

A comprehensive exploitation of instance selection methods for automatic text classification

Washington Cunha¹, Leonardo Rocha² (Co-advisor), Marcos Gonçalves¹ (Advisor)

¹Department of Computer Science – Federal University of Minas Gerais – Brazil

²Department of Computer Science – Federal University of São João del Rei – Brazil

{washingtoncunha,mgoncalv}@dcc.ufmg.br, lcrocha@ufsj.edu.br

Abstract. *Progress in Natural Language Processing (NLP) has been dictated by the rule of more: more data, more computing power and more complexity, best exemplified by the Large Language Models. However, training (or fine-tuning) large dense models for specific applications usually requires significant amounts of computing resources. This **Ph.D. dissertation** focuses on an under-investigated NLP data engineering technique, whose potential is enormous in the current scenario known as Instance Selection (IS). The IS goal is to reduce the training set size by removing noisy or redundant instances while maintaining the effectiveness of the trained models and reducing the training process cost. We provide a comprehensive and scientifically sound comparison of IS methods applied to an essential NLP task – Automatic Text Classification (ATC), considering several classification solutions and many datasets. Our findings reveal a significant untapped potential for IS solutions. We also propose two novel IS solutions that are noise-oriented and redundancy-aware, specifically designed for large datasets and transformer architectures. Our final solution achieved an average reduction of 41% in training sets, while maintaining the same levels of effectiveness in all datasets. Importantly, our solutions demonstrated speedup improvements of 1.67x (up to 2.46x), making them scalable for datasets with hundreds of thousands of documents. This thesis falls under the Information Systems and Artificial Intelligence (Generative, LLM, NLP, among others) topic of the SBSI-CTDGI call for papers.*

Student Level: Ph.D., concluded on August 26th, 2024

Board Members: Franco Maria Nardini (ISTI-CNR); Thierson Couto Rosa (UFG); Rodrygo Luis Teodoro Santos (DCC-UFMG); Anisio Mendes Lacerda (DCC-UFMG)

Dissertation available at: <http://hdl.handle.net/1843/76441>

Publications available at: <http://bit.ly/3WvX1T5>

1. Introduction

The exponential growth in data availability has made it challenging to organize and retrieve content effectively. In this context, Automatic Text Classification (ATC) offers a solution by mapping textual documents (e.g., web pages, emails, reviews, tweets, social media messages) into predefined categories of interest. Indeed, ATC models are crucial for emerging applications [Cunha et al. 2023a] such as fake news detection, hate speech identification, relevance feedback, sentiment analysis, product review analysis,

election vote inference, and assessing public agency satisfaction. As a supervised task, ATC benefits from applications generating large volumes of labeled data, such as social networks like Twitter. Moreover, crowdsourcing [Sindagi et al. 2020] and soft labeling [Roy and Cambria 2022] further reduce the costs of acquiring labeled datasets.

Transformer-based architectures, including small and large language models (SLMs and LLMs) such as RoBERTa, GPT-4, and LLama3, represent the state-of-the-art (SOTA) in ATC, achieving outstanding results through pre-training on large datasets and fine-tuning on domain-specific tasks. While zero-shot capabilities are feasible, fine-tuning remains essential for optimal performance, leveraging pre-trained weights for task-specific adaptations [de Andrade et al. 2023].

According to Andrew Ng, the success of these (language) models is due to extensive pre-training on massive datasets (e.g., 45TB for GPT-3) and the adaptability of pre-trained models via fine-tuning. This approach enables faster task-specific training compared to starting from scratch [Uppaal et al. 2023]. However, fine-tuning remains resource-intensive despite being faster than full training, requiring significant computational power and time. For instance, fine-tuning XLNet (one of the datasets considered in our dissertation) on the MEDLINE dataset took 80 GPU hours. Therefore, practical constraints often delay re-training, impacting model performance in scenarios requiring continuous updates, such as fraud detection and recommendation systems.

Indeed, resource limitations in companies and research groups restrict experimentation with deep learning models, which require extensive computational resources. For example, in this dissertation, we ran 4,000 experiments that took approximately 5,600 hours. Reducing financial, computational, and environmental costs is crucial, as shown by the significant energy consumption and carbon emissions associated with large models. Given increasing data volumes, re-training demands, and environmental concerns, proposing scalable and cost-effective strategies is essential. These include creating efficient deep learning algorithms, using advanced hardware, or improving data preprocessing techniques. This dissertation focuses on the latter, aiming to enhance model performance while reducing costs.

In [Cunha et al. 2020], we proposed exploiting a set of preprocessing techniques in a data transformation pipeline for building cost-effective models. That solution improved effectiveness (e.g., models with higher accuracy) at a much lower cost (e.g., shorter time for ATC model construction). One of the main contributions of this work was the explicit incorporation into the pipeline of an **Instance Selection** (IS) stage, a promising set of techniques and growing research area that helps to deal with many of the aforementioned issues.

Specifically, in contrast to traditional Feature Selection approaches [Lu et al. 2017], in which the main objective is to select the most informative terms (words), **Instance Selection** methods are focused on selecting the most representative instances (documents) for the training set [Garcia et al. 2012]. The intuition behind this type of algorithm is to remove potentially noisy or redundant instances from the original training set and improve performance in terms of total time training time while keeping or even improving effectiveness. More specifically, IS methods have three main goals: (i) to reduce the number of instances by selecting the most representative

ones; (ii) to maintain (or even improve) effectiveness by removing noise and redundancy; and (iii) to reduce the total time for applying an end-to-end model (which includes from traditional preprocessing steps to the model training step). By selecting the most representative instances, IS methods can also potentially remove noise from erroneous annotations. According to these objectives, IS methods must respect three fundamental constraints consisting of *reducing the amount of training without loss of effectiveness and with efficiency gains*. IS methods seek to optimize these three constraints simultaneously.

Despite their potential, IS approaches have been understudied in ATC, particularly in deep learning [Barigou 2018, Tsai et al. 2014]. Traditional IS methods developed for other domains include **CNN** [Hart 1968], **ENN** [Wilson 1972], and **Drop3** [Wilson and Martinez 2000]. Recent methods, such as **LSSm**, **LSBo** [Leyva et al. 2015], and **PSDSP** [Carbonera and Abel 2018], have mainly been tested on small tabular datasets with weak classifiers such as KNN. In contrast, text classification datasets are unstructured, larger, and more complex, featuring high dimensionality and skewness. Given the high computational costs of deep learning transformer models with large training datasets, presenting an ideal scenario for applying IS techniques.

