

Sampling Bias in Deep Active Classification: An Empirical Study

Ameya Prabhu^{1*†}, Charles Dognin^{2*}, Maneesh Singh²

University of Oxford¹ Verisk | AI, Verisk Analytics²

ameya.prabhu@mailfence.com, {charles.dognin, maneesh.singh}@verisk.com

Abstract

The exploding cost and time needed for data labeling and model training are bottlenecks for training DNN models on large datasets. Identifying smaller representative data samples with strategies like active learning can help mitigate such bottlenecks. Previous works on active learning in NLP identify the problem of sampling bias in the samples acquired by uncertainty-based querying and develop costly approaches to address it. Using a large empirical study, we demonstrate that active set selection using the posterior entropy of deep models like FastText.zip (FTZ) is robust to sampling biases and to various algorithmic choices (query size and strategies) unlike that suggested by traditional literature. We also show that FTZ based query strategy produces sample sets similar to those from more sophisticated approaches (e.g ensemble networks). Finally, we show the effectiveness of the selected samples by creating tiny high-quality datasets, and utilizing them for fast and cheap training of large models. Based on the above, we propose a simple baseline for deep active text classification that outperforms the state-of-the-art. We expect the presented work to be useful and informative for dataset compression and for problems involving active, semi-supervised or online learning scenarios. Code and models are available at: <https://github.com/drimpossible/Sampling-Bias-Active-Learning>

1 Introduction

Deep neural networks (DNNs) trained on large datasets provide state-of-the-art results on various NLP problems (Devlin et al., 2019) including text classification (Howard and Ruder, 2018). However, the cost and time needed to get labeled data and to train models is a serious impediment to creating new and/or better models. This problem can be mitigated by creating smaller representative datasets with active learning which can be used for training DNNs to achieve similar test accuracy as

that using the full training dataset. In other words, the smaller sample can be considered a surrogate for the full data.

However, there is lack of clarity in the active learning literature regarding sampling bias in such surrogate datasets created using active learning (Settles, 2009): its dependence on models, functions and parameters used to acquire the sample. Indeed, what constitutes a good sample? In this paper, we perform an empirical investigation using active text classification as the application.

Early work in active text classification (Lewis and Gale, 1994) suggests that greedy query generation using label uncertainty may lead to efficient representative samples (Nonetheless, the same test accuracy). Subsequent concerns regarding sampling bias has led to explicit use of expensive diversity measures (Brinker, 2003; Hoi et al., 2006) in acquisition functions or using ensemble approaches (Liere and Tadepalli, 1997; McCallum and Nigam, 1998) to improve diversity implicitly.

Deep active learning approaches adapt the discussed framework above to train DNNs on large data. However, it is not clear if the properties of deep approaches mirror those of their shallow counterparts and if the theory and the empirical evidence regarding sampling efficiency and bias translates from shallow to deep models. For example, (Sener and Savarese, 2018) and (Ducoffe and Precioso, 2018) find that uncertainty based strategies perform no better than random sampling even if ensembles are used and using diversity measures outperform both. On the other hand, (Beluch et al., 2018; Gissin and Shalev-Shwartz, 2019) find that uncertainty measures computed with ensembles outperform diversity based approaches while (Gal et al., 2017; Beluch et al., 2018; Siddhant and Lipton, 2018) find them to outperform uncertainty measures computed using single models. A recent empirical study (Siddhant and Lipton, 2018) investigating active learning in NLP suggests that Bayesian active learning outperforms classical uncertainty sampling across all settings. However, the approaches have been limited to rel-

* indicates equal contribution

† Work done at Verisk | AI

actively small datasets.

1.1 Sampling Bias in Active Classification

In this paper, we investigate the issues of sampling bias and sample efficiency, the stability of the actively collected query and train sets and the impact of algorithmic factors - i.e. the setup chosen while training the algorithm, in the context of deep active text classification on large datasets. In particular, we consider two sampling biases: label and distributional bias, three algorithmic factors: initial set selection, query size and query strategy along with two trained models and four acquisition functions on eight large datasets.

To isolate and evaluate the impact of the above (combinatorial) factors, a large experimental study was necessary. Consequently, we conducted over 2.3K experiments on 8 popular, large, datasets of sizes ranging from 120K-3.6M. Note that the current trend in deep learning is to train large models on very large datasets. However, the aforementioned issues have not yet been investigated in the literature in such a setup. As shown in Table 1, the datasets used in latest such analysis on active text classification by (Siddhant and Lipton, 2018) are quite small in comparison. The datasets used by us are two orders of magnitude larger, our query samples often being the size of the entire datasets used, and the presented empirical study is more extensive (20x experiments).

Our findings are as follows:

(i) We find that utilizing the uncertainty query strategy using a deep model like FastText.zip (FTZ)¹ to actively construct a representative sample provides query and train sets with remarkably good sampling properties.

(ii) We find that a single deep model (FTZ) used for querying provides a sample set similar to more expensive approaches using ensemble of models. Additionally, the sample set has a large overlap with support vectors of an SVM trained on the entire dataset largely invariant to a variety of algorithmic factors, thus indicating the robustness of the acquired sample set.

(iii) We demonstrate that the actively acquired training datasets can be utilized as small, surrogate training sets with a 5x-40x compression for training large, deep text classification models. In

¹We use FastText.zip (FTZ) to optimize the time and resources needed for this study.

particular, we can train the ULMFiT (Howard and Ruder, 2018) model to state of the art accuracy at 25x-200x speedups.

(iv) Finally, we create a novel, state-of-the-art baseline for active text classification which outperforms recent work (Siddhant and Lipton, 2018), using Bayesian dropout, utilizing 4x less training data. We also outperform (Sener and Savarese, 2018) at all training data sizes. The latter uses an expensive diversity based query strategy (coreset sampling).

