

From Random to Informed Data Selection: A Diversity-Based Approach to Optimize Human Annotation and Few-Shot Learning

Alexandre Alcoforado¹, Thomas Palmeira Ferraz², Lucas Hideki Okamura¹,
Israel Campos Fama¹, Arnold Moya Lavado¹, Bárbara Dias Bueno¹,
Bruno Veloso³, Anna Helena Reali Costa¹

¹Escola Politécnica, Universidade de São Paulo, São Paulo, Brazil

²Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France

³Faculty of Economics, University of Porto and INESC TEC, Porto, Portugal
{alexandre.alcoforado, bruno.miguel.veloso, dev.arn.ml}@gmail.com
thomas.palmeira@telecom-paris.fr lucasokamura@alumni.usp.br
{israelfama, barbarabueno, anna.reali}@usp.br

Abstract

A major challenge in Natural Language Processing is obtaining annotated data for supervised learning. An option is the use of crowdsourcing platforms for data annotation. However, crowdsourcing introduces issues related to the annotator’s experience, consistency, and biases. An alternative is to use zero-shot methods, which in turn have limitations compared to their few-shot or fully supervised counterparts. Recent advancements driven by large language models show potential, but struggle to adapt to specialized domains with severely limited data. The most common approaches therefore involve the human itself randomly annotating a set of datapoints to build initial datasets. But randomly sampling data to be annotated is often inefficient as it ignores the characteristics of the data and the specific needs of the model. The situation worsens when working with imbalanced datasets, as random sampling tends to heavily bias towards the majority classes, leading to excessive annotated data. To address these issues, this paper contributes an automatic and informed data selection architecture to build a small dataset for few-shot learning. Our proposal minimizes the quantity and maximizes diversity of data selected for human annotation, while improving model performance.

1 Introduction

In real-life scenarios, particularly in the realm of Machine Learning (ML) in Natural Language Processing (NLP), annotated data is often a scarce and challenging resource to acquire. In many cases, researchers and practitioners are faced with the daunting task of developing accurate models with extremely limited or even non-existent annotated

training data. To address this challenge, the process is typically initiated by building a small annotated dataset and using it as a basis for training ML models using supervised learning methods. Subsequently, this process can be iterated by creating annotated datasets of increasing size through techniques commonly referred to as Active Learning (AL) (Ren et al., 2021).

As an alternative approach to acquiring annotated data, crowdsourcing platforms like Amazon Mechanical Turk have been used in recent years. However, relying solely on human annotation services from these platforms brings its own set of challenges (Nowak and Rüger, 2010; Karpinska et al., 2021). Variability in expertise among annotators often results in inconsistent annotation criteria and, at times, conflicting annotations. Moreover, human annotators may encounter difficulties when dealing with large datasets, leading to errors and delays in data annotation processes. An additional concern lies in the potential introduction of bias through annotators’ subjectivity and personal biases, which can negatively affect the performance of trained models. To mitigate these challenges, numerous research works have attempted to address these issues, either by selecting high-quality annotators in multiple-annotated-data setups or by employing diverse methods to weight each annotator’s input (Zhang et al., 2023a; Hsueh et al., 2009; Hovy et al., 2013; Basile et al., 2021).

In low-resource settings, a common practice is to randomly sample a subset of the unlabeled data for the annotation process (Tunstall et al., 2022; Beijbom, 2014). This approach involves selecting a few examples at random, which are then annotated to form the initial training dataset. However,

this methodology may be suboptimal since it neglects the specific characteristics of the data and the requirements of the learning model. In other words, randomly sampled data may fail to adequately represent the full spectrum of classes or concepts present within the dataset.

The advent of zero-shot methods has provided an intriguing approach to perform initial annotation without any annotated training data (Alcoforado et al., 2022). Nonetheless, historical shortcomings have often placed zero-shot methods behind their few-shot counterparts in terms of performance. Recent strides in the field of NLP, particularly the emergence of general-purpose Large Language Models (LLMs), have opened up exciting avenues in multi-task learning and zero-shot problem-solving (Ferraz et al., 2023). These models exhibit remarkable skills across various tasks (Brown et al., 2020; Touvron et al., 2023) but still encounter difficulties when adapting to specific domains where highly specialized knowledge may be entirely absent from their training data (Yang et al., 2023; Zhang et al., 2023b).

In the realm of few-shot text classification, the challenge of acquiring annotated data becomes increasingly daunting, particularly when confronted with imbalanced datasets (Ferraz et al., 2021). Common benchmark datasets used for few-shot text classification tasks often exhibit a semblance of balance or slight imbalance. However, such datasets represent rare exceptions in the real-world landscape, where data distributions are typically skewed and imbalanced, mirroring the inherent complexity of practical scenarios. The prevalence of imbalanced data poses a significant challenge, as traditional random sampling strategies become increasingly suboptimal. In scenarios where one class overwhelmingly dominates, random sampling tends to favor the majority class, resulting in data selection that inadequately represents the underrepresented and rare classes.