2. Hypothesis and Research Questions

The main hypothesis (H1) of this Ph.D. dissertation is:

H1: It is possible to simultaneously reduce data, maintain model quality, and improve time for fine-tuning ATC models through IS methods.

In order to confirm this hypothesis, we propose **three** research questions for our Ph.D. dissertation. In sum, **RQ1** aims at evaluating the traditional IS approaches shall be present in Section 4 in the context of the ATC task concerning the posed IS methods tripod constraints: reduction, effectiveness, and efficiency. Next, considering the literature gap, **RQ2** aims at demonstrating the feasibility of proposing a novel IS framework and by showing how it can accommodate different requirements posed by distinct scenarios, mainly those associated with big data. Finally, **RQ3** aims to investigate the capability of each IS method to handle and effectively remove noisy instances. The answer to this question also motivated us to demonstrate the feasibility of proposing a novel extended IS framework capable of removing simultaneously redundant and noisy instances from the training set. Next, we present each of the RQs considered in this Ph.D. dissertation.

- RQ1. What is the impact of applying traditional IS methods in the ATC context regarding the posed constraints?** RQ1 aims to evaluate the traditional IS approaches in the context of the ATC task, focusing on the tripod constraints of the posed IS methods: reduction, effectiveness, and efficiency.
- RQ2. Can a novel instance selection method focused on redundancy removal overcome the limitations of existing IS methods to achieve the tripod restrictions in the ATC scenario?** This RQ aims to demonstrate the feasibility of proposing a novel IS framework and by showing how it can accommodate different requirements posed by distinct scenarios, mainly those associated with big data.
- RQ3. Is it possible to extend the previous proposal to not only remove redundancy but also remove noise, enhancing the level of quality considering all tripod criteria?** The objective of this RQ is to demonstrate the feasibility of proposing a novel extended IS framework capable of removing simultaneously redundant and noisy instances from the training set.

3. Publications

Our work in the Instance Selection field has been validated and published in the main Information Systems conferences and journals in the last four years only, including **two** published papers in the *Information Processing and Management (IP&M)* (*h5-index: 96, Impact Factor: 7.4, A1*), a worldwide leading journal in *Information Retrieval* [Cunha et al. 2021, Cunha et al. 2020]. There is also a publication in *ACM Computing Surveys (CSUR)* (*h5-index: 103, IF: 23.8, A1*) [Cunha et al. 2023c], a leading journal in *Computer Science Theory and Methods*, covering comprehensive literature review of both traditional and SOTA methods in the IS field as well as an extensive experimental comparison between these methods; the *International Conference SIGIR* [Cunha et al. 2023b], covering our first proposal of a novel redundancy-oriented IS framework aimed at large datasets with a particular focus on transformers. Finally, this dissertation also resulted in an *ACM Transactions on Information Systems (TOIS)* [Cunha et al. 2024] (*h5-index: 48, IF: 5.6, A1*), a worldwide leading journal in **Information Systems**, which includes the proposal of an extended noise-oriented and redundancy-aware IS framework for ATC.

1. **Cunha, Washington**, et al. “A Noise-Oriented and Redundancy-Aware Instance Selection Framework.” *ACM Transactions on Information Systems (ACM TOIS)* (2024)– *h5-index: 48.0*
2. **Cunha, Washington**, et al. “An effective, efficient, and scalable confidence-based instance selection framework for transformer-based text classification.” *ACM SIGIR 2023* – *h5-index: 103.0*.
3. **Cunha, Washington**, et al. “A Comparative Survey of Instance Selection Methods applied to Non-Neural and Transformer-Based Text Classification.” *ACM Computing Surveys* (2023) – *h5-idx: 157.0*
4. **Cunha, Washington**, et al. “On the cost-effectiveness of neural and non-neural approaches and representations for text classification: A comprehensive comparative study.” *IP&M* (2021) - *h5-index: 114.0*
5. **Cunha, Washington**, et al. “Extended pre-processing pipeline for text classification: On the role of meta-feature representations, sparsification and selective sampling.” *IP&M* (2020) – *h5-index: 114.0*

Our dissertation contributed **directly** to other journals during the doctorate period:

1. Gonçalves, M., **Cunha, W** et al. (2024) “Mais com menos - processamento de linguagem natural inteligente e sustentável baseado em engenharia de dados e inteligência artificial avançada.” **IV Seminário de Grandes Desafios da Computação no Brasil 2025-2035** - Sociedade Brasileira de Computação.
2. França, C., **Cunha, W** et al. (2024). On representation learning-based methods for effective, efficient, and scalable code retrieval. *Neurocomputing*. – *h5-index: 136.0*
3. Andrade, C., **Cunha, W** (2023) On the class separability of contextual embeddings representations – or “the classifier does not matter when the (text) representation is so good!”. *IP&M*. – *h5-index: 114.0*
4. Zanotto, B. S., **Cunha, W** et al. (2021). Pcv50 automatic classification of electronic health records for a value-based program through machine learning. *Value in Health*. – *h5-index: 57.0*
5. Zanotto, B. S., **Cunha, W** et al. (2021). Stroke outcome measurements from electronic medical records: Cross-sectional study on the effectiveness of neural and nonneural classifiers. *JMIR Med* – *h5-idx: 52.0*
6. Viegas, F. et al., **Cunha, W** (2024). Exploiting contextual embeddings in hierarchical topic modeling and investigating the limits of the current evaluation metrics. *Computational Linguistics*. – *h5-idx: 38*
7. Felix, L. et al., **Cunha, W** (2024). Why are you traveling? Inferring trip profiles from online reviews and domain-knowledge. *Online Social Networks and Media*. – *h5-index: 28.0*
8. Viegas, F., **Cunha, W** et al. (2024). Pipelining semantic expansion and noise filtering for sentiment analysis of short documents – clusent method. *Journal on Interactive Systems*. – *h5-index: 9.0*
9. Cunha, W et al. (2021) Extended Pre-Processing Pipeline For Text Classification: On the Role of Meta-Features, Sparsification and Selective Sampling. *CTDBD - SBB'D21* – *h5-index: 7.0*
10. Gomes, C., **Cunha, W** et al. (2019). CluWords: Explorando Clusters Semânticos entre Palavras para Aprimorar Modelagem de Tópicos. *CTIC - Iniciação Científica em Computação*. – *h5-index: 1.0*

Our dissertation contributed **directly** to other conference papers during the Ph.D. period:

1. Viegas, F., **Cunha, W** et al. (2020). Cluhtm - semantic hierarchical topic modeling based on cluwords. Meeting of the Association for Computational Linguistics (ACL) *ACL'20*. – *h5-index: 215.0*
2. Mendes, L., **Cunha, W** et al. (2020). “Keep it Simple, Lazy”—Metalazy: A new metastrategy for lazy text classification. *CIKM'20*. – *h5-index: 91.0*

3. Viegas, F., **Cunha, W** et al. (2019). Cluwords: Exploiting semantic word clustering for enhanced topic modeling. **WSDM'19**. – h5-index: 77.0
4. Andrade et al., **Cunha, W** (2024) Explaining the hardest errors of contextual embedding-based classifiers. Conference on Computational Natural Language Learning (**CoNLL'24**). – h5-index: 39.0
5. Pasin, A., **Cunha, W** et al. (2024). A quantum annealing instance selection approach for efficient and effective transformer fine-tuning. **ACM SIGIR ICTIR**. – h5-index: 24.0
6. Fonseca, G., **Cunha, W** et al. (2024). Estratégias de Undersampling para Redução de Viés em Classificação de Texto Baseada em Transformers. **WebMedia'24**. – h5-index: 13.0
7. Vasconcelos, N., **Cunha, W** et al. (2024). Integrando Avaliações Textuais de Usuários em Recomendação baseada em Aprendizado por Reforço. **WebMedia'24**. – h5-index: 13.0
8. Viegas, F., **Cunha, W** et al. (2023). Clusent–combining semantic expansion and de-noising for dataset-oriented sentiment analysis of short texts. **WebMedia'23**. – h5-index: 13.0
9. Cecilio, P., **Cunha, W** et al. (2023). Um framework para extração automática de informações em patentes farmacêuticas. **WebMedia'23**. – h5-index: 13.0
10. Júnior, A., **Cunha, W** et al. (2022) Evaluating topic modeling pre-processing pipelines for portuguese texts. **WebMedia'22** – h5-index: 13.0
11. Pasin, A., **Cunha, W** et al. (2024). A quantum annealing-based instance selection approach for transformer fine-tuning. **Italian Information Retrieval Workshop**. – h5-index: 7.0
12. Santos, W., **Cunha, W** et al. (2023). Uma metodologia para tratamento do viés da maioria em modelos de stacking via identificação de documentos difíceis. **SBB'D'23**. – h5-index: 7.0
13. Macul, V., **Cunha, W** et al. (2024). Inteligência Artificial Generativa para Personalização do Cuidado de Saúde Integral: Processo de caracterização dos dados. **II Simpósio CI-IA Saúde da UFMG**.

4. Systematic Literature Review of Instance Selection Methods

In this section, we present a critical analysis (*a.k.a., rapid (systematic-based) literature review*) of the most traditional and/or recent (state-of-the-art) proposals in the Instance Selection (IS) area. This review aims to comprehensively assess the most relevant works related to IS strategies applied in different scenarios. In particular, we focus on experimentally oriented studies, that is, studies that have strong experimental and empirical components to support their findings.

To achieve our objective, we collected a set of 100 publications¹ that included the most cited articles related to IS. We assume that highly cited articles are potentially influential as they have received much attention. From those, we selected the most popular methods to include in our experimental assessment of IS methods applied to ATC. We also selected a set of recently proposed methods (considered state-of-the-art) to complete our experimental comparison. At the end of this literature review process, which is further detailed next, we end up with a mix of the traditional and state-of-the-art set of methods, comprising 13 IS strategies to be evaluated in the next section in the ATC scenario.

A simplified version of the rapid review procedure is shown in Figure 1. First, we collected articles returned by a set of four queries submitted to Google Scholar. Second, we used each article's unique URL to remove duplicates (deduplication phase). This procedure resulted in 1,740 unique articles. Third, we ranked the remaining articles by the number of citations. Fourth, we classified the articles, in ranked order, according to the desired criteria of (i) being related to IS and (ii) conducting experimental comparisons among methods. Fifth, we filtered out articles that did not match the criteria, up to the point that we achieved 100 relevant articles. Up to this point, we had inspected 702 articles in ranked order. From these articles, we selected the most popular and most recent methods to be compared in our experimental assessment.

¹The list of articles can be found at <https://shorturl.at/zCLW7>




Figure 1. Workflow of study selection and analysis of the (rapid) systematic literature review of instance selection methods.

The results of our searches and analyses reinforced our perception that IS methods are almost exclusively applied to tabular structured data — and their application in NLP tasks is rare, which is odd since this is one of the areas that could benefit most from this type of method. In sum, from the 100 considered IS papers, 92 of them considered just tabular data. We propose to investigate the use of IS methods along with ATC models, which have been highly popular in applications as diverse as the detection of fake news and hate speech, sentiment analysis, revision of product characteristics, inferring opinions, and assessing the satisfaction of products and services, among many others.

We also propose a new and innovative taxonomy of IS strategies by extending a 10-year-old taxonomy proposed in [Garcia et al. 2012]. This previous IS taxonomy was proposed in a different context, considering different types of datasets (small tabular ones) and learning methods. For instance, deep learning neural network methods were not considered in the envisioned scenarios by the time that taxonomy was proposed.

As mentioned, the field has significantly evolved since the proposal of the original taxonomy. These advances, which we have carefully documented in our new extended taxonomy, are reflected in our new extended taxonomy with three new categories – represented in green in Figure 2 – and eight new, recently proposed methods – written with black color letters in Figure 2. The new categories refer to approaches based on density, spatial hyperplanes, and clustering-based approaches proposed since 2015.

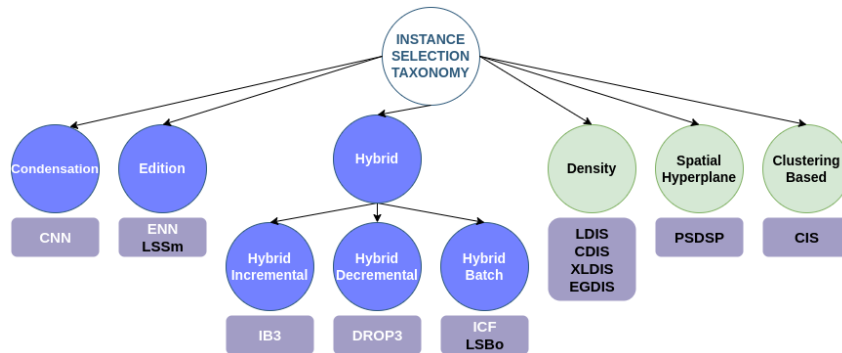


Figure 2. New Extended Instance Selection Taxonomy Proposal

Last, in our Ph.D. Dissertation, we also provided detailed information about the thirteen IS approaches evidenced in our critical analysis (not presented here due to space constraints) that we will consider in our experimentation evaluation for the ATC context. Therefore, in the next section, we provide a complete experimental evaluation comparing the IS methods detailed above, combining them with SOTA ATC strategies. Our objective with this set of experiments is to answer the first posed research question (RQ1).