The rest of the paper is organized as follows: in Section 2, the experimental methodology and setup are described. Section 3 presents the experimental study on sampling biases as well as the impact of various algorithmic factors. In Section 4, we compare with prior literature in active text classification. Section 5 presents a downstream use case - fast bootstrapping of the training of very large models like ULMFiT. Finally, we discuss the current literature in light of our work in Section 6 and summarize the conclusions in Section 7.

2 Methodology

This section describes the experimental approach and the setup used to empirically investigate the issues of (i) sampling bias and (ii) sampling efficiency in creating small samples to train deep models.

2.1 Approach

A labelled training set is incrementally built from a pool of unlabeled data by selecting & acquiring labels from an oracle in sequential increments. In this, we follow the standard approach found in the active learning literature. We use the following terminology:

Queries & Query Strategy: We refer to the (incremental) set of points selected to be labeled and added to the training as the query and the (acquisition) function used to select the samples as the query strategy.

Pool & Train Sets: The pool is the unlabeled data from which queries are iteratively selected, labeled and added to the (labeled) train set.

Let $\mathcal{D}_S = (\mathbf{x}_i, y_i)$ denote a dataset consisting of $|\mathcal{S}| = n$ i.i.d samples of data/label pairs, where $|\cdot|$ denotes the cardinality. Let $\mathcal{S}_0 \subset \mathcal{S}$ denote an initial randomly drawn sample from the initial pool. At each iteration, we train the model on the current train set and use a model-dependent query strategy

to acquire new samples from the pool, get them labeled by an oracle and add them to the train set. Thus, a sequence of training sets: $[\mathcal{S}_1, \mathcal{S}_2 \dots, \mathcal{S}_b]$ is created by sampling b queries from the pool set, each of size K . The b queries are given by $[\mathcal{S}_1 - \mathcal{S}_0, \mathcal{S}_2 - \mathcal{S}_1 \dots, \mathcal{S}_b - \mathcal{S}_{b-1}]$. Note that $|\mathcal{S}_i| = (|\mathcal{S}_0| + i \times K)$ and $\mathcal{S}_1 \subset \mathcal{S}_2 \dots \subset \mathcal{S}_b \subset \mathcal{S}$.

In this paper, we investigate the efficiency and bias of sample sets $\mathcal{S}_b^1, \mathcal{S}_b^2, \dots, \mathcal{S}_b^t$ obtained by different query strategies Q^1, Q^2, \dots, Q^t . We exclude the randomly acquired initial set and perform comparisons on the actively acquired sample sets defined as $\hat{\mathcal{S}}_j^i = (\mathcal{S}_j^i - \mathcal{S}_0^i)$.

2.2 Experimental Setup

In this section, we share details of the experimental setup, and present and explain the choice of the datasets, models and query strategies used.

Datasets: We used eight, large, representative datasets widely used for text classification: AG-News (AGN), DBPedia (DBP), Amazon Review Polarity (AMZP), Amazon Review Full (AMZF), Yelp Review Polarity (YRP), Yelp Review Full (YRF), Yahoo Answers (YHA) and Sogou News (SGN). Please refer to Section 4 of (Zhang et al., 2015) for details regarding the collection and characteristics of these datasets. Table 1 provides a comparison regarding the choice of datasets, models and number of experiments between our study and (Siddhant and Lipton, 2018) which investigates a variety of NLP tasks including text classification while we focus only on the latter.

Models: We reported two text classification models as representatives of classical and deep learning approaches respectively which were fast to train and also had good performance on text classification: Multinomial Naive Bayes (MNB) with TF-IDF (Wang and Manning, 2012) and FastText.zip (FTZ) (Joulin et al., 2016). The FTZ model provides results competitive with VDC-NNs (a 29 layer CNN) (Conneau et al., 2017) but with over $15,000\times$ speedup (Joulin et al., 2017). This allowed us to conduct a thorough empirical study on large datasets. Multinomial Naive Bayes (MNB) with TF-IDF features is a popularly claimed baseline for text classification (Wang and Manning, 2012).

Query Strategies: Uncertainty based query strategies are widely used and well studied in the active learning literature. Those strategies typically use a scoring function on the (softmax) out-

Paper	#Exp	Datasets (#Train)	Models (Full Acc)
DAL	120	TREC-QA (6k), MAREview (10.5k)	SVM (89%), CNN (91%), LSTM (92%)
Ours	2.3K	AGN (120k), SGN (450k), DBP (560k), YRP (560k), YRF (650k), YHA (1400k), AMZP (3600k), AMZF (3000k)	FTZ (97%), MNB (90%)

Table 1: Comparison of active text classification datasets and models (Acc on Trec-QA) used in (Siddhant and Lipton, 2018) and our work. We use significantly larger datasets (two orders larger), perform 20x more experiments, and use more efficient and accurate models.

put of a single model. We evaluate the following ones: Least Confidence (LC) and Entropy (Ent). Independently training ensembles of models (Lakshminarayanan et al., 2017) is another principled approach to obtain uncertainties associated with the output estimate. Then, we tried four query strategies - LC and Ent computed using single and ensemble models and evaluated them against random sampling (chance) as a baseline. For ensembles, we used five FTZ ensembles (Lakshminarayanan et al., 2017). In contrast, (Siddhant and Lipton, 2018) used Bayesian ensembles using Dropout, proposed in (Gal et al., 2017). Please refer to Section 4 for a comparison.

Implementation Details: We performed 2304 active learning experiments. We obtained our results on three random initial sets and three runs per seed (to account for stochasticity in FTZ) for each of the eight datasets. The query sizes were 0.5% of the dataset for AGN, AMZF, YRF and YHA and 0.25% for SGN, DBP, YRP and AMZP respectively for $b = 39$ sequential, active queries. We also experimented with different query sizes keeping the size of the final training data $b \times K$ constant. The default query strategy uses a single model with output Entropy (Ent) unless explicitly stated otherwise. Results in the chance column are obtained using random query strategy.

