Short text classification of social groups focused on courses

Antonio Leandro Martins Candido[1,2]
Jose Everardo Bessa Maia[1]

[1]Universidade Estadual do Ceará - UECE, 60714-903, Fortaleza-CE, Brazil

[2]Instituto Federal do Ceará - IFCE, 61609-090, Caucaia-CE, Brazil

leandro.candido@aluno.uece.br, jose.maia@uece.br

Abstract. *The ubiquitous possession of communication devices enhances a Social IoT whose most visible face is Social Networks. Social networks, also called relationship networks, are an important source for collecting unbiased feelings or mental states of groups of students in relation to a course in which they participate. These groups in social networks, often isolated and without the knowledge of the school administration, function as a discussion forum and the analysis of the texts posted there can enhance educational management with preventive measures to mitigate factors of course failure. This work develops the classification of sentiment in texts of groups of students with a view to detecting the need for intervention by educational management. The accuracy of the method is reinforced by a domain adaptation technique to take advantage of labeled data from other related domains. The results obtained are compatible with the state of the art of sentiment classification in other application domains.*

1. Introduction

Unsupervised Domain Adaptation (UDA) has emerged as a promising approach within machine learning and natural language processing, particularly in contexts where labeled data is scarce or expensive to obtain. It is especially relevant for text classification tasks that require the transfer of knowledge from well-labeled source domains to underrepresented or low-resource target domains. This study investigates the application of two widely recognized UDA techniques, DeepCORAL [Sun and Saenko 2016] and Domain-Adversarial Neural Networks (DANN) [Ganin et al. 2016] to enhance the performance of binary label classification in partially labeled target datasets derived from the Stanford MOOC platform. The source datasets employed include sentiment-labeled corpora from Amazon, Chat, IMDb, and Yelp, chosen for their diverse linguistic features and domain characteristics.

The central hypothesis underpinning this work is that sentiment-related knowledge extracted from rich source domains can be effectively transferred to target educational domains, thereby improving the accuracy of binary sentiment label classification even in the absence of labeled examples in the target domain. Such sentiment classification is particularly valuable in educational contexts, where student feedback on social platforms may contain latent indicators of engagement, frustration, or confusion.

To benchmark model performance, the study contrasts intra-domain classification, where training and testing are conducted within the same domain and typically yield the best-case performance, with cross-domain classification scenarios, which illustrate the generalization gap encountered when models are applied to unseen domains without adaptation. UDA techniques such as DeepCORAL and DANN aim to mitigate this performance drop by aligning feature distributions between source and target domains, thereby enabling models to generalize more effectively in the absence of target-domain labels.

Textual data were vectorized using both traditional TF-IDF representations and contextual embeddings generated by BERT, allowing for comparative analysis of embedding effectiveness in the context of domain transfer. Experimental results indicate that DeepCORAL consistently outperformed DANN in terms of both accuracy and stability, particularly in sentiment-based tasks. DeepCORAL's alignment of second-order statistics appears to offer a more robust mechanism for domain adaptation in cases where source and target domains exhibit moderate similarity. In contrast, DANN, which relies on adversarial training to minimize domain discrepancy, demonstrated greater sensitivity to domain shifts, resulting in more variable performance.

These findings underscore the practical viability of UDA for text classification tasks, particularly in domains such as education where labeled data is often limited. The use of UDA not only improves classification accuracy but also facilitates model generalization, a key requirement for real-world deployments across heterogeneous domains.

In parallel, this research aligns with the growing significance of the Social Internet of Things (Social IoT) in educational environments. The widespread interconnectivity of devices and users has enabled the passive collection of large volumes of textual data from student interactions on social networks. These interactions, while often informal and unmoderated, contain valuable insights into students' emotional and cognitive states. By integrating advanced natural language processing methods with UDA techniques, this study demonstrates how it is possible to automatically detect early signs of dissatisfaction, confusion, or disengagement within student communities. Such insights can empower educational administrators to implement timely and proactive interventions, thereby improving learning outcomes, increasing student retention, and reducing failure rates.