5. A Comparative Survey of Instance Selection Methods applied to NonNeural and Transformer-Based Text Classification

In this section, we propose to assess the trade-off among reduction, efficiency, and effectiveness of these 13 most representative traditional IS methods (Section 4) applied to the ATC task using large and varied datasets. It is essential to notice that the selected IS methods presented in the previous section have been tested only with small structured tabular datasets (such as those from the UCI repository). Regarding ATC methods, we considered the current SOTA in the ATC field: the transformer-based architectures, such as BERT, XLNet, RoBERTa, and others. These methods have a high cost regarding computational resources, mainly when dealing with large labeled training data. Therefore, they constitute an ideal scenario for the application of IS techniques. Last, but not least important, we should emphasize that our work, as far as we know, is the first to apply IS as a preprocessing step before using transformer-based architectures in the ATC context. This contribution is realized using an experimental study whose rigor and magnitude (7 transformer methods and 13 IS approaches) have not yet been reported in the literature on IS.

Next, we present the comparative Scenario and experimental setup proposed to assess the trade-off among reduction, efficiency, and effectiveness of these 13 most representative traditional IS methods applied to the ATC.

Task	Dataset	Size	Dim.	# Classes	Density	Skewness
Topic	DBLP	38,128	28,131	10	141	Imbalanced
	Books	33,594	46,382	8	269	Imbalanced
	ACM	24,897	48,867	11	65	Imbalanced
	20NG	18,846	97,401	20	96	Balanced
	OHSUMED	18,302	31,951	23	154	Imbalanced
	Reuters90	13,327	27,302	90	171	Extremely Imbalanced
	WOS-11967	11,967	25,567	33	195	Balanced
	WebKB	8,199	23,047	7	209	Imbalanced
	Twitter	6,997	8,135	6	28	Imbalanced
	TREC	5,952	3,032	6	10	Imbalanced
Sentiment	WOS-5736	5,736	18,031	11	201	Balanced
	SST1	11,855	9,015	5	19	Balanced
	pang_movie	10,662	17,290	2	21	Balanced
	Movie Review	10,662	9,070	2	21	Balanced
	vader_movie	10,568	16,827	2	19	Balanced
	MPQA	10,606	2,643	2	3	Imbalanced
	Subj	10,000	10,151	2	24	Balanced
	SST2	9,613	7,866	2	19	Balanced
Large	yelp_reviews	5,000	23,631	2	132	Balanced
	AGNews	127,600	39,837	4	37	Balanced
	Yelp_2013	335,018	62,964	6	152	Imbalanced
	MEDLINE	860,424	125,981	7	77	Extremely Imbalanced

Tabela 1. Datasets Statistics, including: size, dimensionality, classes, document density, and class distribution.

Datasets, Data Representation, and Preprocessing We adopted the 22 real-world datasets collected from various sources in two broad ATC tasks [Li et al. 2022]: i) *topic classification*; and ii) *sentiment analysis*. The datasets cover a range of domains, diversity in size, dimensionality, classes, document density, and class distribution.

The TFIDF representation is input to all IS methods. Other text representations could have been adopted as input to IS methods, such as *static embeddings* (e.g., FastText [Joulin et al. 2017]) or *contextualized embeddings* built by transformer architectures (whether by forwarding documents through fine-tuned model or a zero-shot approach).

However, as previously demonstrated in [Andrade et al. 2024]: (i) static embeddings can slow down classification methods significantly; and (ii) using contextualized embeddings directly as IS input is inefficient and ineffective, probably because it sacrifices sparsity. In this context, given the goal of finding cost-effective solutions, the TF-IDF representation is still the best alternative as input for the instance selection methods.

We removed stopwords and kept features appearing in at least two documents. We normalized the TF-IDF product result using the L2-norm. In practice, as illustrated in **Figure 3**, we first split the dataset employing the Stratified K-Fold cross-validation methodology – the smaller datasets were executed using k=10-fold partition, while for the larger ones, we adopted 5 folds due to the cost of the procedure –, then we construct the TFIDF matrix representation of the documents for the IS stage, and then, we use the corresponding raw document chosen as input for the Transformers classifiers.

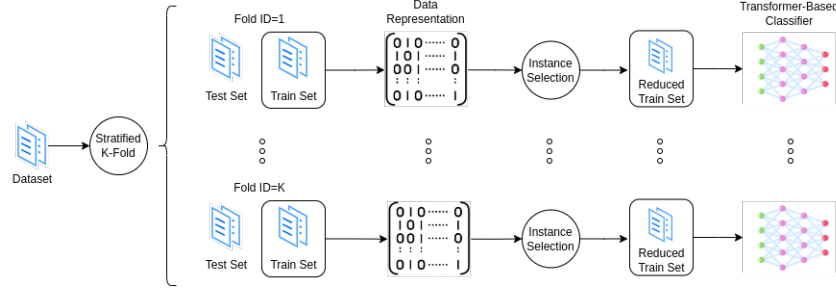


Figure 3. Data Representation and Preprocessing Procedure

Text Classification Methods As mentioned, our goal is to study and compare our proposed method against the SOTA IS techniques in the context of **Transformers** architectures – notably the SOTA in classification in several domains² [Gasparetto et al. 2022, de Andrade et al. 2024]. We compare the effectiveness among the latest version of the following Transformers – **RoBERTa**, **BERT**, **DistilBERT**, **BART**, **AlBERT**, and **XLNet**) – applied to all tested datasets.

Given the large number of hyperparameters to be tuned, performing a grid search with cross-validation is not feasible for all of them. As a result, to determine the optimum hyperparameter, we applied the methodology from [Cunha et al. 2021]. We fixed the initial learning rate as 5×10^{-5} , the max number of epochs as 20, and 5 epochs as patience. Finally, we perform a grid search on max_len (150 and 256) and batch_size (16 and 32) since these specified values directly impact efficiency and effectiveness.