We used Scikit-Learn (Pedregosa et al., 2011) implementation for MNB and original implementation for FastText.zip (FTZ)². We required 3 weeks of running time for all FTZ experiments on a x1.16xlarge AWS instance with Intel Xeon E7-8880 v3 processors and 1TB RAM to obtain results presented in this work. The experiments are deterministic beyond the stochasticity involved in training the FTZ model, random initialization and

²<https://github.com/facebookresearch/fastText>

Dsets	Limit	FTZ ($\cap Q$)	MNB ($\cap Q$)	FTZ ($\cap S$)	MNB ($\cap S$)
SGN	1.61	1.56 \pm 0.03	1.15 \pm 0.32	1.59 \pm 0.01	1.57 \pm 0.01
DBP	2.64	2.50 \pm 0.02	2.27 \pm 0.11	2.51 \pm 0.0	2.58 \pm 0.01
YHA	2.30	2.25 \pm 0.01	2.22 \pm 0.02	2.25 \pm 0.0	2.28 \pm 0.0
YRP	0.69	0.69 \pm 0.0	0.56 \pm 0.13	0.69 \pm 0.0	0.69 \pm 0.01
YRF	1.61	1.56 \pm 0.02	1.42 \pm 0.21	1.56 \pm 0.0	1.57 \pm 0.01
AGN	1.39	1.33 \pm 0.04	1.13 \pm 0.17	1.33 \pm 0.0	1.35 \pm 0.01
AMZP	0.69	0.69 \pm 0.0	0.69 \pm 0.0	0.69 \pm 0.0	0.69 \pm 0.0
AMZF	1.61	1.58 \pm 0.02	1.6 \pm 0.01	1.59 \pm 0.0	1.61 \pm 0.0

Table 2: Label entropy with a large query size ($b = 9$ queries). $\cap Q$ denotes averaging across queries of a single run, $\cap S$ denotes the label entropy of the final collected samples, averaged across seeds. Naive Bayes ($\cap Q$) has biased (inefficient) queries while FastText ($\cap Q$) shows stable, high label entropy showing a rich diversity in classes despite the large query size. Overall, the resultant sample ($\cap S$) becomes balanced in both cases.

SGD updates. The entire list of hyperparameters and metrics affecting uncertainty such as calibration error (Guo et al., 2017) is given in the supplementary material. The experimental logs and models are available on our github link³.

3 Results

In this section, we study several aspects of sampling bias (class bias, feature bias) and the impact of relevant algorithmic factors (initial set selection, query size and query strategy).

We evaluated the actively acquired queries and sample set for sampling bias, and for the stability as measured by %intersection of collected sets across a critical influencing factor. Higher sample intersections indicate more stability increase to the chosen influencing factor.

3.1 Aspects of Sampling Bias

We study two types of sampling biases: (a) Class Bias and (b) Feature Bias.

3.1.1 Class Bias

Greedy uncertainty based query strategies are said to pick disproportionately from a subset of classes per query (Sener and Savarese, 2018; Ebert et al., 2012), developing a lopsided representation in each query. However, its effect on the resulting sample set is not clear. We test this by measuring the Kullback-Leibler (KL) divergence between the ground-truth label distribution and the distribution obtained per query as one experiment ($\cap Q$),

³<https://github.com/drimpossible/Sampling-Bias-Active-Learning>

and over the resulting sample ($\cap S$) as the second. Let us denote P as the true distribution of labels, \hat{P} the sample distribution and C the total number of classes. Since P follows a uniform distribution, we can use Label entropy instead ($L = -KL(P||\hat{P}) + \log(C)$). Label entropy L is an intuitive measure. The maximum label entropy is reached when sampling is uniform, $\hat{P}(x) = P(x)$, i.e. $L = \log(C)$.

We present our results in Table 2. We observe that across queries ($\cap Q$), FTZ with entropy strategy has a balanced representation from all classes (high mean) with a high probability (low std) while Multinomial Naive Bayes (MNB) results in more biased queries (lower mean) with high probability (high std) as studied previously. However, we did not find evidence of class bias in the resulting sample ($\cap S$) in both models: FastText and Naive Bayes (column 5 and 6 from Table 2).

We conclude that entropy as a query strategy can be robust to class bias even with large query sizes.

3.1.2 Feature Bias

Uncertainty sampling can lead to undesirable sampling bias in feature space (Settles, 2009) by repeating redundant samples and picking outliers (Zhu et al., 2008). Diversity-based query strategies (Sener and Savarese, 2018) are used to address this issue, by selecting a representative subset of the data. In the context of active classification, it is good to pick the most informative samples to be the ones closer to class boundaries⁴.

Indeed, recent work suggests that the learning in deep classification networks may focus on small part of the data closer to class boundaries, thus resembling support vectors (Xu et al., 2018; Toneva et al., 2019). To investigate whether uncertainty sampling also exhibits this behavior, we perform below a direct comparison with support vectors from a SVM. For this, we train a FTZ model on the full training data and train a SVM on the resulting features (sentence embeddings) to obtain the support vectors and compute the intersection of support vectors with each selected set. The percentage intersections are shown in Table 3. The high percentage overlap is a surprising result which shows that the sampling is indeed biased but in

⁴In this work, we assume ergodicity in the setup. We do not consider incremental, online modeling scenarios where new modes or new classes are sequentially encountered.