To address these challenges, in this paper we introduce an innovative automatic data selection architecture for few-shot learning. Our approach is designed to identify the most informative and representative data points that should be annotated by humans in low-resource, annotation-scarce scenarios. It leverages a framework that systematically orders data points based on their likelihood to (i) belong to distinct classes, thereby avoiding unnecessary redundancy in human annotation efforts, and (ii) enhance the overall performance of the

learning model. Our evaluation of this approach encompasses various low-resource natural language processing datasets, demonstrating its capacity to minimize redundancy in human annotation efforts and improve model performance compared to traditional random sampling or manual data selection strategies, particularly in cases with a limited number of annotated examples.

In summary, this work presents two primary contributions:

1. The introduction of an automatic data selection architecture for few-shot learning that leverages active learning principles to identify the most informative and representative data points for annotation.
2. An extensive analysis of various implementations of our architecture, highlighting its effectiveness to build the first version of a dataset in the context of low-resource text classification.

Our results emphasize the benefits of informed data selection, which not only streamlines the annotation process but also results in a more diverse set of annotated data. Furthermore, models trained with these diverse datasets exhibit improved performance, which may benefit subsequent iterations of the dataset with Active Learning techniques. Our experiments unveil the potential of informed data selection strategies in addressing the challenges of few-shot learning in low-resource NLP scenarios.

2 Background

In low-resource NLP settings, where annotated data is scarce and expensive to obtain, Active Learning (AL) (Ren et al., 2021) methods show themselves as a very promising approach. AL attempts to maximize the performance gain of a model by annotating the smallest number of samples. AL algorithms select data from an unlabeled dataset and query a human annotator only on this selected data, which aims to minimize human efforts in annotation by using only the most informative data.

Uncertainty sampling (Zhu et al., 2010) is among the most used method to select which points to be annotated. It employs a single classifier to pinpoint unlabeled instances where the classifier exhibits the lowest confidence. Other approaches include query-by-committee (Kee et al., 2018), where a pool of models is used to find diverse disagreements, margin sampling (Ducoffe and Precioso, 2018), and

entropy sampling (Li et al., 2011). The first one looks for points where models disagree the most on the predicted labels; while the second selects data points with the highest entropy, indicating the lowest classification probability across all potential classes

An essential aspect of AL involves the allocation of annotation budgets. Given that human effort is dedicated to annotating data, it is crucial to maximize its utility and minimize human effort. Various strategies have emerged to address this challenge. Recent research suggests optimizing directly for human effort, while others combine model uncertainty with diverse data representation through diversity sampling. A holistic approach combines these factors with cost-effectiveness, weighting data based on anticipated reductions in loss, classification entropy, and acquisition cost. These approaches collectively aim to minimize redundancy, which occurs when a human annotates a data point that the model would predict the correct label in subsequent iterations.

In this work, we deal with the very first version of a dataset, which will serve as the foundation for iterative model improvement using AL methods. Consequently, our primary focus is not on optimizing cost-effectiveness, as the data was obtained through random sampling. Instead, we are exploring alternative data selection strategies to ensure that the initial data pool closely resembles a “near-ideal” random sample. This selection should not only minimize unnecessary annotations but also elevate the model’s performance above the random average. To achieve this goal, we employ an uncertainty-based strategy to address two distinct challenges: identifying data points that are distant from the decision boundary and selecting examples that offer a more diverse and informative perspective on the dataset.

In addition to uncertainty estimation, various strategies are available for actively selecting data points to enhance low-resource NLP models. Diversity-based methods place their focus on achieving a balance between informativeness and the diversity of concepts or linguistic structures within the selected subset. This approach aims to prevent the model from learning biased information. Such balance can be achieved through techniques like calculating pairwise distances between data points and employing sampling strategies to select diverse examples. For instance, Sener and Savarese (2018) employed the cosine similarity be-

tween word vector representations and a k-center greedy algorithm to identify the most diverse subset of data. Meanwhile, Zhang et al. (2021) utilized a mutual information-based criterion to ensure that the selected data points are positioned far apart from each other in the embedding space. Additionally, there are works that combine diversity and uncertainty sampling in order to enhance the model’s performance.