Instance Selection Methods We consider in this chapter a set of 13 IS methods described in Section 4, namely: *Condensed Nearest Neighbor (CNN)*; *Edited Nearest Neighbor (ENN)*; *Iterative Case Filtering (ICF)*; *Instance Based 3 (IB3)*; *Decremental Reduction Optimization Procedure (Drop3)*; *Local Set-based Smoother (LSSm)*; *Local Set Border Selector (LSBo)*; *Local Density-based Instance Selection (LDIS)*; *Central Density-based Instance Selection (CDIS)*; *eXtended Local Density-based Instance Selection (XLDIS)*; *Prototype Selection based on Dense Spatial Partitions (PSDSP)*; *Enhanced Global Density-based Instance Selection (EGDIS)*; and *Curious Instance Selection (CIS)*. All parameters for the IS methods were defined with grid-search, using cross-validation in the training set during an initial empirical experiments round.

²LLMs such as GPT and LLama are built on top of Transformer architectures.

Metrics and Experimental Protocol All experiments were executed on an Intel Core i7-5820K with 6-Core and 12-Threads, running at 3.30GHz, 64Gb RAM, and a GeForce GTX TITAN X (12GB) and Ubuntu 19.04. Due to skewness in the datasets, we evaluated the classification effectiveness using Macro Averaged F1 (MacroF1)[Sokolova and Lapalme 2009]. We employed the paired t-test with a 95% confidence level to compare the average outcomes from our cross-validation experiments. Finally, we applied the Bonferroni correction [Hochberg 1988] to account for multiple tests. We consider reduction mean by defined as $\bar{R} = \frac{\sum_{i=0}^k \frac{|T_i| - |S_i|}{|T_i|}}{k}$, where T is the original training set, and S is the solution set containing the selected instances by the IS method being evaluated. Last, in order to analyze the cost-effectiveness tradeoff, we also evaluate each method’s cost in terms of the total time required to build the model. The Speedup is calculated as $S = \frac{T_{wo}}{T_w}$, where T_w is the total time spent on model construction using the IS approach, and T_{wo} is the total time spent on execution without the IS phase.

Experimental Results: Our findings reveal a significant potential of IS. In some cases, these methods can reduce the training set by up to 90% while maintaining their effectiveness. The IS approaches we studied achieved average reductions between 15.6% (LSSm) to 91.1% (XLDIS). However, despite the motivation for noise removal, selection methods were not able to improve the effectiveness of the text classification models. This underscores the need for further studies to investigate this issue, which could have profound implications for the field of data science and machine learning.

We demonstrated that three traditional IS methods (**LSSm**, **CNN**, **LSBo**) were able to reduce the total text classification models’ construction time while keeping the effectiveness in 12 (out of 19) considered datasets – with speedups between 1.04x (CNN - Books) and 5.69x (LSBo - Reuters90). In the other datasets, we observed that the introduction of IS approaches caused an overhead³ in terms of the total time to generate the model (running the IS methods + model construction), making the whole process more costly from a computational cost perspective. Overall, considering the three tripod constraints altogether and all datasets, the best traditional IS method was **CNN**.

Answering RQ1, our evaluation of the tripod constraints (reduction - efficiency - effectiveness), we showed that, in some datasets, specific selection methods can reduce the training set without loss of effectiveness and with efficiency improvements. **However, our experiments revealed that no IS method can respect all restrictions in all cases.** Also, in some situations, the application of the IS methods can incur additional overheads in the time to construct the model. Our results present a partial answer to the posed question regarding the need for large training sets for performing fine-tuning. In some cases, it is possible to use IS methods to reduce the training set without loss of effectiveness and efficiency gains – meaning that we do not always need a lot of data. Therefore, these findings neither totally support nor completely refute our posed hypothesis that traditional IS methods are capable of simultaneously respecting all posed restrictions.

In Table 2, we provide pros (**P**), cons(**C**) and recommendations (**R**) for the six best IS methods. The remaining seven methods were ineffective in the ATC context, producing statistically worse results in 116 out of 133 (7 methods x 19 datasets). Particularly our focus is on NLP tasks, especially ATC ones.

³We fully explore this time overhead behaviour in our Dissertation (Chapter 3.3.3)

CNN	(P): Best overall method, achieving the best trade-off on the tripod Reduction-Effectiveness-Efficiency (C): Limited effectiveness when applied to the largest datasets (R): Medium-to-small topic-related data or sentiment analysis tasks
LSSm	(P): Achieved the best overall Effectiveness (C): Lowest overall Efficiency when considering the three best IS approaches (R): General tasks that cannot deal with effectiveness losses
LSBo	(P): Considering the three best IS methods in terms of effectiveness, achieved the best reduction rate (C): Limited results considering effectiveness when applied to the topic-related tasks (R): Sentiment analysis tasks that need performance and scalability
EGDIS	(P): Considering the IS methods in this list, achieved the best reduction rate and speed up trade-off (C): Limited results considering effectiveness when applied to the medium-to-large topic-related task (R): Tasks that need performance and scalability and may afford some (up to 5%) effectiveness losses
CIS	(P): Good effectiveness results in sentiment analysis tasks (C): Lowest overall Efficiency (R): Not recommended for large NLP tasks due to high associated computational cost
IB3	(P): 4 th best IS method considering effectiveness (C): Limited results considering effectiveness when applied to large topic-related tasks (R): Medium-to-small topic-related data or sentiment tasks that can deal with a limited effectiveness loss

Tabela 2. Instance Selection Pros (P), Cons (C) and Recommendations (R).

Our results motivate further investigations on exploiting IS methods in the ATC context, especially regarding new transformers. Our study concerning RQ1 also opens space for designing new, more efficient, effective, and scalable IS methods for the current ATC and the big data scenarios in general. To help close this gap, in the next chapter, we introduce the E2SC framework, which is a new two-step framework that satisfies all the constraints of the tripod and can be used in real-world situations, even with datasets containing thousands of instances, with a special focus on transformer-based architectures.

Final Remark: Both sections 4 and 5 played a crucial role in producing an impactful publication in the *ACM Computing Surveys (CSUR)* [Cunha et al. 2023c].

6. An Effective, Efficient, and Scalable Confidence-Based Instance Selection Framework for Transformer-Based Text Classification

In the previous section, we found that no one method was able to meet all restrictions in all cases. Additionally, in some scenarios, using these methods increased the time it takes to construct the model, which corresponds to a gap in the literature on methods capable of respecting the posed restriction simultaneously.




Figure 4. The proposed E2SC Framework.

To help close this gap, the main contribution of this section is the proposal of **E2SC** – **E**ffective, **E**fficient, and **S**calable **C**onfidence-based **I**nstance **S**election – a novel two-step framework⁴ aimed at large datasets with a special focus on transformer-based

⁴To guarantee the reproducibility of our solutions, all the code, the documentation of how to run it and datasets are available on: <https://github.com/waashk/>

architectures, our first redundancy-oriented IS solution. E2SC is a technique that satisfies the tripod’s constraints and is applicable in real-world scenarios, including datasets with thousands of instances. E2SC’s overall structure can be seen in Figure 4.

