Transactions on Computer Science and Intelligent Systems Research
ISSN: 2960-1800, eISSN: 2960-2238 | Volume 9
AIDML 2025

Expanding Movie Sentiment Analysis with Fastai

Shunqian Mao *

College of Education Science, Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China

* Corresponding Author Email: a2945105434@outlook.com

Abstract. Sentiment analysis, a crucial component of natural language processing (NLP), isolates personal information from intellectual information, such as digital responses. Moving learning-based deep learning systems like Fastai have performed incredibly well in general text classification. However, they haven't been tested on their adaptability to domain-specific tasks like video sentiment analysis. Applying the Kaggle videos, this research examines Fastai's success. Using worksheet data to make sentiment labels implementing textual inputs and user ratings. The paper assesses the capabilities and limitations of Fastai's AWD- Classification structures through broad data preprocessing, type great-tuning, and several multifaceted testing situations. Experimental results reveal an accuracy of 94 % on standardized tests; however, they expose threats in addressing extended or quiet functions. The paper proposes employing impressive structures and data augmentation strategies to address these issues. The paper's results show Fastai's skills and areas for improvement and provide helpful recommendations for its effectiveness in practical programs.

Keywords: Sentiment Analysis; Fastai; Natural Language Processing; Deep Learning; Text Classification.

1. Introduction

Natural language processing (NLP) is a key artificial intelligence component, which makes innovations possible in everything from customer feedback to social media analytics. Companies can use sentiment analysis, a unit of NLP, to analyze the public's opinions of products, services, or multimedia content because it is necessary to process individual opinions from intellectual data [1]. Deep learning, primarily transfer learning, has revolutionized the field in contrast to traditional methods that rely on lexicon-based techniques or large machine learning concepts. Using large-scale historical and fine-setting techniques, pre-trained language types like ULMFiT and BERT have achieved state-of-the-art results [2, 3].

Despite these changes, persistent domain-specific issues persist. For instance, picture sentiment analysis requires certain options because of the inherent beauty of cinematic reviews, which frequently combine story reports, secret viewpoints, and symbolic language [4]. Existing information primarily focuses on standardized datasets like IMDb thoughts, which lack the heterogeneity of real-world data, such as varying expression lengths, sound, and relaxed speech [5]. This discrepancy highlights the need for domain-specific assessments to join effective options.

A simplified plan for NLP stuff is provided by Fastai, a high-level deep learning range built on PyTorch and featuring tools for transfer learning, automated preprocessing, and streamlined education pipelines. The effectiveness of internet filtering and feeling classification has been demonstrated [6]. However, its impact on picture sentiment analysis is also unfamiliar, particularly regarding its resilience to changing language measures and silent sources. For this study, a Fastai-based form was created using Kaggle videos to create a movie review style. Data from a calculator will be used to study how intellectual traits like simplicity, bravado, and noise affect classification accuracy.

Following a good AWD: Classification design setting within Fastai's framework, our methodology emphasizes thorough data preprocessing to tackle class imbalance and academic contradictions. To evaluate the performance of the principle under different situations, the paper create controlled experiments using a variety of diction cases. The results reveal Fastai's suitability for short-text classification, but they also address issues with how tough or broken sources are treated. These findings provide a framework for potential changes and support the discussion about increasing NLP equipment for domain-specific tasks.

2. Methodology

2.1. Fastai’s Framework for NLP

By creating a light API that makes it simpler to create designs using pre-built parts, Fastai reduces PyTorch's issues. The fastai. Drastically reducing the training fees, the term "package" provides options for tokenization, numericalization, and incorporating initialization. Pre-trained models like AWD-LSTM have a wealth of information in task-specific information, which enables robust performance and softly labelled samples, which are major advantages of their support in transfer learning. Also, Fastai's TextDataLoaders improve coaching productivity by automating evaluation systems and seats, while the Learner program uses innovative marketing methods, such as one-time understanding level planning [7].