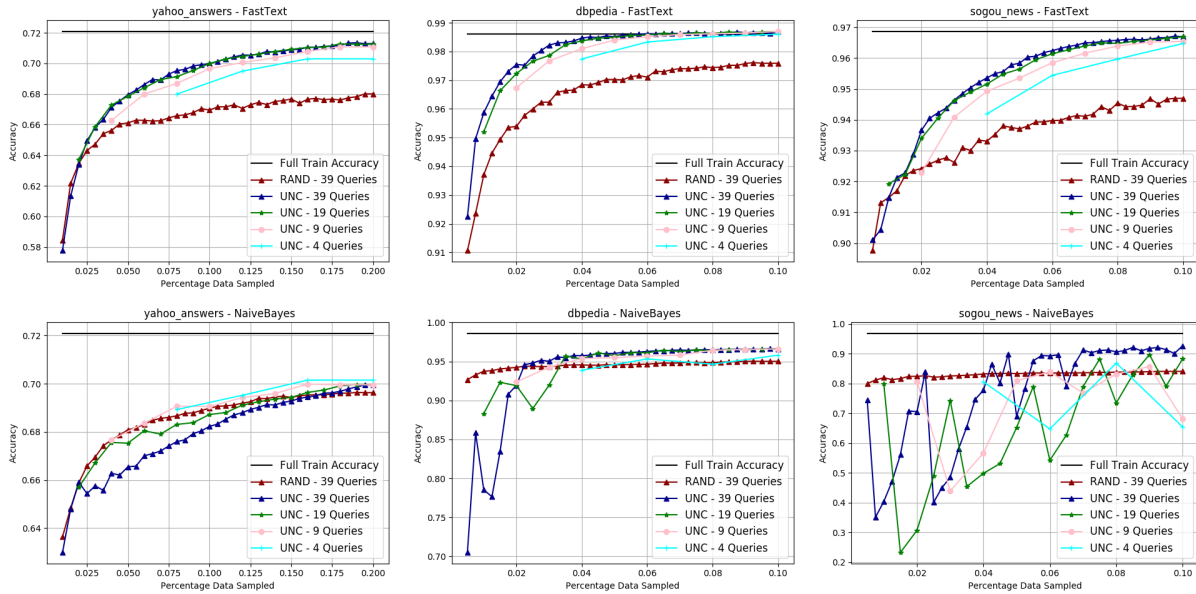


Figure 1: Accuracy across different number of queries b for FastText and Naive Bayes, with $b \times K$ constant. FastText is robust to increase in query size and significantly outperforms random in all cases. Naive Bayes: (Left) All including $b=39$ perform worse than random, (Center) All including $b=9$ eventually perform better than random (Right) $b = 39$ performs better than random but larger query sizes perform worse than random. Uncertainty sampling with Naive Bayes suffers from sampling size bias.

Dsets	Common%	Chance%	#SV
SGN	71.3 ± 0.5	9.3 ± 0.5	13184
DBP	86.3 ± 0.5	9.7 ± 0.5	1479
YRP	57.3 ± 0.5	9.7 ± 0.5	31750
AGN	45.0 ± 0.8	21.0 ± 1.6	1032

Table 3: Proportion of Support Vectors intersecting with our actively selected set calculated by $\frac{|S_{SV} \cap S_b|}{|S_{SV}|}$. Actively selected sets share large overlap with supports of an SVM (critical for classification).

a desirable way. Since the support vectors represent the class boundaries, a large percentage of selected data consists of samples around the class boundaries. This overlap indicates that the actively acquired training sample covers the support vectors well which are important for good classification performance. The overlap with the support vectors of an SVM (a fixed algorithm) also suggests that uncertainty sampling using deep models might generalize beyond FastText, to other learning algorithms.

Experimental Details: We used a fast GPU implementation for training an SVM with a linear kernel (Wen et al., 2018) with default hyperparameters. Please refer to supplementary material for additional details. We ensured the SVM achieves similar accuracies as original FTZ model.

Dsets	Chance	FTZD	FTZS	MNBD	MNBS
SGN	0.8	77.8	81.0	55.5	100.0
DBP	0.9	79.7	81.3	79.7	100.0
YHA	3.7	69.0	73.6	89.5	100.0
YRP	0.9	42.9	43.7	16.0	100.0
YRF	3.6	67.7	71.6	13.6	100.0
AGN	3.7	68.7	70.1	79.8	100.0
AMZP	0.9	48.4	48.8	15.0	100.0
AMZF	3.6	56.8	63.1	57.8	100.0

Table 4: % Intersection of samples obtained with different seeds (ModelD) compared to same seeds (ModelS) and chance intersection for $b = 39$ queries. We see that FastText is initialization independent (FTZD \approx FTZS \gg Chance). NaiveBayes shows significant dependency on the initial set sometimes, while other times performs comparable to FastText.

3.2 Algorithmic Factors

We analyze three algorithmic factors of relevance to sampling bias: (a) Initial set selection (b) Query size, and, (c) Query strategy.

3.2.1 Initial Set Selection

To investigate the dependence of the actively acquired train set on the initial set, we compare the overlap (intersection) of the incrementally constructed sets from different random initial sets versus the same initial set. The results are shown in Table 4. We first observe that chance overlaps (column 2) are very low - less than 4%. Columns

Dsets	Chance	FTZ 9 \cap 19 \cap 39	FTZ 39 \cap 39 \cap 39	MNB 9 \cap 19 \cap 39	MNB 39 \cap 39 \cap 39
SGN	0.83 \pm 0.0	77.0 \pm 0.5	77.9 \pm 0.2	31.9 \pm 0.0	55.5 \pm 0.0
DBP	0.9 \pm 0.0	80.0 \pm 0.1	79.6 \pm 0.2	82.3 \pm 0.0	79.7 \pm 0.0
YHA	3.7 \pm 0.0	68.3 \pm 0.1	69.0 \pm 0.0	92.1 \pm 0.0	89.5 \pm 0.0
YRP	0.9 \pm 0.0	46.0 \pm 0.9	42.7 \pm 1.0	10.8 \pm 0.0	16.0 \pm 0.0
YRF	3.6 \pm 0.0	68.4 \pm 0.2	67.6 \pm 0.1	14.2 \pm 0.0	13.6 \pm 0.0
AGN	3.7 \pm 0.0	70.3 \pm 0.2	68.7 \pm 0.1	81.6 \pm 0.0	79.8 \pm 0.0
AMZP	0.9 \pm 0.0	45.8 \pm 0.1	48.2 \pm 0.2	11.5 \pm 0.0	15.0 \pm 0.0
AMZF	3.6 \pm 0.0	55.2 \pm 0.4	57.0 \pm 0.2	28.4 \pm 0.0	57.8 \pm 0.0

Table 5: Intersection of samples obtained with different values of b . We see the intersection of samples selected with different number of intersections comparable to highest possible (different seeds) in FastText, far higher compared to chance intersection. This indicates similar samples are selected regardless of sample size. NaiveBayes does not show clear trends but occasionally the queried percentage drops significantly when increasing iterations, occasionally it remains unaffected.