E2SC’s **first** step – Figure 4(a) – aims to assign a probability to each instance being removed from the training set (α parameters). We adopt an exact KNN model solution⁵ to estimate the probability of removing instances, as it is considered a **calibrated**⁶[Rajaraman et al. 2022] and computationally cheap (fast) model [Cardoso et al. 2017]. Our **first hypothesis (H1)** is that high confidence (if the model is calibrated to the correct class, known in training) positively correlates with redundancy for the sake of building a classification model. Accordingly, we keep the hard-to-classify instances (probably located in the decision border regions), weighted by confidence, for the next step, in which we partially remove only the easy ones.

As the **second** step of our method – Figure 4(b) – we propose to estimate a near-optimal reduction rate (β parameter) that does not degrade the deep model’s effectiveness by employing a validation set and a weak but fast classifier. Our **second hypothesis (H2)** is that we can estimate the effectiveness behavior of a robust model (deep learning) through the analysis and variation of selection rates in a weaker model. Again, we explore KNN for this. More specifically, we introduce an iterative method that statistically compares, using the validation set, the KNN model’s effectiveness without any data reduction against the model with iterative data reduction rates. In this way, we can estimate a reduction rate that does not affect the KNN model’s effectiveness. Last, considering the output of these two steps together (Figure 4(c)), $\beta\%$ instances are randomly sampled, weighted by the α distribution, to be removed from the training set.

Experimental Results:

We compare E2SC proposal with **six** robust state-of-the-art instance selection baseline methods considering as input of the best of **seven** deep learning text classification methods in a large benchmark with **19** datasets. Our experimental evaluation shows that **E2SC** managed to significantly reduce the training sets (by **27%** on average; varying between 10% and 60% of reduction) while maintaining the same levels of effectiveness in **18** (out of 19) considered datasets. Also, we found that **E2SC** was able to reduce the total ATC models’ construction time while keeping the effectiveness in all (19) considered datasets – with speedups of **1.25** on average, varying between 1.02x (Books) and 2.04x (yelp_reviews). Overall, considering the three tripod constraints altogether and all datasets, the best IS method so far was our first proposed framework. Finally, to demonstrate the flexibility of our framework to cope with large datasets, we propose two modifications. Our enhanced solution managed to increase the reduction rate of the training sets (to **29%** on average) while maintaining the same levels of effectiveness in **all** datasets, with speedups of **1.37** on average. The framework scaled to large datasets, reducing them by up to 40% while statistically maintaining the same effectiveness with speedups of **1.70x**.

⁵We depart from the premise that the exact KNN model solution is a reasonable proxy for redundancy. We test and confirm this premise in our Dissertation.

⁶A calibrated classifier is one whose probability class predictions correlate well with its accuracy, e.g., for those instances predicted with 80% of confidence the classifier is correct in roughly 80% of the cases.

Despite being innovative and achieving significant results in terms of effectiveness, efficiency, and reduction, the E2SC framework focused only on **redundancy**, leaving some other aspects that may help to reduce training further untouched. One such aspect is **noise**, here defined as instances incorrectly labeled by humans in the dataset [Martins et al. 2021] as well as (possible) outliers that do not contribute (or even get in the way) to model learning (tuning). Indeed, according to [Martins et al. 2021], users, whether regular individuals or experts, make a reasonable amount of mistakes while labeling complex instances – between 56% and 64% of the time. Other few noisy instances (a.k.a., outliers) may be correctly labeled, but they differ so significantly from other instances from the same class that they are either useless for the sake of learning (tuning) or may even be detrimental to the process.

Noisy instances can potentially constitute a significant portion of available data in these contexts. Noisy (training) instances may not only degrade the model’s effectiveness by incorporating misleading patterns in the model but may also be detrimental to performance as they need to be processed to extract and incorporate these patterns into the model. If the amount of noise is significant, there will certainly be negative impacts on effectiveness and efficiency. However, in a simulated scenario designed to evaluate the capability of the IS baseline methods and our previous solution to remove noise, none of the IS solutions satisfactorily performed the task.

Final Remark: This section was summarized in a conference paper in the *ACM SIGIR Conference on Research and Development in Information Retrieval* [Cunha et al. 2023b].

7. An Extended Noise-Oriented and Redundancy-Aware Instance Selection Framework for Transformer-Based Automatic Text Classification

The main contribution of this section is the proposal of an *extended bi-objective instance selection (biO-IS)* framework built upon our first one aimed at removing both redundant and noisy instances simultaneously.

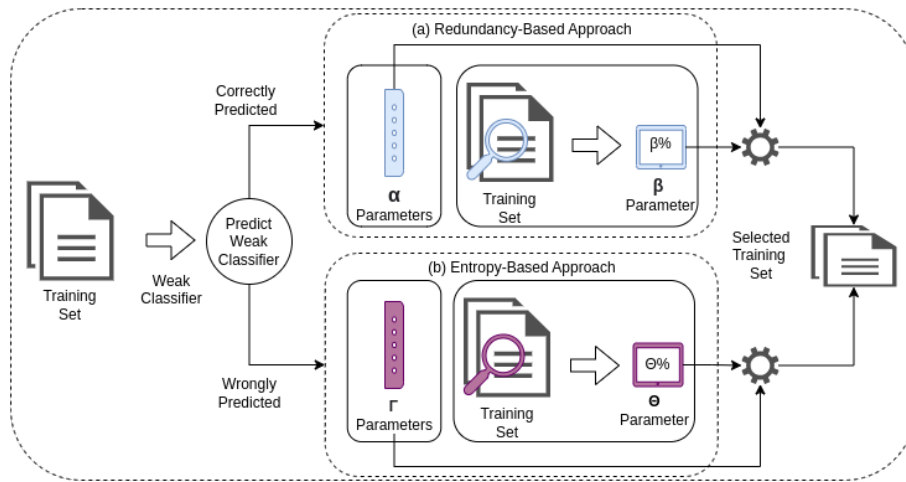


Figure 5. Bi-objective Instance Selection Framework

The main contribution of this section is the proposal of an *extended bi-objective instance selection (biO-IS)* framework built upon our first one aimed at removing both

redundant and noisy instances simultaneously. As depicted in Figure 5, our extended framework encompasses three main components: a weak classifier, a redundancy-based approach, and an entropy-based approach. In Figure 5 (in blue), we depart from our original solution proposed in the previous section considering the Logistic Regression as the calibrated weak classifier instead of KNN, as in our original work. An in-depth comparative analysis of several possibilities of the weak classifiers (see more in the Dissertation), including Decision Tree (DT), Logistic Regression (LR), XGBoost, LightGBM (LGBM), and Linear SVM, demonstrated LR as the best option in terms of a trade-off effectiveness-calibration-cost.