2.2. Design for exploration

The Kaggle-starring films. CSV dataset, comprising 8 551 entries, was filtered to retain only the overview (textual input) and vote_average (sentiment label). Sentiment labels were binarized using a threshold of 6.5, with ratings ≥ 6.5 classified as positive (1) and <6.5 as negative (0). During training-test splits (80: 20 ratio), stratified sampling was used to reduce class imbalance. Short Texts: Truncated overviews (such as" A gripping thriller") were the product of three test categories.

Explains documents that are more than 200 thoughts long. Scripts with intentionally altered grammatical errors and useless claims. Fastai's fine-tune technique was used to create wonderful-music in the AWD- Classification fashion over five civilizations, with training metrics recorded for each generation. Grid search optimised hyperparameters like batch size (32) and dropout rate (0.5) for overfitting.

3. Results and Discussion

3.1. Dataset and Test Sample Selection

The films. Title, overview, and vote_average are some of the movie-related features in the computer selection. Data filtering: replacing lost rules and keeping only overview and vote_average spots are among the data preprocessing methods. The vote_average challenge is used to assign positive and negative sentiment labels. 20 % of the information is used for education and 20 % for tests. Additionally, three distinct types of test samples were created: Short Texts, which are succinct descriptions of movie content (such as" A thrilling adventure"). Longer texts contain total video outlines. Language vocabulary errors

As shown in table 1, in the first epoch, the model overmatched its peak validation accuracy (94.4%) with little overfitting despite decreasing training loss. This is in line with Merity et cetera's results [8]. Where AWD- Classification demonstrated rapid participation in text classification issues. However, accuracy decreased over time, indicating that longer education may have yielded a lower income.

Table 1. Training Metrics Across Epochs

Epoch	Train Loss	Valid Loss	Accuracy
0	0.2165	0.2366	94.43%
1	0.2084	0.2090	94.43%
2	0.1651	0.2364	93.75%
3	0.1082	0.2484	93.75%
4	0.0664	0.2474	93.65%

3.2. Analysis of a Test Incident

The model's 92.2 % accuracy demonstrated that it was appropriate for specific circumstances. The accuracy of longer texts dropped to 88.7 % due to repeated terminology, which reduced significant romantic signs. Noisy Texts' performance decreased to 82.1 %, indicating a lack of stylistic accuracy. These results corroborate research by Zhang et al., where character-level noise hurt convolutional networks [9]. Raphael and Charles suggest using a self attention system that prioritizes key currencies to improve resilience [3]. Additionally, noise-related issues can be resolved using data augmentation techniques like [10-12].

4. Discussion

AWD- LSTM's limited ability to promote sentiment-laden statements within extended options may be to blame for the lower accuracy of more texts. Regular LSTMs struggle to focus on significant currencies amid superior intellectual patterns, in contrast to Transformer-based designs, which use self-attention systems to capture long-range relationships and filter record information. Due to this challenge, hybrid architectures that combine Transformers ' social consciousness with Fastai's reduced training valves are required [2]. Incorporating BERT's political attention layers might make it easier for the woman to discern subtle romantic signals in long reviews. Additionally, consistency may be maintained across various data sets by using domain-specific measures like GLUE [5].

Quiet scriptures, spelling mistakes, and inconsistent vocabulary made classification challenges even more difficult. Because of misaligned semantic representations, this vibration interferes with tokenization and connecting techniques. This is in collection with Raffel and colleagues ' analysis [3]. stressed the value of effective preparation in reducing performance degradation brought on by noise. In future applications, advanced data augmentation methods like back-translation or word replacement may be used to increase training data collection while maintaining sentiment labels. Additionally, Zhang et cetera. On the character level, convolutional networks are suggested [6]. Might make subword designs more appealing and reduce the need for actual sign matches [8].