Dsets	Chance	FTZ Ent-Ent	FTZ Ent-LC	FTZ Ent-DelEnt	FTZ DelEnt-DelLC	FTZ DelEnt-DelEnt
SGN	9.4 \pm 0.0	84.6 \pm 0.2	83.1 \pm 0.3	81.7 \pm 0.1	82.6 \pm 0.1	84.2 \pm 0.1
DBP	9.3 \pm 0.0	85.7 \pm 0.2	85.5 \pm 0.3	83.3 \pm 0.1	83.0 \pm 0.4	83.2 \pm 0.2
YHA	19.0 \pm 0.0	79.0 \pm 0.0	71.6 \pm 0.2	76.3 \pm 0.1	69.6 \pm 0.7	75.6 \pm 3.9
YRP	9.3 \pm 0.0	58.4 \pm 0.6	59.0 \pm 0.3	59.0 \pm 0.6	61.6 \pm 0.7	62.1 \pm 0.1
YRF	19.0 \pm 0.0	77.8 \pm 0.2	66.6 \pm 0.3	75.8 \pm 0.1	65.4 \pm 0.3	80.1 \pm 0.2
AGN	19.1 \pm 0.0	78.3 \pm 0.1	77.3 \pm 0.1	77.1 \pm 0.3	78.2 \pm 0.4	79.0 \pm 0.3
AMZP	9.5 \pm 0.0	63.5 \pm 0.2	63.5 \pm 0.3	66.1 \pm 0.4	70.0 \pm 0.1	70.0 \pm 0.1
AMZF	19.0 \pm 0.0	70.3 \pm 0.1	64.3 \pm 0.2	69.6 \pm 0.1	65.6 \pm 0.2	72.6 \pm 0.2

Table 6: Intersection of query strategies across acquisition functions. We observe that the % intersection among samples in the Ent-LC is comparable to those Ent-Ent. Similarly, the Ent-DelEnt (entropy with deletion) is comparable to both DelEnt-DelLC and DelEnt-DelEnt showing robustness of FastText to query functions (beyond minor variation). DelEnt-DelEnt obtains similar intersections as compared to Ent-Ent, showing the robustness of the acquired samples to deletion.

3 and 5 present overlaps from different initial sets while 4 and 6 from same initial sets. We note from column 4 and 6 that due to the stochasticity of training in FTZ, we expect non-identical final sets even with same initial samples as well. The results demonstrate that samples obtained using FastText are largely initialization independent (low variation between columns 3 and 4) consistently across datasets while the samples obtained with Naive Bayes can be vastly different showing relatively heavy dependence on the initial seed. This indicates the relative stability of train set obtained with the posterior uncertainty of the actively trained FTZ as an acquisition function.

3.2.2 Query size

Since the sampled data is sequentially constructed by training models on previously sampled data, large query sizes were expected to impact samples collected by uncertainty sampling and the performance thereof (Hoi et al., 2006). We experiment with various query sizes - (0.25%, 0.5%, 1%) for DBP, SGN, YRP and AMZP and (0.5%, 1%, 2%) for the rest corresponding to 9, 19 and 39 iterations. Figure 1 shows that FastText (top

row) has very stable performance across sample sizes while MNB (bottom row) show more erratic performance. Table 5 presents the intersection of samples obtained with different query sizes across multiple runs. We observe a high overlap of the acquired samples across different query sizes indicating that the performance is independent of the query size (compare column 3 to column 4 where the size is held constant) while MNB results in lower overlap with more erratic behavior due to change in the query size (compare column 5 compared to column 6).

3.2.3 Query strategy

We now investigate the impact of various query strategies using FastText by evaluating and comparing the correlation between the respective actively selected sample sets.

Acquisition Functions: We compare four uncertainty query strategies: Least Confidence (LC) and Entropy (Ent), with and without deletion of least uncertain samples from the training set. Deletion of least uncertain samples reduces the dependence on the initial randomly selected set. The results are presented in Table 6. We present five of

Dsets	Chance	FTZ-FTZ Ent	FTZ-5F TZ Ent	5FTZ-5FTZ Ent-LC	5FTZ-5FTZ Ent-Ent
SGN	9.4 ± 0.0	84.6 ± 0.2	86.3 ± 0.2	85.4 ± 0.4	85.8 ± 0.0
DBP	9.3 ± 0.0	85.7 ± 0.2	86.6 ± 0.3	86.78 ± 0.1	87.8 ± 0.2
YRP	9.3 ± 0.0	58.4 ± 0.6	58.1 ± 0.7	58.3 ± 0.3	58.2 ± 0.2
YRF	19.0 ± 0.0	77.8 ± 0.2	79.0 ± 0.3	68.5 ± 1.1	77.6 ± 0.3
AGN	19.1 ± 0.0	78.3 ± 0.1	79.0 ± 0.2	79.1 ± 0.2	77.9 ± 0.2

Table 7: Intersection of query strategies across single and ensemble of 5FTZ models. We observe that the % intersection of samples selected by ensembles and single models is comparable to intersection among either. The 5-model committee does not seem to add any additional value over selection by a single model.

the ten possible combinations and again observe the high degree of overlap in the collected samples. It can be concluded that the approach is fairly robust to these variations in the query strategy.