In sum, this work contributes to the intersection of domain adaptation, sentiment analysis, and educational data mining, offering practical and methodological insights for future research and applications in low-resource, real-world educational settings.

In the following sections, related work is presented, highlighting relevant approaches in text classification and domain adaptation. Then, the methods used, including the DeepCORAL and DANN domain adaptation techniques, and the experiments performed are detailed. Subsequently, the results obtained are discussed, comparing the performance of the proposed approaches with existing techniques. Finally, the conclusions of the study are presented, along with suggestions for future work.

2. Related works

Recent research has increasingly emphasized the multifaceted challenges associated with text classification and domain adaptation, particularly in the context of short text data prevalent in the Internet of Things (IoT) and social network environments. These challenges stem primarily from high variability in linguistic context, sparse or noisy data, and frequent domain shifts, which significantly complicate the development of robust and generalizable models [Lima and Maia 2018]. Short texts, such as tweets, status updates, or IoT-generated messages, often lack sufficient semantic cues, making it difficult to infer sentiment, intent, or topic without contextual enrichment.

To address these issues, researchers such as Kumar and Zhang et al. have explored the use of Hierarchical Attention Networks (HAN), which offer promising performance for both sentiment and document classification tasks. These architectures are particularly effective in capturing complex linguistic features such as sarcasm, negation, and contextual dependencies, which are often present in informal and user-generated texts from social media platforms. By modeling textual data at both the word and sentence levels, HANs can extract semantically rich representations, enabling finer-grained classification even in the presence of noisy or ambiguous inputs [Kumar 2022, Zhang and Rao 2020].

In parallel, Li et al. and Dogra et al. have tackled the scarcity of labeled data in IoT applications by integrating semi-supervised learning approaches and domain-specific natural language processing (NLP) architectures. These strategies have proven effective in tasks such as email filtering and real-time message classification, enhancing not only accuracy but also security and operational efficiency in distributed IoT systems [Dogra et al. 2022]. Their work underscores the growing importance of label-efficient methods that can function in environments where supervised learning is impractical due to privacy concerns or annotation costs.

Complementary efforts by Meng et al. and Luo et al. demonstrate how convolutional neural networks (CNNs) and attention mechanisms can be adapted for specific applications such as sentiment analysis and query classification in smart tourism environments. These models leverage both local and global textual features to understand user intent, personalize responses, and optimize user experience in real-time systems [Meng 2024, Luo et al. 2022]. Their results suggest that domain-specific model tuning and attention-based architectures are critical for achieving scalability and contextual understanding in specialized IoT sectors.

In a broader systems context, fog computing architectures have also been explored to support efficient NLP processing. Mohammed et al. propose data replication strategies within fog nodes to minimize latency and reduce synchronization costs across the network, improving the feasibility of deploying NLP models at the edge [Mohammed et al. 2021]. Meanwhile, Qiang et al. address the need for computational efficiency by developing compact recurrent neural network (RNN) models tailored for real-time, resource-constrained IoT environments. These lightweight models maintain classification accuracy while meeting the latency and power requirements of embedded devices [Qiang et al. 2022].

For industrial IoT (IIoT) applications, Wang demonstrates how enhanced Random Forest classifiers can be employed to improve classification reliability on structured sensor and maintenance data, particularly in environments with evolving data streams and concept drift [Wang 2024]. In a complementary approach, Antunes et al. utilize semantic feature extraction to boost classification accuracy and resilience, showing that integrating domain knowledge with machine learning can yield more interpretable and robust models in industrial settings [Antunes et al. 2018].

Further advancing the state of the art, Chen et al. explore zero-shot learning frameworks combined with knowledge graph embeddings to classify data in emergency scenarios where labeled examples are unavailable. These methods allow for on-the-fly generalization to novel categories, a critical requirement in dynamic and unpredictable social IoT contexts. Additionally, emotion-aware classification models proposed by Luo and Zhang apply unsupervised and semi-supervised techniques to detect sentiment and emotional tone in unlabeled social media content. Their results indicate that such hybrid approaches are highly effective for extracting actionable insights from large-scale, unlabeled textual data [Luo and Zhang 2022].