3 Methods

To tackle the challenge of determining which data to annotate, we have devised Informed Data Selection methods, which, in practice, can be thought of as ordering algorithms when executed to completion. Random data selection can sometimes result in an imbalanced distribution of labels for human annotation, leading to an overabundance of certain labels while leaving others underrepresented. Our proposed methods also address this issue since the labels are not known before the annotation process. However, our findings indicate that our approach may be conducive to achieving a more equitable distribution of documents across various labels. We contend that our method is particularly well-suited for situations where humans are faced with a complete lack of labeled data. Here, a dataset consists of words, phrases or documents that must be labeled, and will be referred to in this paper as “documents”.

We have selected random sampling as our baseline method and have developed three additional methods for comparison against this baseline. These methods are constructed using distinct heuristics: (i) The first method assesses semantic similarity and prioritizes documents with low similarity to those already selected; (ii) The second method involves clustering embeddings and systematically selects one document from each cluster based on cluster size; and (iii) The third method employs random sampling to choose documents with lower lexical similarity, excluding those that share too many common n-grams. Further elaboration on these methods is provided below.

Let \mathcal{D} be a set of documents d_i . Let E be the set of embeddings for each document $d_i \in \mathcal{D}$. We define $C = \{c_1, \dots, c_{n_{classes}}\}$, $|C| = n_{classes}$, as the set of target classes for the classification task in supervised training. $\mathcal{D}_{selected}$ is the set \mathcal{D} rearranged by f according to the Informed Data Selection meth-

ods proposed here, with $|\mathcal{D}_{selected}| = |\mathcal{D}|$,

$$\mathcal{D}_{selected} = f(n_{classes}, \mathcal{D}, E), \quad (1)$$

Elements from $\mathcal{D}_{selected}$ are then selected to constitute D_a with the most relevant documents for labeling. Let n_{shots} be the target number of annotated documents per class. The set of annotated documents is $D_a = \{D_a^{c_1}, D_a^{c_2}, \dots, D_a^{c_{n_{classes}}}\}$, with $|D_a| = |D_a^{c_1}| + |D_a^{c_2}| + \dots + |D_a^{c_{n_{classes}}}|$. Ideally, we want $|D_a^{c_i}| = n_{shots}$.

The **overannotation rate** θ is defined as the excess of documents annotated with the respective method used up to the target number n_{shots} of annotated documents for each class $c_i \in C$, with:

$$\theta = |D_a| / (n_{classes} * n_{shots}). \quad (2)$$

It measures the excess of annotated documents generated by the method until the desired target n_{shots} is achieved for each specific class c_i .

We now describe the three Informed Data Selection methods proposed in this paper.

1) Reverse Semantic Search (RSS): Given a set of documents \mathcal{D} , its respective set of embeddings E , and a similarity function between pairs of embeddings $sim(x_1, x_2)$, RSS calculates the similarity matrix between all embeddings of E . The similarity matrix S is an $|\mathcal{D}| \times |\mathcal{D}|$ matrix whose (i, j) element equals the similarity $sim(e_i, e_j)$ between $e_i, e_j \in E$, with e_i and e_j being the embeddings of $d_i, d_j \in \mathcal{D}$, $d_i \neq d_j$. RSS initially selects the two documents with the least similarity and puts both in a new set named $\mathcal{D}_{selected}$. Then, iteratively, RSS continues to select the next most dissimilar element from the rest of the set $\{\mathcal{D} - \mathcal{D}_{selected}\}$. RSS stops when $|\mathcal{D}_{selected}| = |\mathcal{D}|$. In fact, RSS sorts the documents in \mathcal{D} based on their dissimilarity. The idea is that the annotation process is performed for each document in the new set generated $\mathcal{D}_{selected}$, in order, until at least n_{shots} are obtained for each of the $n_{classes}$.

2) Ordered Clustering (OC): Given a set of documents \mathcal{D} and its respective set of embeddings E , OC applies a hierarchical and density-based clustering algorithm that assigns a membership probability to each document in relation to each cluster, indicating the probability of that document being in that cluster. Then, OC orders the clusters based on their size, i.e., based on the number of documents that belong to a given cluster. Finally, OC exhaustively selects the document with the lowest membership probability from each cluster, from

largest to smallest cluster, and removes it from the cluster, placing it, in removal order, in $\mathcal{D}_{selected}$. The OC iterative process stops when all clusters are empty. Here too, the annotation process is performed for each document in the new set generated $\mathcal{D}_{selected}$, in order, until at least n_{shots} are obtained for each of the $n_{classes}$.

3) Limited Lexical Similarity (LLS): Given a set of documents \mathcal{D} , a lexical comparison function $g(d_1, d_2)$ (based on BLEU score, ROUGE score or other metrics) and a threshold value β , LLS chooses the first document d_i randomly and inserts it into the initially empty set $\mathcal{D}_{selected}$. LLS then proceeds by choosing the next document d_{i+1} at random, discarding it if $g(d_{i+1}, d_i) > \beta$ and keeping it otherwise. LLS stops when there are no more documents to select. Similar to the RSS and OC methods, the generated set can have many elements. Note that in this case, $|\mathcal{D}_{selected}|$ may be smaller than $|\mathcal{D}|$, given that some documents were discarded. Thus, the annotation takes place by removing documents from $\mathcal{D}_{selected}$, in the order in which they were inserted in $\mathcal{D}_{selected}$, until at least n_{shots} are obtained for each of the $n_{classes}$.