Ensembles versus Single Models: A similar experiment was conducted to investigate the overlap between a single FTZ model and a probabilistic committee of models (5-model ensemble with FTZ (Lakshminarayanan et al., 2017)) to identify comparative advantages of using ensemble methods. The results are presented in Table 7 showing little to no difference in sample overlaps.⁵ We conclude that more expensive sampling strategies commonly used, like ensembling, may offer little benefit compared to using a single FTZ model with posterior uncertainty as a query function.

The experiments in this section demonstrate that uncertainty based sampling using deep models like FTZ show no class bias or an undesirable feature bias (and favorable bias to class boundaries). There is also a high degree of robustness to algorithmic factors, especially query size, a surprisingly high degree of overlap in the resulting training samples and stable performances (classification accuracy). Additionally, all uncertainty query strategies perform well, and expensive sampling strategies like ensembling offer little benefit. We conclude that sampling biases demonstrated in active learning literature do hold well with traditional models, however, they do not seem to translate to deep models like FTZ using (posterior) uncertainty.

4 Application: Active Text Classification

Experimental results from the previous sections suggest that entropy function with a single FTZ

⁵The ensembles were too costly to run on larger datasets, so the results for YHA, AMZP and AMZF could not be obtained.

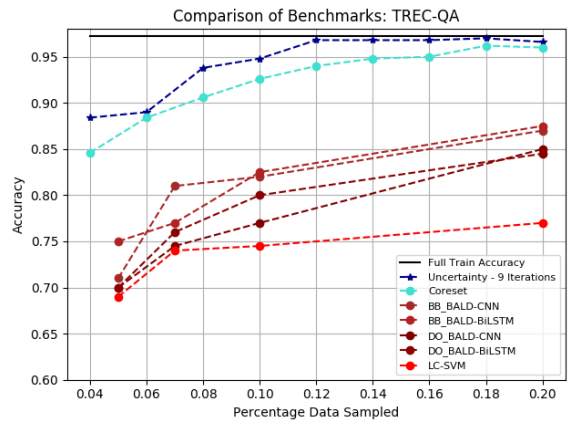


Figure 2: Active text classification: Comparison with K-Center Coreset, BALD and SVM algorithms. Accuracy is plotted against percentage data sampled. We reach full-train accuracy using 12% of the data, compared to BALD which requires 50% data and perform significantly worse in terms of accuracy. We also outperform K-center greedy Coreset at all sampling percentages without utilizing additional diversity-based augmentation.

model would be a good baseline for active text classification. We compare our baseline with the latest work in deep active learning for text classification - BALD (Siddhant and Lipton, 2018) and with the recent diversity based Coreset query function (Sener and Savarese, 2018) which uses a costly K-center algorithm to build the query. Experiments are performed on TREC-QA for a fair comparison (used by (Siddhant and Lipton, 2018)). Table 8 shows that the results of our study generalize to small datasets like TREC-QA.

The results are shown in Figure 2 using the baseline with the query size of 2% of the full dataset (b=9 queries). Note that uncertainty sampling converges to full accuracy using just 12% of the data, whereas (Siddhant and Lipton, 2018) required 50% of the data. There is also a remarkable accuracy improvement over (Siddhant and Lipton, 2018) which can be largely attributed to the models used (FastText versus 1-layer CNN/BiLSTM). Also, uncertainty sampling outperforms diversity-based augmentations like Coreset Sampling (Sener and Savarese, 2018) before convergence. Thus, we establish a new state-of-the-art baseline for further research in deep active text classification.

5 Application: Training of Large Models

The cost and time needed to get and label vast amounts of data to train large DNNs is a serious

Dsets	Chance	FTZ-Ent-Ent	FTZ Ent-LC	SV Chce%	SV Com%
TQA	15.1 ± 0.0	59.7 ± 0.5	56.3 ± 1.4	18.7 ± 6.1	79.0 ± 3.6

Table 8: Results of sample selection from previous investigations on small datasets (Trec-QA).

Model	AGN	DBP	SGN	YRF	YRP	YHA	AMZP	AMZF
VDCNN (Conneau et al., 2017)	91.3	98.7	96.8	64.7	95.7	73.4	95.7	63.0
DPCNN (Johnson and Zhang, 2017)	93.1	99.1	98.1	69.4	97.3	76.1	96.7	65.2
WC-Reg (Qiao et al., 2018)	92.8	98.9	97.6	64.9	96.4	73.7	95.1	60.9
DC+MFA (Wang et al., 2018)	93.6	99.2	-	66.0	96.5	-	-	63.0
DRNN (Wang, 2018)	94.5	99.2	-	69.1	97.3	70.3	96.5	64.4
ULMFiT (Howard and Ruder, 2018)	95.0	99.2	-	70.0	97.8	-	-	-
EXAM (Du et al., 2019)	93.0	99.0	-	-	-	74.8	95.5	61.9
Ours: ULMFiT (Small data)	93.7 (20)	99.2 (10)	97.0 (10)	67.6 (20)	97.1 (10)	74.3 (20)	96.1 (10)	64.1 (20)
Ours: ULMFiT (Tiny data)	91.7 (8)	98.6 (2.3)	97.4 (6.3)	66.3 (8)	96.7 (4)	73.3 (8)	95.8 (4)	62.9 (8)

Table 9: Comparison of accuracies with state-of-the-art approaches (earliest-latest) for text classification (%dataset in brackets). We are competitive with state-of-the-art models while using 5x-40x compressed datasets.

