect the proportion between weights of two of the remaining samples, it is expected that the performance of PosProb will peak at some specific imbalance ratio and decline as the imbalance ratio changes further away from this peak.

Another important aspect of the results is the number of positive samples accumulated by each solution. For the balanced dataset, PosProb over-skews the selection toward the positive class, leading to more than 60% of positive selected samples. As the imbalance grows, this stabilizes, and for imbalance ratio of 50 and 100 InvProp is able to identify a larger number of positive samples. The higher AUC score of InvProp points toward PosProb finding less diverse positive samples due to the large weight it assigns to positive samples which the model is certain about.

## 5.7 Extreme imbalance

To further test the limits of our methods, we extend the imbalance ratios to 200 and 1000. Results are presented in Table 6. Interestingly, even though the InvProp and PosProb weighting schemes manage to produce more positive examples than the uniform scheme, the Balanced Accuracy and the AUC metrics are better for the uniform as the imbalance ratio reaches extremes such as 1000. This could be explained by the weighting methods choosing positive examples that are too similar which do not contribute to the efficacy of the model. In general, we believe that in such extreme cases, where positive examples are so scarce, the trained model is unable to learn and is under-performing, hindering the efficacy of any AL approach. The performance of all AL algorithms with the Uniform weighting scheme exemplifies this, where we see that random

	Balanced Accuracy			AUC			Positives Ratio		
	InvProp	PosProb	Uniform	InvProp	PosProb	Uniform	InvProp	PosProb	Uniform
200.0	<b>0.6619</b>	0.6543	0.6377	0.8244	0.8251	<b>0.8302</b>	<b>0.1263</b>	0.0886	0.0267
1000.0	0.6181	0.6154	0.6160	0.7821	0.7831	<b>0.8081</b>	<b>0.0535</b>	0.0304	0.0103

Table 6: Comparing the weighting schemes against extreme imbalance ratios.

sampling produces the highest AUC score all of the AL algorithms, with statistical significance (see Table 9, Appendix D).

### 5.8 Sensitivity to number of bins

One drawback of InvProp-weighting compared to Uniform and PosProb-weighting is the addition of a new parameter, the number of bins. We evaluate the effect of the number of bins on the performance of different AL algorithms (with InvProp-weights) on the Binary-Imbalanced variants of the AG\_news dataset. We consider InvProp with 2, 5, 10, 15 and 20 bins. Results appear in table 11 in Appendix D. In table 12 (Appendix D), we evaluate the sensitivity to the number of bins by measuring significance with a relative t-test. This table lists the p-value  $p(m)$  of a test comparing AUC and Balanced Accuracy for the selection of  $m$  bins (for  $m \in \{2, 5, 15, 20\}$ ) to that of 10, as 10 was our choice throughout the paper. The results demonstrate that our method is insensitive to the bin count hyper-parameter, as except for one outlier, all p-values fall in  $[0.05, 0.95]$ . This motivates our choice to avoid optimizing this parameter, thus fixing it to a single value throughout the paper. This experiment was performed for all imbalance ratios discussed in this paper (2, 10, 20, 50, 100). For brevity, we report only on an imbalance ratio of 100, since for this value the affect of the number of bins parameter was the largest. Since we conclude that this parameter is not significant even for this imbalance ratio, this conclusion is relevant across all imbalance ratios.

### 5.9 Warm start

Recall that in lieu of starting the learning process with no examples (cold-start), we start the process with 5 examples from each class. In Table 10 (Appendix) we show the size of the confidence intervals for the BADGE algorithm over the different datasets, showing an increase of at least  $1.5X$ , sometimes over  $10X$  when comparing the cold-start to the chosen warm-start scenario. A major cause of the cold-start variance is the time it takes the learner to achieve a handful of positive exam-

ples. Indeed, when the imbalance ratio is 50, a random selection is expected to obtain only 1 positive sample in each iteration, and in many runs the learner fails simply as it does not find a sufficient initial amount of positive examples. A full solution should discuss realistic methods for obtaining an initial seed of examples, but this challenge is outside the scope of this paper. In order to remove the noise in the evaluation process originating from the time it takes to obtain a handful of examples, we have all methods initialize with the mentioned warm-start. We chose to work with particularly 5 examples as this is a small enough number of samples for it to be easy to obtain even in the imbalance scenario, and on the other hand, the variance is significantly reduced compared to the ‘cold-start’ setting.

## 6 Conclusions

In this paper, we tackled the problem of AL in imbalanced datasets. We propose a novel weighting-technique, InvProp, and apply it to three recent AL algorithms (BADGE, APLS and DAL) as well as to naive random sampling. We show results for the PosProb weighting scheme in addition to InvProp and compare to a uniform weighting baseline on four datasets. We show that InvProp-weighting consistently finds more positive examples, and leads to better or comparable performance compared to other weighting schemes on all tested datasets and AL algorithms. We also test various imbalance settings showing that InvProp outperforms the Uniform baseline across all imbalance ratios which represent some imbalance. In addition, on a balanced dataset the Uniform baseline outperforms our method in only one metric, while being comparable in the rest. Compared to PosProb, our solution proved to be more robust to changes in the imbalance ratio. When testing extreme scenarios where the imbalance is less than 0.5%, all tested algorithms perform worse than the simple random sampling baseline, suggesting that AL algorithms are not advantageous given an uninformative model. Concluding, we show our novel weighting scheme improves several SoTA AL algorithms on various



datasets and imbalance ratios.