4 Experimental Setup

This section outlines the experimental setup for evaluating our proposed Informed Data Selection architecture. The evaluation is conducted on five text classification datasets, selected to explore varying degrees of data imbalance, class diversity, language, and domain. In this section, we present the datasets used and describe two key experimental settings: *Human Annotation* and *Few-shot learning with selected data*.

4.1 Datasets

We use the following datasets in our experiments:

- **AgNews (Zhang et al., 2015):** A news dataset with 4 classes and balanced data distribution. It consists of 120,000 training examples and 7,600 test examples, available only in English.
- **SST5 (Socher et al., 2013):** A sentiment analysis dataset with 5 classes and a slightly imbalanced data distribution. It contains 8,544 training examples, 1,101 validation examples, and 2,210 test examples, available in English.
- **Emotion (Saravia et al., 2018):** An emotion analysis dataset with 5 classes and imbalanced data distribution. It includes 16,000 training




Figure 1: Full Architecture of our Settings. Results from RQ 1 are evaluated with metric Overannotation Rate. Results from RQ 2 use metrics Accuracy and Macro-F1 Score.

examples and 2,000 test examples, available in English.

- **Multilingual Sentiment Analysis (MSA)**¹: A multilingual sentiment analysis dataset with 3 classes and balanced data distribution. We make use of the Portuguese subset of this dataset, that contains 1,839 training examples and 870 test examples.
- **BRNews**²: A Brazilian Portuguese news dataset with 19 classes and imbalanced data distribution. It comprises 176,114 training examples and 176,114 test examples, available only in Portuguese.

The train and test splits are utilized for training and evaluation, unless specified otherwise. An overview of these datasets is provided in Table 1. The choice of these datasets aims at isolating and scrutinizing key data distribution variables. Our focus centers on examining the impact of factors such as the number of samples per class, the quantity of classes within each dataset, the extent of data imbalance, and the language (English or Portuguese) on the outcomes of Informed Data Selection methods.

Table 1: Datasets Characteristics

Dataset	# docs	classes	Balancing	Lang
AgNews	127600	4	balanced	En
SST5	11855	5	slightly imbalanced	En
Emotion	18000	6	imbalanced	En
MSA	3033	3	balanced	Pt
BRNews	352228	19	very imbalanced	Pt

4.2 Research Questions

In our study, we aim to address specific research questions through distinct experimental settings, each designed to provide insights into the efficacy of our Informed Data Selection methods. These experimental settings are detailed below.

¹Available on <https://huggingface.co/datasets/tyqiangz/multilingual-sentiments>

²Available on <https://huggingface.co/datasets/iara-project/news-articles-ptbr-dataset>

4.2.1 RQ1: Which method allows for more efficient human annotation?

To tackle this question, we simulate a real-life scenario where no annotated data is initially available, and human annotators are required to annotate the data. We compare different sorting methods designed to prioritize annotation and, leveraging known ground-truth, we quantify the overannotation rate (see Eq. 1) that each method might entail. In this context, we compare the performance of our Informed Data Selection methods with that of a random sampling strategy, referred to as **Random**.

4.2.2 RQ2: Which method yields better few-shot learning?

To address this second question, we turn our attention to models trained on the dataset created in the context of RQ1. The goal is to determine whether the more efficient annotation process comes with a price, and could potentially lead to biased models, resulting in decreased performance compared to conventional random sampling. Conversely, our initial hypothesis suggests that Informed Data Selection, by increasing data diversity, will lead to model improvement, as it provides more knowledge with same amount of training data.

4.3 Evaluation Metrics

Within the context of RQ1 setting, the primary evaluation metric is the **overannotation rate** θ (Eq. 2). This metric is relevant as in resource-constrained scenarios, the imperative lies in the minimization of excessive annotation. For this metric **lower values** mean more efficiency.

As for the RQ2 setting, we employ conventional metrics commonly used in text classification. These include **Accuracy**, which measures the percentage of correctly classified instances, and, exclusively for the very imbalanced dataset, the **Macro F1-score**, a metric that calculates the harmonic mean of precision and recall for each class and then averages these values across all classes.