ULMFiT	AGN	DBP	YRP	YRF
Full	95.0	99.2	97.8	70.0
Ours-Small	93.7 (20)	99.2 (10)	97.1 (10)	67.6 (20)
Ours-Tiny	91.7 (8)	98.6 (2.3)	96.7 (4)	66.3(8)

Table 10: ULMFiT: Resulting sample \hat{S}_b compared to reported accuracies in (Howard and Ruder, 2018) (%dataset in brackets). We observe that using our cheaply obtained compressed datasets, we can achieve similar accuracies with 25x-200x speedup (5x less epochs, 5x-40x less data). Transferability to other models is evidence of the generalizability of the subset collected using FTZ to other deep models.

impediment to creating new and/or better models. Our study suggests that the training samples collected with uncertainty sampling (entropy) on a single model FTZ may provide a good representation (surrogate) for the entire dataset. Buoyed by this, we investigate if we can speedup training of ULMFiT (Howard and Ruder, 2018) using the surrogate dataset. We show these results in Table 10. We achieve 25x-200x speedup⁶ (5x fewer epochs, 5x-40x smaller training size). We also benchmark the performance against the state-of-the-art on text classification as shown in Table 9. We conclude that we can significantly compress the training datasets and speedup classifier training time with little tradeoff in accuracy.

Implementation Details: We use the official github repository for ULMFiT⁷, use default hyperparameters and train on one NVIDIA Tesla V100 16GB GPU. Further details are provided in sup-

⁶The cost of acquiring the training data using FTZ-Ent is negligible in comparison.

⁷https://github.com/fastai/fastai/tree/master/courses/dl2/imb_scripts

plementary material.

6 Related Work

We now expand on the brief literature review in Section 1 to better contextualize our work. We divide the past works into (i) Traditional Models and (ii) Deep Models.

Sampling Bias in Classical AL in NLP: Active learning (AL) in text classification started with greedy uncertainty query strategy from a pool using decision trees (Lewis and Gale, 1994), which was shown to be effective and led to widespread adoption with classifiers like SVMs (Tong and Koller, 2001), Naive Bayes (Roy and McCallum, 2001) and KNN (Fujii et al., 1998). This strategy was also applied to other NLP tasks like parse selection (Baldrige and Osborne, 2004), sequence labeling (Settles and Craven, 2008) and information extraction (Thompson and Mooney, 1999). These early papers popularized two greedy uncertainty query methods: Least Confident and Entropy.

Issues of lack of diversity (large redundancy in sampling) (Zhang and Oles, 2000) and lack of robustness (high variance in sample quality)(Krogh and Vedelsby, 1994) guided subsequent efforts. The two most popular directions were: (i) augmenting uncertainty with diversity measures (Hoi et al., 2006; Brinker, 2003; Tang et al., 2002) and (ii) using query-by-committee (McCallum and Nigam, 1998; Liere and Tadepalli, 1997). For a comprehensive survey of classical AL methods for NLP, please refer to (Settles, 2009).

Sampling Bias in Deep AL: Deep active learning approach adapt the above framework to the

training of DNNs on large data. Two main query strategies are used: (i) ensemble based greedy uncertainty, which represents a probabilistic query-by-committee paradigm (Gal et al., 2017; Beluch et al., 2018), and (ii) diversity based measures (Sener and Savarese, 2018; Ducoffe and Precioso, 2018). Papers proposing diversity based approaches find that greedy uncertainty based sampling (using ensemble and single model) perform significantly worse than random (See Figures 4 and 2 respectively in (Sener and Savarese, 2018; Ducoffe and Precioso, 2018)). They attribute the poor performance to redundant, highly correlated sampling selected using uncertainty based methods and justify the need for prohibitively expensive diversity-based approaches (Refer section 2 of (Sener and Savarese, 2018) for details on the expensiveness of various diversity sampling methods). However, K-center greedy coresets sampling scales poorly: we were only able to use it on TREC-QA (a small dataset). On the other hand, ensemble-based greedy uncertainty methods find that probabilistic averaging from a committee (Gal et al., 2017; Beluch et al., 2018) performs better than single model as with on diversity based methods like coresets (Gissin and Shalev-Shwartz, 2019; Beluch et al., 2018). Current approaches in text classification literature mostly adopt the ensemble based greedy uncertainty framework (Siddhant and Lipton, 2018; Lowell et al., 2018; Zhang et al., 2017).

However, our work demonstrates the problems of sampling bias and efficiency may not translate from shallow to deep approaches. Recent evidence from image domain (Gissin and Shalev-Shwartz, 2019) demonstrates atleast a subset of our findings generalize to other DNNs (class bias and query functions). Uncertainty sampling using a deep model like FTZ demonstrates surprisingly good sampling properties without using ensembles or bayesian methods. Ensembles do not seem to significantly affect sampling. Whether this behavior generalizes to other deep models and tasks is yet to be seen.

Other Related Works: An interesting set of papers (Soudry et al., 2018; Xu et al., 2018) show that deep neural networks trained with SGD converge to the maximum margin solution in the linearly separable case. Several works investigate the possibility that deep networks give high importance to a subset of the training dataset (Toneva

et al., 2019; Vodrahalli et al., 2018; Birodkar et al., 2019), resembling supports in support vector machines. In our experiments, we find that active learning with uncertainty sampling with deep models like FTZ has a (surprisingly) large overlap with the support vectors of an SVM. Thus, it seems to have a inductive bias for class boundaries, similar to the above works. Whether this property generalizes to other deep models is yet to be seen.

7 Conclusion

We conducted a large empirical study of sampling bias and efficiency, along with algorithmic factors which impacting active text classification. We conclude that uncertainty sampling with deep models like FastText.zip exhibits negligible class bias, seems to be favorably biased to sampling data points near class boundaries, is robust to various algorithmic factors and expensive sampling strategies like ensembling offer little benefit. Also, we find a surprisingly large overlap of actively acquired points with supports of a SVM. We additionally show that uncertainty sampling can be effectively used to bootstrap the training of large DNN models by generating compact surrogate datasets (5x-40x compression). Finally, FTZ-Ent provides a strong baseline for deep active text classification, outperforming previous results by a margin of 4x less data.

The current work opens up several directions for future investigations. To list a few: (a) a deeper look into the nature of sampled data - their distribution in the feature space, as well as their importance for the task at hand; (b) the creation of surrogate datasets for a variety of applications, including hyperparameter and architecture search, etc; (c) an extension to other deep models (beyond FTZ) and beyond classification models; and, (d) an extension to semi-supervised, online and continual learning.