To address the second objective of noise removal (the lower part of Figure 5 – in purple), we propose a new step based on the entropy and a novel iterative process to estimate near-optimum reduction rates. Considering the instances wrongly predicted by the weak classifier, the main objective is to assign a probability to each of them being removed from the training set based on the probability of the instance being noise. For this, we propose using the entropy function as a proxy to determine the reduction behavior caused by these instances for the sake of training an ATC model. The intuition behind this new step is that when the prediction provided by the calibrated weak classifier is incorrect, the entropy of the posterior probabilities negatively correlates with the classifier’s confidence. This means that low entropy occurs when the classifier assigns an instance with absolute certainty to a wrong class, while high entropy occurs when the classifier is uncertain among several classes. Therefore, we consider the chance of a noisy instance being removed by the inverse of the entropy of the prediction, so when an incorrect prediction is accompanied by low entropy, it is more likely to be removed, and, otherwise, it is more likely to be kept. Accordingly, the proposed biO-IS framework provides a comprehensive solution to address redundancy and noise removal simultaneously.

Experimental Results:

We compare **biO-IS** with **seven** robust SOTA IS baseline methods, including our first proposal, E2SC, in the text classification domain considering the same benchmark covering **22** datasets. Our experimental evaluation reveals that, in a simulated scenario designed to evaluate the capability of the IS baseline methods and our previous solution to remove noise, none of the IS solutions were capable of satisfactorily performing the task. On the other hand, **biO-IS** managed to remove up to 66.6% of the manually inserted noise. Moreover, **biO-IS** managed to significantly reduce the training sets (by **40.1%** on average; varying between 29% and 60% of reduction) while maintaining the same levels of effectiveness in **all** of the considered datasets. Also, **biO-IS** managed to consistently provide speed-ups of **1.67x** on average (maximum of **2.46x**). No baseline, not even our previous SOTA solution, was capable of achieving results with this level of quality, considering all tripod criteria. Indeed, the only other method capable of maintaining the effectiveness on all datasets was E2SC; **biO-IS** improves over E2SC by 41% regarding reduction rate and from 1.42 to 1.67 (on average) regarding speedup, achieving the SOTA in the Instance Selection field.

Final Remark: This section was summarized in a journal paper in the *ACM Transactions on Information Systems* [Cunha et al. 2024], a leading journal in **Information Systems**.

8. Conclusion and Future Work

This dissertation surveyed classical and recent IS approaches, revealing significant advances but limited application scope. Most traditional methods target small tabular datasets, with rare applications in NLP despite its potential benefits.

To address this gap, we conducted a rigorous comparative study of classical and SOTA IS methods applied to ATC. The study evaluated the trade-offs among reduction, effectiveness, and cost, motivated by the rising costs of new ATC solutions due to deep learning embeddings and increasing data volumes. The main findings from over 5,000 experiments include: (i) IS methods rarely improved ATC model performance, with a few exceptions like XLNet. (ii) Among 13 IS methods, LSSm, CNN, and LSBo performed best, achieving significant reductions (46.6% average) while maintaining effectiveness in 12 datasets. (iii) Fine-tuning transformers proved critical for effectiveness, and IS methods helped reduce training costs in many cases (161 out of 190). Contrary to common beliefs, neural networks often require representative—not large—data to perform well in ATC. Overall, IS techniques effectively reduce training set sizes without compromising effectiveness, especially in fine-tuning transformer models. However, traditional approaches often fall short of meeting all trade-offs, underscoring the need for more efficient, scalable IS techniques in big data scenarios.

To address these challenges, we proposed E2SC, a redundancy-oriented two-step IS framework designed for large datasets and transformer-based architectures, introducing: (i) Calibrated weak classifiers to estimate data usefulness during transformer training. (ii) Iterative processes and heuristics to determine optimal reduction rates. E2SC reduced training sets by 30% on average while maintaining effectiveness across 22 datasets, achieving speedups of up to 70%.

To overcome E2SC limitations regarding handling noise effectively, we developed biO-IS, an extended framework that simultaneously removes redundant and noisy instances. Experimental results showed biO-IS reduced training sets by 41% on average (up to 60%), removed 66.6% of manually inserted noise, and maintained effectiveness across datasets. It provided consistent speedups (1.67x average, up to 2.46x), outperforming E2SC in reduction and efficiency.

These findings confirm the dissertation hypothesis: small and large language models can be trained with less data without sacrificing effectiveness. This not only enables significant cost and energy savings but also contributes to reducing carbon emissions. Such promising results instill hope for a more sustainable and efficient future in NLP, where advancements in IS techniques can lead to substantial environmental and economic benefits.

By proposing and evaluating the proposed enhancements (modifications and extensions) that resulted in the proposed *biO-IS* framework, our work opens several new opportunities. In future work, we intend to evaluate our proposed framework on other models, such as T5 and DeBERTa, assessing its generalizability. We also plan to apply IS concepts in other fields beyond ATC, such as searching, ranking, recommendation, and other NLP tasks, such as question answering and supervised topic modeling. So far, we have measured the importance of a whole document for the training (tuning) process, but in the future, we could do this at a more fine-grained level, such as for reviews'

aspects [Luiz et al. 2018] and documents’ passages. We also plan to investigate using our framework in an unlabeled deep-learning pre-training stage, e.g., to build a large language model from scratch more efficiently.

We also want to investigate the introduction of IS techniques in the context of AutoML solutions as a step in a pipeline of transformations as in [Cunha et al. 2020]. Furthermore, considering the vast number of dimensions in textual datasets, assessing how *feature selection* methods interact with *instance selection* can be very interesting. Essentially, this involves working with both the rows (documents) and the columns (features) of a document-feature matrix representing textual data. Last but not least, we intend to perform a cost-benefit analysis on LLMs (effectiveness vs. cost) to potentially pursue IS for fine-tuning LLMs, also exploiting new scalable computational paradigms such as quantum computing [Ferrari Dacrema et al. 2024].

Acknowledgements This work was partially supported by CNPq, CAPES, FAPEMIG, AWS, Google, NVIDIA, CIIA-Saúde, and FAPESP.