4.4 Implementation Details

For addressing RQ1, our chosen embedding model for RSS and OC is paraphrase-multilingual-mpnet-base-v2³. To perform clustering in OC, we employ the HDBSCAN algorithm (Campello et al., 2013). We employ BLEU score (Papineni et al., 2002) as comparison function in LLS. The entire process for LLS and Random is executed identically 10 times, and results are reported as mean values along with confidence intervals.

Regarding RQ2, we train models under two distinct configurations to isolate the influence of the training algorithm for few-shot learning. We utilize the HuggingFace Transformers library (Wolf et al., 2020) and employ the following methods:

- **FINETUNE:** We fine-tune the XLM-Roberta-large (Conneau et al., 2020), a pre-trained encoder-based Language Model, following conventional fine-tuning procedures for Sequence Classification. The training process spans 30 epochs with a learning rate of 2×10^{-5} .
- **SETFIT:** For this method, we utilize Sentence Transform fine-tuning (SetFit) (Tunstall et al., 2022), an efficient approach for few-shot learning in encoder-based models. SetFit dynamically generates training pairs from annotated data and leverages contrastive loss for training the model on the classification task. As the base model, we also use paraphrase-multilingual-mpnet-base-v2.

Results in RQ2 for LLS and Random, which exhibit stochastic behavior, are presented in terms of mean values and standard deviations across 10 runs. The experiments are conducted across a range of n_{shots} values, specifically 8, 16, 32, and 64, with a batch size of 16 for the training process.

5 Results

We compare the performance of our proposed Informed Data Selection methods with random sampling strategy on the five datasets.

5.1 Efficiency in Human Annotation (RQ1)

Charts in Figure 2 show results for experiments where we measure the overannotation rate θ as

³Available on <https://huggingface.co/sentence-transformers/paraphrase-mpnet-base-v2>




Figure 2: Overannotation Rate θ per Dataset and Method.

a function of the number of samples per class in the Dataset, for each method (RSS, OC, LLS and Random as baseline). Methods LLS and Random are executed 10 times, and averaged results (along with confidence interval) are shown.

In **balanced datasets**, we observe that **no method consistently outperforms** the random baseline. This is seen in AgNews and MSA. It can be explained, as we have mentioned before, by the distribution of classes in these datasets: both are heavily balanced, which tend to favor random sampling methods. So, when it comes to balanced data distributions, the human may not worry about over-annotation of the random method. It is interesting to note that in MSA, when n_{shots} is in the range of 30 to 60, RSS would indeed be a better choice than random sampling. Also, our methods are slightly more competent in MSA than in AgNews. The language factor may play a minor role here: because our embedding model, although multilingual, was trained on more English than Portuguese data, its embeddings are less tuned to the Portuguese language, which might explain why RSS promotes variety for a longer range of n_{shots} , but eventually converges with most other methods. Aside from this possible model-related factor, language does not seem to be a relevant factor for our selection methods.

For **imbalanced data distributions**, two of our methods consistently outperform random sampling: RSS and OC. We observe a lower overannotation rate θ in SST5 and Emotion when $n_{shots} < 30$, indicating that both RSS and OC are a better fit than random sampling in imbalanced distributions. As we increase n_{shots} further from 30, only RSS in the Emotion dataset worsens, but methods are overall more efficient in choosing which data to annotate, generating less excess of annotations.

For a **heavily imbalanced distribution**, we see a different behavior. We observe that as **number of classes and data imbalance grow, overannotation rate θ increases** for every method tested (BRNews has 10 times more overannotation rate than balanced datasets). In turn, OC generates too much overannotation rate (more than 6 times than Random baseline), and is thus considered an outlier and excluded from the chart.

Results show that **RSS considerably outperforms Random baseline** for $n_{shots} < 40$. This is once again due to the fact that this dataset has a much higher number of classes, with very imbalanced distribution of documents per each class,

much closer to a real-life scenario humans find themselves. In these scenarios, our method thrives, generating as few as half excess annotations when compared to the Random method. However, as observed for every dataset, our methods and Random baseline also converge when n_{shots} increases further away from around 50.

5.2 Model Performance (RQ2)

Figure 3 shows results for experiments where we compare the performance of classifiers trained with data selected by our methods and the Random method. Because OC fails to generate a feasible excess of annotation for BRNews, it is deemed as not applicable and therefore excluded from reports.

As a general result, we observe that our methods **OC and LLS fail to consistently outperform** the Random baseline. However, **RSS outperforms random sampling in almost every scenario**. For both FINETUNE and SETFIT, RSS is better than random sampling for every dataset with the exception of the AgNews, where random sampling yields higher accuracy. A mix of many factors may be responsible for this: first, AgNews is balanced, which favors random sampling when selecting training data; second, the task of AgNews is simple when compared to other datasets, because classes in it have distinct traits (ie. they refer to distinct themes, such as Sports, Technology, etc) which may help with decision boundaries of the model. The other balanced dataset, MSA, does not have these distinct traits for its classes, which instead express a kind of gradation (ie. Positive, Neutral, Negative). In other words, the classification task in MSA is tougher, which means that selecting data with more variability can effectively boost model performance.