Taken together, these studies illustrate the dynamic intersection of IoT, natural language processing, and domain adaptation. They highlight a clear trajectory toward developing intelligent, adaptive, and resource-efficient systems capable of processing complex textual data across diverse and evolving domains.

3. Method

3.1. Deep CORAL

The Deep CORAL (Correlation Alignment for Deep Domain Adaptation) method is a widely adopted approach in the field of domain adaptation, particularly effective in scenarios where there exists a domain shift a distributional mismatch between the labeled training data (source domain) and the unlabeled or differently distributed test data (target domain). Domain shift is a common challenge in real-world machine learning applications, such as cross-domain sentiment analysis, medical diagnosis across populations, or object recognition across varying imaging conditions. Deep CORAL addresses this by learning domain-invariant representations, making it possible to transfer learned knowledge from one domain to another with minimal performance degradation [Sun and Saenko 2016].

3.1.1. Covariance Alignment

Deep CORAL extends the classic CORAL algorithm by integrating it directly into deep learning models. Rather than requiring handcrafted feature engineering, Deep CORAL leverages the representational power of deep neural networks while simultaneously aligning the second-order statistics (i.e., covariance) of source and target domain features. This alignment is achieved by computing the covariance matrices of features extracted from both domains:

$$\mathbf{C} = \frac{1}{n-1}(\mathbf{F}-\mu)^T(\mathbf{F}-\mu) \tag{1}$$

In this formulation, $\mathbf{F}$ denotes the matrix of feature representations obtained from a deep network, and μ is the mean vector of the features. Specifically, $\mathbf{F_s}$ and $\mathbf{F_t}$ refer to the features extracted from the source and target domains, respectively, and μ_s, μ_t are their corresponding mean vectors. The number of samples in the source and target domains are denoted by n_s and n_t.

The core idea of Deep CORAL is to minimize the difference between the covariance matrices of the source and target features, effectively reducing domain mismatch. This is quantified using the Frobenius norm, leading to the CORAL loss:

$$\mathcal{L}_{\text{CORAL}} = \frac{1}{4d^2}|\mathbf{C}_s - \mathbf{C}_t|_F^2 \tag{2}$$

where d is the dimensionality of the feature representations, and $|\cdot|_F$ denotes the Frobenius norm. This loss encourages the deep model to learn representations whose statistical structure is similar across domains.

To incorporate domain adaptation into model training, the CORAL loss is combined with the standard classification loss (typically cross-entropy) into a joint objective:

$$\mathcal{L} = \mathcal{L}\text{class} + \lambda\mathcal{L}\text{CORAL} \tag{3}$$

Here, λ is a hyperparameter that balances the contribution of the classification loss and the CORAL domain adaptation loss during training. This composite objective allows the model to not only classify the labeled source data correctly but also generalize to the target domain by reducing domain discrepancy.

3.2. Adversarial Domain Adaptation with DANN

An alternative and influential approach to domain adaptation is the Domain-Adversarial Neural Network (DANN), which applies adversarial training to learn domain-invariant features. While Deep CORAL aligns statistical moments explicitly, DANN uses a domain classifier trained to distinguish between source and target domain samples, and a feature extractor trained to fool this domain classifier through a gradient reversal layer (GRL). The result is a minimax game in which the network learns features that are both discriminative for the main task and indistinguishable across domains [Ganin et al. 2016].