Acknowledgements: We thank Prof. Vineeth Balasubramium, IIT Hyderabad, India for the many helpful suggestions and discussions.

References

- Jason Baldridge and Miles Osborne. 2004. Active learning and the total cost of annotation. In *EMNLP*.
- William H. Beluch, Tim Genewein, Andreas Nurnberger, and Jan M. Kohler. 2018. The power of ensembles for active learning in image classification. In *CVPR*.
- Vighnesh Birodkar, Hossein Mobahi, and Samy Bengio. 2019. Semantic redundancies in image-classification datasets: The 10% you don't need. *arXiv preprint arXiv:1901.11409*.
- Klaus Brinker. 2003. Incorporating diversity in active learning with support vector machines. In *ICML*.
- Alexis Conneau, Holger Schwenk, Loïc Barrault, and Yann Lecun. 2017. Very deep convolutional networks for text classification. In *EACL*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Cunxiao Du, Zhaozheng Chin, Fuli Feng, Lei Zhu, Tian Gan, and Liqiang Nie. 2019. Explicit interaction model towards text classification. In *AAAI*.
- Melanie Ducoffe and Frederic Precioso. 2018. Adversarial active learning for deep networks: a margin based approach. In *ICML*.
- Sandra Ebert, Mario Fritz, and Bernt Schiele. 2012. Ralf: A reinforced active learning formulation for object class recognition. In *CVPR*.
- Atsushi Fujii, Takenobu Tokunaga, Kentaro Inui, and Hozumi Tanaka. 1998. Selective sampling for example-based word sense disambiguation. In *Computational Linguistics*.
- Yarin Gal, Riashat Islam, and Zoubin Ghahramani. 2017. Deep bayesian active learning with image data. In *ICML*.
- Daniel Gissin and Shai Shalev-Shwartz. 2019. Discriminative active learning. *ArXiv preprint arxiv:1907.06347v1*.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *ICML*.
- Steven C. H. Hoi, Rong Jin, and Michael R. Lyu. 2006. Large-scale text categorization by batch mode active learning. In *WWW*.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *ACL*.
- Rie Johnson and Tong Zhang. 2017. Deep pyramid convolutional neural networks for text categorization. In *ACL*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, H erve J egou, and Tomas Mikolov. 2016. Fasttext. zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *EACL*.
- Anders Krogh and Jesper Vedelsby. 1994. Neural network ensembles, cross validation and active learning. In *NeurIPS*.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. In *NeurIPS*.
- David D. Lewis and William A. Gale. 1994. A sequential algorithm for training text classifiers. In *SIGIR*.
- Ray Liere and Prasad Tadepalli. 1997. Active learning with committees for text categorization. In *AAAI*.
- David Lowell, Zachary C Lipton, and Byron C Wallace. 2018. How transferable are the datasets collected by active learners? *arXiv preprint arXiv:1807.04801*.
- Andrew McCallum and Kamal Nigam. 1998. Employing em and pool-based active learning for text classification. In *ICML*.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. In *JMLR*.
- Chao Qiao, Bo Huang, Guocheng Niu, Daren Li, Daxiang Dong, Wei He, Dianhai Yu, and Hua Wu. 2018. A new method of region embedding for text classification. In *ICLR*.
- Nicholas Roy and Andrew McCallum. 2001. Toward optimal active learning through sampling estimation of error reduction. In *ICML*.
- Ozan Sener and Silvio Savarese. 2018. Active learning for convolutional neural networks: A core-set approach. In *ICLR*.
- Burr Settles. 2009. Active learning literature survey. Technical report, University of Wisconsin-Madison.
- Burr Settles and Mark Craven. 2008. An analysis of active learning strategies for sequence labeling tasks. In *EMNLP*.
- Aditya Siddhant and Zachary C Lipton. 2018. Deep bayesian active learning for natural language processing: Results of a large-scale empirical study. In *EMNLP*.

- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. 2018. The implicit bias of gradient descent on separable data. *JMLR*.
- Min Tang, Xiaoqiang Luo, and Salim Roukos. 2002. Active learning for statistical natural language parsing. In *ACL*.
- Cynthia A Thompson and Raymond J Mooney. 1999. Active learning for natural language parsing and information extraction. In *ICML*.
- Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J. Gordon. 2019. An empirical study of example forgetting during deep neural network learning. In *ICLR*.
- Simon Tong and Daphne Koller. 2001. Support vector machine active learning with applications to text classification. *JMLR*.
- Kailas Vodrahalli, Ke Li, and Jitendra Malik. 2018. Are all training examples created equal? an empirical study. *arXiv preprint arXiv:1811.12569*.
- Baoxin Wang. 2018. Disconnected recurrent neural networks for text categorization. In *ACL*.
- Shiyao Wang, Minlie Huang, and Zhidong Deng. 2018. Densely connected CNN with multi-scale feature attention for text classification. In *IJCAI*.
- Sida Wang and Christopher D Manning. 2012. Baselines and bigrams: Simple, good sentiment and topic classification. In *ACL*.
- Zeyi Wen, Jiashuai Shi, Qinbin Li, Bingsheng He, and Jian Chen. 2018. ThunderSVM: A fast SVM library on GPUs and CPUs. In *JMLR*.
- Tengyu Xu, Yi Zhou, Kaiyi Ji, and Yingbin Liang. 2018. Convergence of SGD in learning ReLU models with separable data. *arXiv preprint arXiv:1806.04339*.
- Tong Zhang and F Oles. 2000. The value of unlabeled data for classification problems. In *ICML*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *NeurIPS*.
- Ye Zhang, Matthew Lease, and Byron C Wallace. 2017. Active discriminative text representation learning. In *AAAI*.
- Jingbo Zhu, Huizhen Wang, Tianshun Yao, and Benjamin K Tsou. 2008. Active learning with sampling by uncertainty and density for word sense disambiguation and text classification. In *ACL*.