We note that **the higher the degree of data imbalance, the more consistently RSS will outperform random sampling**. However, reporting only accuracy in a heavily imbalanced dataset is insufficient to adequately represent performance of a classifier. Thus, Table 2 shows results of **Macro-F1 Score** for both training methods in BRNews. We see that for FINETUNE, **RSS performs consistently better**, while SETFIT also shows a slight improvement when compared to Random, falling above the confidence interval only for $n_{shots} = 8$. This is an indicative that **both RSS and Random methods perform almost equally well across classes**, disconsidering imbalance among classes. This indicates that both methods succeed at selecting diverse data for model training. Still, **RSS**

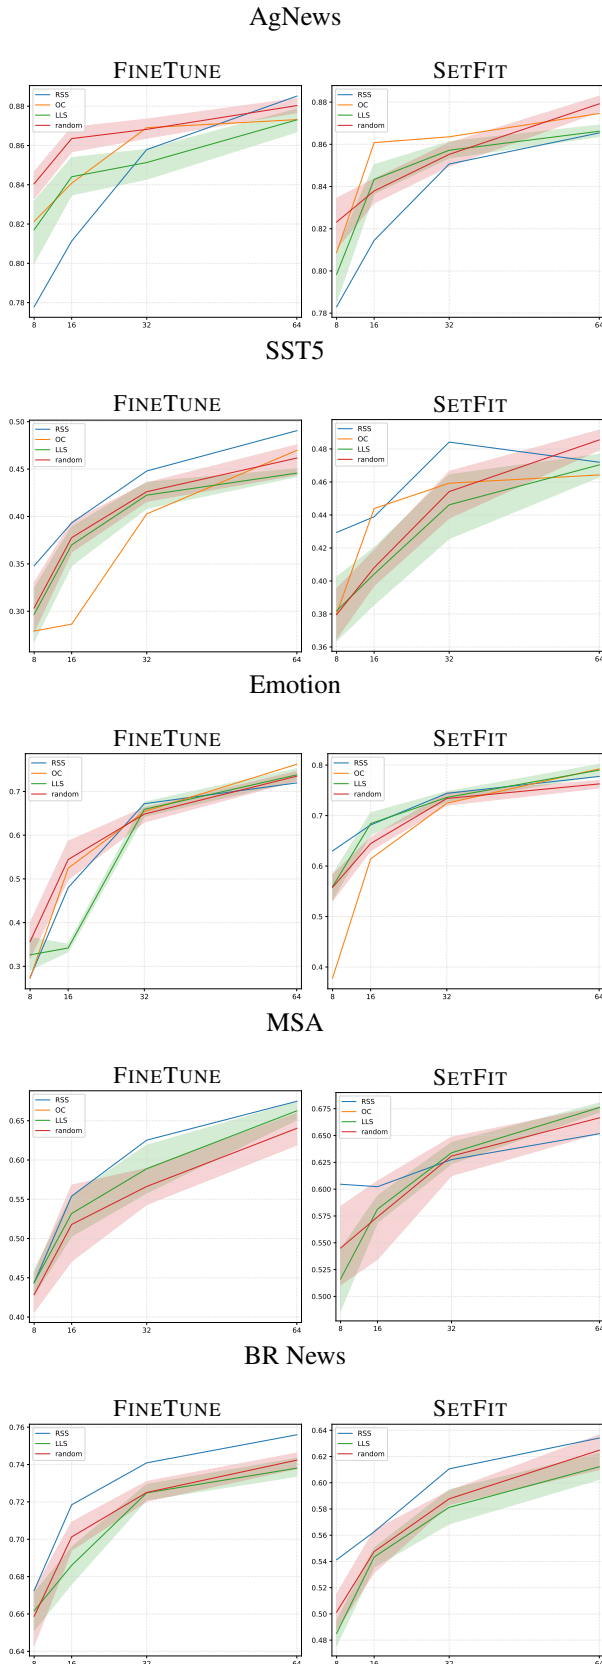


Figure 3: Accuracy over evaluation datasets.

provides **higher accuracy**, turning it into the **recommended method for low-resource setups**.

Table 2: FineTune and SetFit F1-Score Macro for Random Selection on BRNews dataset.