The high-level architecture of DANN is shown in Figure 1:

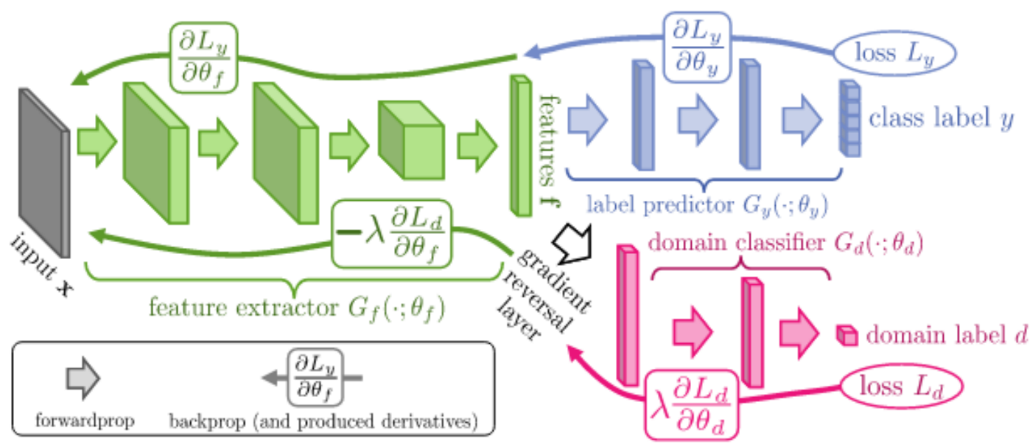

Figure 1. The DANN architecture (Ganin, et. al. 2016)

In DANN, the total loss function is composed of two parts:

- A label predictor loss (L_y) computed on labeled source domain samples.
- A domain classification loss (L_d) computed on both source and target samples.

The feature extractor is optimized to minimize L_y while maximizing L_d (via gradient reversal), ensuring the learned features are informative for classification but uninformative for domain discrimination. The overall training objective is:

$$\underbrace{E(\theta_f, \theta_y, \theta_d)}_{L_{total}} = \underbrace{\frac{1}{n}\sum_{i=1}^{n} L_y^i(\theta_f, \theta_y)}_{L_y} - \lambda \underbrace{\left(\frac{1}{n}\sum_{i=1}^{n} L_d^i(\theta_f, \theta_d) + \frac{1}{n'}\sum_{i=1}^{n'} L_d^i(\theta_f, \theta_d)\right)}_{L_d}. \quad (4)$$

In this expression:

- θ_f are the parameters of the feature extractor,
- θ_y are the parameters of the label classifier,
- θ_d are the parameters of the domain classifier,
- n and n' are the number of samples from the source and target domains, respectively,
- λ is a weighting factor controlling the trade-off between classification and domain confusion.

3.3. Comparative Insights

Both Deep CORAL and DANN have demonstrated strong performance across multiple domain adaptation benchmarks, including image classification, sentiment analysis, and other NLP tasks. While Deep CORAL is simpler to implement and more stable in training due to its non-adversarial nature, DANN is more expressive, capable of learning complex domain-invariant mappings in highly nonlinear feature spaces.

In practice, the choice between these methods depends on factors such as:

- The degree of similarity between domains (Deep CORAL is more effective when domains share structure),
- The availability of computational resources (DANN can be more demanding due to adversarial training),
- The sensitivity to hyperparameters and training stability.

When applied to tasks such as sentiment classification in cross-domain educational datasets, both techniques offer promising results by enabling models trained on richly annotated domains (e.g., Amazon reviews, IMDb, Yelp) to generalize to under-resourced target domains (e.g., MOOC discussion forums) without requiring new labels.

4. Experiments and discussion

4.1. Datasets

The purpose of this work is to investigate how Unsupervised Domain Adaptation (UDA) techniques, specifically DeepCORAL [Sun and Saenko 2016] and DANN [Ganin et al. 2016], can improve binary label classification on parcials target datasets of Stanford MOOC Dataset (Med5000, Hum3000 and Hum2000) using sentiment datasets

Table 1. Datasets

Dataset	negative Sentiment	positive Sentiment	Total
Amazon	500	500	1000
Chat	406	178	584
Hum2000	403	1738	2141
Hum3000	196	2833	3029
Imdb	362	386	748
Med500	1277	3904	5181
Yelp	500	500	1000