Referências

- Andrade, C., Cunha, W., Fonseca, G., Pagano, A., Santos, L., Pagano, A., Rocha, L., and Gonçalves, M. (2024). Explaining the hardest errors of contextual embedding based classifiers. In *Proceedings of the 28th Conference on Computational Natural Language Learning*.
- Barigou, F. (2018). Impact of instance selection on knn-based text categorization. *Journal of Information Processing Systems*, 14(2).
- Carbonera, J. L. and Abel, M. (2018). Efficient instance selection based on spatial abstraction. In *2018 IEEE 30th International Conference on Tools with Artificial Intelligence (ICTAI)*.
- Cardoso, T. N., Silva, R. M., Canuto, S., Moro, M. M., and Gonçalves, M. A. (2017). Ranked batch-mode active learning. *Information Sciences*, 379:313–337.
- Cunha, W., Canuto, S., Viegas, F., Salles, T., Gomes, C., Mangaravite, V., Resende, E., Rosa, T., Gonçalves, M. A., and Rocha, L. (2020). Extended pre-processing pipeline for text classification: On the role of meta-feature representations, sparsification and selective sampling. *IP&M*, 57(4):102263.
- Cunha, W., França, C., Rocha, L., and Gonçalves, M. A. (2023a). Tpdrr: A novel two-step transformer-based product and class description match and retrieval method. *arXiv preprint arXiv:2310.03491*.
- Cunha, W., França, C., Fonseca, G., Rocha, L., and Gonçalves, M. A. (2023b). An effective, efficient, and scalable confidence-based instance selection framework for transformer-based text classification. In *ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’23*.
- Cunha, W., Mangaravite, V., Gomes, C., Canuto, S., Resende, E., Nascimento, C., Viegas, F., França, C., Martins, W. S., Almeida, J. M., Rosa, T., Rocha, L., and Gonçalves, M. A. (2021). On the cost-effectiveness of neural and non-neural approaches and representations for text classification: A comprehensive comparative study. *Information Processing & Management*.
- Cunha, W., Moreo, A., Esuli, A., Sebastiani, F., Rocha, L., and Gonçalves, M. (2024). A noise-oriented and redundancy-aware instance selection framework. *ACM Transactions on Information Systems*.
- Cunha, W., Viegas, F., França, C., Rosa, T., Rocha, L., and Gonçalves, M. A. (2023c). A comparative survey of instance selection methods applied to nonneural and transformer-based text classification. *ACM Comput. Surv.*
- de Andrade, C., Cunha, W., Reis, D., Pagano, A. S., Rocha, L., and Gonçalves, M. A. (2024). A strategy to combine 1stgen transformers and open llms for automatic text classification. *arXiv preprint arXiv:2408.09629*.

- de Andrade, C. M., Belém, F. M., Cunha, W., França, C., Viegas, F., Rocha, L., and Gonçalves, M. A. (2023). On the class separability of contextual embeddings representations – or “the classifier does not matter when the (text) representation is so good!”. *IP&M*, 60(4):103336.
- Ferrari Dacrema, M., Pasin, A., Cremonesi, P., and Ferro, N. (2024). Quantum computing for information retrieval and recommender systems. In *European Conference on Information Retrieval*, pages 358–362.
- Garcia, S., Derrac, J., Cano, J., and Herrera, F. (2012). Prototype selection for nearest neighbor classification: Taxonomy and empirical study. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Gasparetto, A., Marcuzzo, M., Zangari, A., and Albarelli, A. (2022). A survey on text classification algorithms: From text to predictions. *Inf.*, 13:83.
- Hart, P. (1968). The condensed nearest neighbor rule (corresp.). *IEEE transactions on information theory*, 14(3):515–516.
- Hochberg, Y. (1988). A sharper bonferroni procedure for multiple tests of significance. *Biometrika*, 75(4).
- Joulin, A., Grave, E., Bojanowski, P., and Mikolov, T. (2017). Bag of tricks for efficient text classification. In *Proceedings of the Conference European Chapter Association Computational Linguistics (EACL)*.
- Leyva, E., González, A., and Pérez, R. (2015). Three new instance selection methods based on local sets: A comparative study with several approaches from a bi-objective perspective. *Pattern Recognition*.
- Li, Q., Peng, H., Li, J., Xia, C., Yang, R., Sun, L., Yu, P. S., and He, L. (2022). A survey on text classification: From traditional to deep learning. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 13(2):1–41.
- Lu, X., Duan, X., Mao, X., Li, Y., and Zhang, X. (2017). Feature extraction and fusion using deep convolutional neural networks for face detection. *Mathematical Problems in Engineering*, 2017.
- Luiz, W., Viegas, F., Alencar, R., Mourão, F., Salles, T., Carvalho, D., Gonçalves, M. A., and Rocha, L. (2018). A feature-oriented sentiment rating for mobile app reviews. In *Proceedings of the WebConf’18*.
- Martins, K., Vaz de Melo, P., and Santos, R. (2021). Why do document-level polarity classifiers fail? In *Proceedings of the 2021 Conference of the NAACL: Human Language Technologies*.
- Rajaraman, S., Ganesan, P., and Antani, S. (2022). Deep learning model calibration for improving performance in class-imbalanced medical image classification tasks. *PloS one*.
- Roy, A. and Cambria, E. (2022). Soft labeling constraint for generalizing from sentiments in single domain. *Knowledge-Based Systems*, 245:108346.
- Sindagi, V. A., Yasarla, R., Babu, D. S., Babu, R. V., and Patel, V. M. (2020). Learning to count in the crowd from limited labeled data. In *Computer Vision – ECCV*, pages 212–229, Cham.
- Sokolova, M. and Lapalme, G. (2009). A systematic analysis of performance measures for classification tasks. *Information Processing & Management (IP&M)*, 45(4):427–437.
- Tsai, C.-F., Chen, Z.-Y., and Ke, S.-W. (2014). Evolutionary instance selection for text classification. *J. Syst. Softw.*, 90(C):104–113.
- Uppaal, R., Hu, J., and Li, Y. (2023). Is fine-tuning needed? pre-trained language models are near perfect for out-of-domain detection. *arXiv preprint arXiv:2305.13282*.
- Wilson, D. L. (1972). Asymptotic properties of nearest neighbor rules using edited data. *IEEE Transactions on Systems, Man, and Cybernetics*, pages 408–421.
- Wilson, D. R. and Martinez, T. R. (2000). Reduction techniques for instance-based learning algorithms. *Machine learning*, 38(3):257–286.