Training	n_{shots}	RSS	Random
FINETUNE	8	58.6	56.7 ± 2.3
FINETUNE	16	62.0	60.6 ± 0.8
FINETUNE	32	65.2	62.8 ± 0.9
FINETUNE	64	66.8	63.6 ± 0.5
SETFIT	8	46.83	45.0 ± 1.7
SETFIT	16	48.8	48.8 ± 1.9
SETFIT	32	52.3	52.2 ± 1.0
SETFIT	64	55.9	55.6 ± 1.2

Another important result is **the convergence of all methods when n_{shots} grows**. Because our methods are suited to the construction of a very first version of a dataset for Active Learning, both overannotation rate and model performance converge when $n_{shots} > 64$. A reason is that, as the number of selected data grows, diversity will also grow. Although results show that our selection methods promote more diversity for lower n_{shots} , any selection method that does not apply oversampling will bring diversity if n_{shots} keeps increasing. Thus, other methods outpace ours in promoting diversity when we leave the realm of few-shot – i.e. when we annotate too much data. This means that **when the desire is to annotate lots of data per class, most methods evaluated in this work are not suited to the selection**, with random sampling being a better strategy.

6 Conclusion

This work has proposed an automatic Informed Data Selection architecture which aims to select which data should be annotated by a human to build a first dataset. We simulated 2 scenarios, and experimental results we report show our architecture is a better option than random sampling methods for few-shot learning. We have shown that the higher the imbalance in the dataset, the more competent our method is – both in generating less excess of annotation and in improving model performance. As far as we know, there are few works that address the imbalance problem as a variable of a supervised dataset.

In particular, the Reverse Semantic Search (RSS) method has shown to be the most competent in experiments across different languages, number and imbalance across classes. However, it should be noted that for Limited Lexical Similarity (LLS), a numeric threshold is specified. This work has not

fine-tuned this hyperparameter, instead it was chosen by manually inspecting results with different thresholds.

Results indicate that fine-tuning this threshold can improve the overannotation rate and accuracy of LLS, as its standard deviation in many datasets is smaller than that of the Random method. The same can be said about using another comparison function, such as ROUGE score. The need for fine-tuning may be amplified in very specialized domains with unusual vocabulary – however, more experiments are needed to confirm this observation.

We also note that the Ordered Clustering (OC) method did not provide consistent results across many datasets. Because our method relies on picking one document from each cluster, when the number of identified clusters is high, OC fail to select quality data. This can be addressed by combining clustering with RSS or LLS, and is an attractive direction for future work.

Our concluding remark underscores a frequently overlooked aspect in the realm of few-shot learning, particularly in scenarios where labeled data is scarce. Many studies in few-shot learning often assess their methods across a range of datasets, typically characterized by a balance or slight imbalance in class distribution. While there are exceptions, those works that do evaluate on imbalanced datasets often fail to adequately address the consequences of such imbalance. The reality is that in real-world applications, balanced data distributions are a rarity. Hence, we advocate that authors engaged in few-shot learning techniques should be cognizant of this reality, and whenever feasible, report metrics that account for imbalance in their evaluations.

Acknowledgements

This work is financed by National Funds through the Portuguese funding agency, FCT - Fundação para a Ciência e a Tecnologia, within project LA/P/0063/2020. DOI 10.54499/LA/P/0063/2020 | <https://doi.org/10.54499/LA/P/0063/2020>. This work is also supported in part by the Brazilian National Council for Scientific and Technological Development (CNPq) under grant numbers 310085/2020-9, and the Brazilian Coordination for the Improvement of Higher Education Personnel (CAPES, Finance Code 001). Israel Campos Fama is funded by the Secretaria de Fazenda do Estado do Rio Grande do Sul (Sefaz-RS). Bárbara Dias

Bueno is funded by the *Programa Unificado de Bolsas (PUB)* undergraduate research scholarship from Universidade de São Paulo, under project 2117/2023.

References

- Alexandre Alcoforado, Thomas Palmeira Ferraz, Rodrigo Gerber, Enzo Bustos, André Seidel Oliveira, Bruno Miguel Veloso, Fabio Levy Siqueira, and Anna Helena Reali Costa. 2022. [ZeroBERTo: Leveraging Zero-Shot Text Classification by Topic Modeling](#). In *International Conference on Computational Processing of the Portuguese Language (PROPOR 2022)*, pages 125–136. Springer.
- Angelo Basile, Guillermo Pérez-Torró, and Marc Franco-Salvador. 2021. [Probabilistic Ensembles of Zero- and Few-Shot Learning Models for Emotion Classification](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 128–137, Held Online.
- Oscar Beijbom. 2014. [Random Sampling in an Age of Automation: Minimizing Expenditures through Balanced Collection and Annotation](#). *arXiv preprint arXiv:1410.7074*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901.
- Ricardo J. G. B. Campello, Davoud Moulavi, and Joerg Sander. 2013. [Density-based clustering based on hierarchical density estimates](#). In *Advances in Knowledge Discovery and Data Mining*, pages 160–172, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Melanie Ducoffe and Frederic Precioso. 2018. [Adversarial Active Learning for Deep Networks: a Margin Based Approach](#). *arXiv preprint arXiv:1802.09841*.