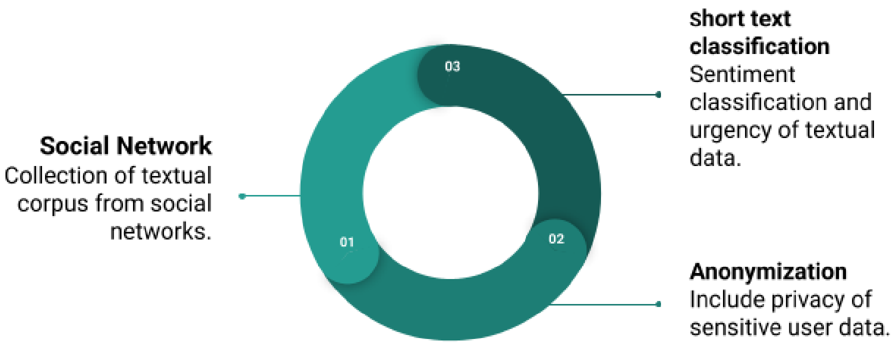

Figure 2. Short Text Classification Stages

as knowledge sources and data augmentation [Agrawal et al. 2015]. The data includes three target datasets, each with three binary labels, and four labeled source datasets for sentiment classification, evaluating whether sentiment information helps to improve classifiers on the target domains.

The Amazon, Chat, Imdb and Yelp sentiment datasets contain binary labels, as per the target datasets, but the chat dataset had textual labels: negative, neutral, and positive. This dataset was formatted for binary labels (negative = 0; neutral-positive = 1). The distribution of the samples is represented in the Table 1. Data privacy has been appropriately addressed, with measures implemented to ensure proper handling of sensitive information. While there is always room for continuous improvement, the current approach demonstrates a clear commitment to safeguarding data, the steps are shown in Figure 2.

4.2. Metrics

Evaluating predictive models is essential to validate their effectiveness. This work uses accuracy, F1-score, precision, and recall metrics in classification tasks. While accuracy reflects correct predictions, it may mislead in imbalanced datasets. The F1-score addresses this by balancing precision and recall, ensuring a comprehensive assessment of overall efficiency and performance in challenging scenarios [Lima and Maia 2018].

4.3. Results

The results obtained in this study demonstrate the effectiveness of unsupervised domain adaptation (UDA) techniques in classifying binary labels in short texts, using the Deep-

CORAL and DANN approaches. The experiments considered vector representations based on TF-IDF and BERT embeddings, allowing a comparative analysis of performances.

In the intra-domain classification scenario, the BERT representation technique combined with SVM (Table 3) outperformed TF-IDF (Table 2) on most datasets, achieving, for example, an accuracy of 94% on the Chat dataset and an F1-score of 95% on the Hum3000 dataset. BERT's robustness is particularly evident in scenarios with greater semantic complexity, reflecting its ability to capture contextual nuances of texts.

Table 2. Intradomain classification with TF-IDF and SVM

Dataset	Acc	F1	Precision	Recall
Amazon	0.79	0.78	0.82	0.74
Chat	0.90	0.79	0.95	0.68
Hum2000	0.82	0.90	0.83	0.99
Hum3000	0.93	0.96	0.93	0.99
Imdb	0.76	0.78	0.74	0.84
Med5000	0.83	0.89	0.87	0.92
Yelp	0.78	0.78	0.81	0.75

The results in the tables show a significant drop in performance compared to the intradomain scenarios. For example, when training the model on the Amazon dataset and testing it on the Hum3000 dataset using TF-IDF, the accuracy was only 47%, with an F1-score of 56%. Similar results were observed for other pairs of domains, reflecting the difficulty of the models in generalizing to different domains without an alignment strategy between their distributions.