## References

- Jordan Ash and Ryan P Adams. 2020. On warm-starting neural network training. *Advances in Neural Information Processing Systems*, 33:3884–3894.
- Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. 2020. [Deep batch active learning by diverse, uncertain gradient lower bounds](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Ricardo Barata, Miguel Leite, Ricardo Pacheco, Marco OP Sampaio, João Tiago Ascensão, and Pedro Bizarro. 2021. Active learning for imbalanced data under cold start. In *Proceedings of the Second ACM International Conference on AI in Finance*, pages 1–9.
- William H Beluch, Tim Genewein, Andreas Nürnberger, and Jan M Köhler. 2018. The power of ensembles for active learning in image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9368–9377.
- Mausam C Lin. 2018. Active learning with unbalanced classes & example-generated queries. In *AAAI Conference on Human Computation*.
- Franck Dernoncourt and Ji Young Lee. 2017. [PubMed 200k RCT: a dataset for sequential sentence classification in medical abstracts](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing, IJCNLP 2017, Taipei, Taiwan, November 27 - December 1, 2017, Volume 2: Short Papers*, pages 308–313. Asian Federation of Natural Language Processing.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Liat Ein Dor, Alon Halfon, Ariel Gera, Eyal Shnarch, Lena Dankin, Leshem Choshen, Marina Danilevsky, Ranit Aharonov, Yoav Katz, and Noam Slonim. 2020. Active learning for bert: an empirical study. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7949–7962.
- Liat Ein-Dor, Alon Halfon, Ariel Gera, Eyal Shnarch, Lena Dankin, Leshem Choshen, Marina Danilevsky, Ranit Aharonov, Yoav Katz, and Noam Slonim. 2020. [Active learning for BERT: an empirical study](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 7949–7962. Association for Computational Linguistics.
- Andrew Estabrooks and Nathalie Japkowicz. 2001. A mixture-of-experts framework for learning from imbalanced data sets. In *International Symposium on Intelligent Data Analysis*, pages 34–43. Springer.
- Yifan Fu, Xingquan Zhu, and Bin Li. 2013. A survey on instance selection for active learning. *Knowledge and information systems*, 35(2):249–283.
- Daniel Gissin and Shai Shalev-Shwartz. 2019. [Discriminative active learning](#). *CoRR*, abs/1907.06347.
- Abbas Kazerouni, Qi Zhao, Jing Xie, Sandeep Tata, and Marc Najork. 2020. Active learning for skewed data sets. *arXiv preprint arXiv:2005.11442*.
- Gwangsu Kim and Chang D Yoo. 2022. Blending query strategy of active learning for imbalanced data. *IEEE Access*.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer. 2015. [Dbpedia - A large-scale, multilingual knowledge base extracted from wikipedia](#). *Semantic Web*, 6(2):167–195.
- David D Lewis and William A Gale. 1994. A sequential algorithm for training text classifiers. In *SIGIR'94*, pages 3–12. Springer.
- Yijing Li, Haixiang Guo, Qingpeng Zhang, Mingyun Gu, and Jianying Yang. 2018. [Imbalanced text sentiment classification using universal and domain-specific knowledge](#). *Knowl. Based Syst.*, 160:1–15.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- David Lowell, Zachary C Lipton, and Byron C Wallace. 2019. Practical obstacles to deploying active learning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 21–30.
- Cristian Padurariu and Mihaela Elena Breaban. 2019. [Dealing with data imbalance in text classification](#). In *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 23rd International Conference KES-2019, Budapest, Hungary, 4-6 September 2019*, volume 159 of *Procedia Computer Science*, pages 736–745. Elsevier.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke

- Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zihui Li, Brij B Gupta, Xiaojiang Chen, and Xin Wang. 2021. A survey of deep active learning. *ACM computing surveys (CSUR)*, 54(9):1–40.
- Sarang Shaikh, Sher Muhammad Daudpota, Ali Shariq Imran, and Zenun Kastrati. 2021. [Towards improved classification accuracy on highly imbalanced text dataset using deep neural language models](#). *Applied Sciences*, 11(2).
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIG-DAT, a Special Interest Group of the ACL*, pages 1631–1642. ACL.
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. [How to fine-tune BERT for text classification?](#) *CoRR*, abs/1905.05583.
- Shuohang Wang, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2021. [Want to reduce labeling cost? GPT-3 can help](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 4195–4205. Association for Computational Linguistics.
- Michelle Yuan, Hsuan-Tien Lin, and Jordan L. Boyd-Graber. 2020. [Cold-start active learning through self-supervised language modeling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 7935–7948. Association for Computational Linguistics.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 649–657.