- Thomas Palmeira Ferraz, Alexandre Alcoforado, Enzo Bustos, André Seidel Oliveira, Rodrigo Gerber, Naíde Müller, André Corrêa d’Almeida, Bruno Miguel Veloso, and Anna Helena Reali Costa. 2021. [DEBACER: a method for slicing moderated debates](#). In *Anais do XVIII Encontro Nacional de Inteligência Artificial e Computacional*, pages 667–678. SBC.
- Thomas Palmeira Ferraz, Marcely Zanon Boito, Caroline Brun, and Vassilina Nikoulina. 2023. [Multilingual DistilWhisper: Efficient Distillation of Multi-task Speech Models via Language-Specific Experts](#). *arXiv preprint arXiv:2311.01070*.
- Dirk Hovy, Taylor Berg-Kirkpatrick, Ashish Vaswani, and Eduard Hovy. 2013. [Learning whom to trust with MACE](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1120–1130, Atlanta, Georgia. Association for Computational Linguistics.
- Pei-Yun Hsueh, Prem Melville, and Vikas Sindhwani. 2009. [Data quality from crowdsourcing: A study of annotation selection criteria](#). In *Proceedings of the NAACL HLT 2009 Workshop on Active Learning for Natural Language Processing*, pages 27–35, Boulder, Colorado. Association for Computational Linguistics.
- Marzena Karpinska, Nader Akoury, and Mohit Iyyer. 2021. [The perils of using Mechanical Turk to evaluate open-ended text generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1265–1285, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Seho Kee, Enrique del Castillo, and George Runger. 2018. [Query-by-committee improvement with diversity and density in batch active learning](#). *Information Sciences*, 454–455:401–418.
- X. Li, E. Gavves, C. G. M. Snoek, M. Worring, and A. W. M. Smeulders. 2011. [Personalizing automated image annotation using cross-entropy](#). In *ACM International Conference on Multimedia*, pages 233–242.
- Stefanie Nowak and Stefan Rügger. 2010. [How reliable are annotations via crowdsourcing: a study about inter-annotator agreement for multi-label image annotation](#). In *Proceedings of the International Conference on Multimedia Information Retrieval, MIR ’10*, page 557–566, New York, NY, USA. Association for Computing Machinery.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B. Gupta, Xiaojiang Chen, and Xin Wang. 2021. [A Survey of Deep Active Learning](#). *ACM Computing Surveys*, 54(9).
- Elvis Saravia, Hsien-Chi Toby Liu, Yen-Hao Huang, Junlin Wu, and Yi-Shin Chen. 2018. [CARER: Contextualized affect representations for emotion recognition](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3687–3697, Brussels, Belgium. Association for Computational Linguistics.
- Ozan Sener and Silvio Savarese. 2018. [Active learning for convolutional neural networks: A core-set approach](#). In *International Conference on Learning Representations*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew 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*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Lewis Tunstall, Nils Reimers, Unso Eun Seo Jo, Luke Bates, Daniel Korat, Moshe Wasserblat, and Oren Pereg. 2022. [Efficient few-shot learning without prompts](#). *arXiv preprint arXiv:2209.11055*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Bing Yin, and Xia Hu. 2023. [Harnessing the power of llms in practice: A survey on chatgpt and beyond](#). *arXiv preprint arXiv:2304.13712*.
- Dequan Zhang, Pengfei Zhou, Chen Jiang, Meide Yang, Xu Han, and Qing Li. 2021. [A stochastic process discretization method combining active learning kriging model for efficient time-variant reliability analysis](#). *Computer Methods in Applied Mechanics and Engineering*, 384:113990.

- Lining Zhang, Simon Mille, Yufang Hou, Daniel Deutsch, Elizabeth Clark, Yixin Liu, Saad Mahamood, Sebastian Gehrmann, Miruna Clinciu, Khyathi Raghavi Chandu, and João Sedoc. 2023a. [A needle in a haystack: An analysis of high-agreement workers on MTurk for summarization](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14944–14982, Toronto, Canada. Association for Computational Linguistics.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Yating Zhang, Yexiang Wang, Fei Cheng, Sadao Kurohashi, et al. 2023b. [Reformulating domain adaptation of large language models as adapt-retrieve-revise](#). *arXiv preprint arXiv:2310.03328*.
- Jingbo Zhu, Huizhen Wang, Benjamin K. Tsou, and Matthew Ma. 2010. [Active learning with sampling by uncertainty and density for data annotations](#). *IEEE Transactions on Audio, Speech, and Language Processing*, 18(6):1323–1331.