Table 3. Intradomain classification with BERT and SVM

Dataset	Acc	F1	Precision	Recall
Amazon	0.82	0.82	0.81	0.83
Chat	0.94	0.89	0.90	0.89
Hum2000	0.82	0.89	0.87	0.91
Hum3000	0.91	0.95	0.94	0.96
Imdb	0.84	0.85	0.84	0.85
Med5000	0.84	0.89	0.88	0.90
Yelp	0.83	0.83	0.84	0.83

The BERT embeddings showed slightly better performance in cross-domain scenarios compared to TF-IDF (Table 6), but still faced notable challenges. In the Amazon test for Hum3000, for example, the accuracy reached 58%, with an F1-score of 67%, values still below the desired value. This indicates that, although BERT captures the contextual features of the texts better, the divergence between the domains continues to be a significant barrier.

The results of DeepCORAL, presented in Table 4, show consistent performance, especially in datasets with greater similarity between source and target domains. For example, when using Hum3000 as the source domain and Hum2000 as the target domain, the accuracy reached 81%, with an F1-score of 90%. These values highlight the robustness of DeepCORAL in aligning feature distributions from different domains, ensuring greater generalization.

Table 4. UDA - DeepCORAL with BERT

(Source / Target) -> Val	Hum2000		Hum3000		Med5000	
	Acc	F1	Acc	F1	Acc	F1
Hum3000 / Hum2000	0.81	0.90	0.94	0.97	0.75	0.86
Hum3000+Amazon / Hum2000	0.81	0.90	0.94	0.97	0.75	0.86
Hum3000+Chat / Hum2000	0.81	0.90	0.94	0.97	0.75	0.86
Hum3000+Imdb / Hum2000	0.81	0.90	0.94	0.97	0.75	0.86
Hum3000+Yelp / Hum2000	0.81	0.90	0.94	0.97	0.75	0.86

The DANN approach, as shown in Table 5, also demonstrated relevant gains, but with greater sensitivity to semantic variations between domains. When comparing Hum3000 as the source and Hum2000 as the target, DANN's performance was similar to DeepCORAL's, with an accuracy of 81% and an F1-score of 90%. However, in scenarios involving more heterogeneous domains, such as Med5000, DANN's performance showed greater variation, highlighting its lower stability compared to DeepCORAL.

Table 5. UDA - DANN with BERT

(Source / Target) -> Val	Hum2000		Hum3000		Med5000	
	Acc	F1	Acc	F1	Acc	F1
Hum3000 / Hum2000	0.81	0.90	0.94	0.97	0.76	0.86
Hum3000+Amazon / Hum2000	0.79	0.87	0.94	0.97	0.73	0.84
Hum3000+Chat / Hum2000	0.81	0.90	0.94	0.97	0.75	0.86
Hum3000+Imdb / Hum2000	0.81	0.90	0.94	0.97	0.75	0.86
Hum3000+Yelp / Hum2000	0.81	0.90	0.94	0.97	0.75	0.86

Table 6. Crossdomain classification with BERT and SVM

Train / Test	Amazon		Chat		Hum2000		Hum3000		Imdb		Med5000		Yelp	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
Amazon	-	-	0.77	0.72	0.58	0.67	0.53	0.68	0.79	0.78	0.51	0.57	0.80	0.79
Chat	0.76	0.71	-	-	0.38	0.40	0.24	0.31	0.78	0.76	0.39	0.32	0.77	0.73
Hum2000	0.58	0.69	0.42	0.51	-	-	0.91	0.95	0.66	0.74	0.74	0.84	0.66	0.74
Hum3000	0.60	0.70	0.43	0.52	0.79	0.88	-	-	0.65	0.73	0.74	0.84	0.59	0.71
Imdb	0.79	0.77	0.78	0.71	0.70	0.80	0.32	0.44	-	-	0.48	0.49	0.78	0.79
Med5000	0.60	0.70	0.50	0.54	0.79	0.88	0.87	0.93	0.63	0.73	-	-	0.59	0.70
Yelp	0.79	0.77	0.77	0.71	0.37	0.37	0.40	0.54	0.74	0.71	0.49	0.50	-	-

The results in Table 7 show that DeepCORAL outperforms other cross-domain methods, achieving higher accuracy across the EDU, H&S, and MED datasets. Notably, its F1-scores for EDU (0.94) and MED (0.90) demonstrate its effectiveness in balancing accuracy and sensitivity, even with significant semantic variation. Compared approaches, like [Capuano 2021] and [Candido and Maia 2023], showed lower accuracies, underscoring the importance of advanced domain adaptation methods like DeepCORAL in addressing generalization challenges in short text classification.

Table 7. Comparing DeepCoral to other approaches

Approach	Dataset	EDU		H&S		MED	
		Acc	F1	Acc	F1	Acc	F1
DeepCoral	EDU	-	-	**0.94**	0.97	**0.81**	0.90
Capuano, 2021		-	-	0.74	-	0.74	-
Candido; Maia 2023		-	-	0.67	-	0.73	-
DeepCoral	H&S	**0.83**	0.91	-	-	**0.81**	0.90
Capuano, 2021		0.74	-	-	-	0.74	-
Candido; Maia 2023		0.84	-	-	-	0.70	-
DeepCoral	MED	**0.83**	0.91	**0.94**	0.97	-	-
Capuano, 2021		0.74	-	0.74	-	-	-
Candido; Maia 2023		0.74	-	0.72	-	-	-

4.4. Discussion

The results obtained in this study underscore the practical effectiveness of unsupervised domain adaptation (UDA) techniques, particularly DeepCORAL and Domain-Adversarial Neural Networks (DANN), in addressing the complex task of short text classification in cross-domain scenarios. Both methods demonstrated the ability to improve model generalization when transferring knowledge from labeled source domains to unlabeled or low-resource target domains. However, noteworthy differences were observed between the two techniques in terms of performance stability, robustness to domain divergence, and sensitivity to semantic variability.

Among the two approaches, DeepCORAL exhibited more consistent and stable results, particularly in experiments involving moderately similar domains, such as the Hum3000 and Hum2000 datasets. These domains share contextual and linguistic characteristics, which facilitates the statistical alignment of second-order feature representations the central mechanism of DeepCORAL. The method's robustness in such scenarios highlights the strength of covariance-based alignment for domains with overlapping feature spaces, enabling smoother knowledge transfer and more reliable classification outcomes across domain boundaries.

In contrast, DANN, while still effective, demonstrated greater variability in performance, especially when applied to more heterogeneous or semantically distant domains, such as the Med5000 dataset. These variations can be attributed to the adversarial nature of DANN's training process, which seeks to learn domain-invariant features through a minimax game between the feature extractor and a domain discriminator. While powerful, this approach tends to be more sensitive to semantic shifts, out-of-distribution features, and variations in textual length or vocabulary richness factors frequently encountered in real-world short text classification tasks.

Furthermore, the ablation studies and control experiments conducted without the use of domain adaptation revealed a substantial drop in classification performance, confirming the critical role of UDA in mitigating domain mismatch. In the absence of adaptation, both precision and recall scores deteriorated significantly, underscoring the inadequacy of naive cross-domain generalization. This finding reinforces the necessity of adopting alignment-based or adversarial adaptation strategies in practical deployments, particularly when target domain annotations are unavailable or costly to obtain.

Despite these promising results, the study also illuminated certain limitations of the evaluated methods. One recurring challenge was the presence of severe class imbalance, which adversely affected the classifier's ability to generalize, even when adaptation techniques were applied. Although metrics such as F1-score and recall remained satisfactory in balanced or moderately imbalanced settings, performance degradation was notable when the minority class constituted a small fraction of the dataset. This suggests that class rebalancing strategies, such as data augmentation, cost-sensitive learning, or synthetic oversampling (e.g., SMOTE), may need to be integrated with UDA to improve outcomes in highly skewed distributions.

In summary, the experimental findings confirm that both DeepCORAL and DANN significantly enhance the performance of short text classification under domain shift. However, the choice of technique should be guided by domain similarity, training stability requirements, and the semantic properties of the datasets. Additionally, the observed limitations point to important directions for future work, including the incorporation of hybrid adaptation models, attention-based mechanisms, and imbalance-aware training strategies to further improve the robustness and applicability of domain adaptation in complex real-world scenarios.

5. Conclusion

This study validates the potential of domain adaptation techniques to improve the classification of short texts in domains with limited labeled data. Among the approaches, DeepCORAL stood out for its robustness and consistency across different scenarios, being especially suitable for contexts where there is moderate similarity between domains. DANN, in turn, also obtained relevant results, but demonstrated greater sensitivity to semantic variations. The work has some limitations, such as the use of only one data set and its size.

The results reinforce the importance of text representations such as BERT for classification tasks, outperforming traditional techniques such as TF-IDF in terms of capturing more complex contexts. However, there are still challenges to be faced, particularly in scenarios with high disparity between domains or imbalanced classes.

Finally, the work contributes to the advancement of the field by providing insights into the conditions under which each technique is most effective, in addition to suggesting future paths to improve the generalization of models, such as the use of hybrid techniques or more advanced transfer learning approaches.

Acknowledgment

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

References

Agrawal, A., Venkatraman, J., Leonard, S., and Paepcke, A. (2015). Youedu: Addressing confusion in mooc discussion forums by recommending instructional video clips. In *Proceedings of the 8th International Conference on Educational Data Mining, EDM 2015*, pages 297–304.

Antunes, M., Gomes, D., and Aguiar, R. L. (2018). Towards iot data classification through semantic features. *Future Generation Computer Systems*, 86:792–798.

Candido, A. L. M. and Maia, J. E. B. (2023). Detecting urgent instructor intervention need in learning forums with a domain adaptation. In *Intelligent Systems Design and Applications*, pages 502–512. Springer Nature.

Capuano, N. (2021). Transfer learning techniques for cross-domain analysis of posts in massive educational forums. In *Intelligent Systems and Learning Data Analytics in Online Education*, pages 133–152. Elsevier.

Dogra, V., Verma, S., Kavita, Chatterjee, P., Shafi, J., Choi, J., and Ijaz, M. F. (2022). A complete process of text classification system using state-of-the-art nlp models. *Computational Intelligence and Neuroscience*, 2022(1):1883698.

Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. (2016). Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1):2096–2030.

Kumar, A. (2022). Contextual semantics using hierarchical attention network for sentiment classification in social internet-of-things. *Multimedia Tools and Applications*, 81(26):36967–36982.

Lima, J. M. C. and Maia, J. E. B. (2018). A topical word embeddings for text classification. In *Anais do XV Encontro Nacional de Inteligência Artificial e Computacional*, pages 25–35. SBC.

Luo, W. and Zhang, L. (2022). Question text classification method of tourism based on deep learning model. *Wireless communications and mobile computing*, 2022(1):4330701.

Luo, X., Yu, Z., Zhao, Z., Zhao, W., and Wang, J.-H. (2022). Effective short text classification via the fusion of hybrid features for iot social data. *Digital Communications and Networks*, 8(6):942–954.

Meng, L. (2024). The convolutional neural network text classification algorithm in the information management of smart tourism based on internet of things. *IEEE Access*.

Mohammed, A. H. K., Jebamikyous, H.-H., Nawara, D., and Kashef, R. (2021). Iot text analytics in smart education and beyond. *Journal of Computing in Higher Education*, 33(3):779–806.

Qiang, Y., Kumar, S. T. S., Brocanelli, M., and Zhu, D. (2022). Tiny rnn model with certified robustness for text classification. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.

Sun, B. and Saenko, K. (2016). Deep coral: Correlation alignment for deep domain adaptation. *European conference on computer vision*, pages 443–450.

Wang, Z. (2024). Construction of an internet of things industry chain classification model based on irfa and text analysis. *Journal of Information Processing Systems*, 20(2):215–225.

Zhang, Y. and Rao, Z. (2020). Hierarchical attention networks for grid text classification. In *2020 IEEE International Conference on Information Technology, Big Data and Artificial Intelligence (ICIBA)*, volume 1, pages 491–494. IEEE.