For this assessment, domain-specific changes that go beyond fundamental changes are required [9]. Because target story reports and subjective views are often combined in movie outlines, creating distinct artistic surroundings from standardized datasets like IMDb reviews [4]. Improve the output quality by changing how pipelines are cleaned to support genre-specific wording, symbolic language, and apparent language. For instance, pre-trained on movie-related corpora and domain-specific embeddings might better capture cultural nuances than universal language models [8]. These modifications align with larger NLP analysis changes, where transfer learning and functional structures continue redefining performance metrics [6].

Beyond scientific studies, these reports have practical uses. Businesses that rely on sentiment analysis, such as those in marketing, pleasure insights, and obtain systems, may apply Fastai's success to immediately strengthen it while adding improvements made to improve endurance. Ensemble methods, which combine Fastai and Transformer-based strategies to increase accuracy while reducing style imperfections, can mix quotes of various sorts [5]. Real-time feedback loops could make Standard form refinement possible, where marked tests are slowly added to teaching information [6].

5. Conclusion

This study effectively evaluates the effectiveness of Fastai's AWD- Classification kind in movie sentiment analysis, revealing both its strengths and limitations when applied to domain-specific academic data. The routine demonstrated robust performance in identifying distinct picture information, getting an accuracy of 94.43 % on standardized review designs. This victory aligns with previous studies demonstrating the usefulness of transfer learning systems for text classification work. However, when handling lengthy or loud texts, the model's accuracy dropped to 88.7 % and 82.1 %, respectively. These results echo Zhang et cetera's experiments. And identify the difficulties that academic stability and soundness present for real-world applications, who also discovered flaws in convolutional neural network-tainted tools.

Future studies do include a variety of options. A thorough understanding of sentiment could be gained from multimodal approaches that combine textual information with extra data, such as customer areas, temporal rating trends, or sensory information from movie trailers. Because international systems require types capable of analyzing a range of linguistic and cultural contexts, cross-linguistic and cross-cultural dexterity research is also required. Included sentiment analysis ' ethical implications, such as bias mitigation and transparency in model decisions, may need further investigation to guarantee fair and understandable results.

In conclusion, Fastai's AWD- Classification style is a promising beginning stage for video sentiment analysis, but its difficulties with handling long or difficult texts necessitate targeted optimizations. The tenacity and relevance of Fastai-based systems may be greatly improved by incorporating Transformer styles, improving cleaning pipelines, and adopting clothes strategies. These developments contrast with wider swings in NLP study, where extensive improvements and dynamic devices continue to push the boundaries of what can be done in sentiment analysis. Addressing these issues may be important in bridging the gap between theoretical progress and practical applications as the demand for appropriate and functional thinking grows.

References

[1] Howard J, & Ruder S. Universal language model fine-tuning for text classification. Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL), 2018, 328–339.

[2] Devlin J, Chang M W, Lee K, & Toutanova K. BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), 2019, 4171–4186.

[3] Raffel C, Shazeer N, Roberts A, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. Journal of Machine Learning Research, 2020, 21(140), 1–67.

[4] Pang B, & Lee L. Opinion mining and sentiment analysis. Foundations and Trends in Information Retrieval, 2008, 2(1–2), 1–135.

[5] Wang A, Singh A, Michael J, et al. GLUE: A multi-task benchmark and analysis platform for natural language understanding. Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2018, 353–355.

[6] Howard J, & Gugger S. Deep learning for coders with Fastai and PyTorch. O'Reilly Media, 2020.

[7] Ruder S, Peters M E, Swayamdipta S, & Wolf T. Transfer learning in natural language processing. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2019, 15–18.

[8] Merity S, Keskar N S, & Socher R. Regularizing and optimizing LSTM language models. arXiv preprint arXiv:1708.02182, 2017.

[9] Zhang X, Zhao J, & LeCun Y. Character-level convolutional networks for text classification. Advances in Neural Information Processing Systems (NeurIPS), 2015, 649–657.

[10] Joulin A, Grave É, Bojanowski P, & Mikolov T. Bag of tricks for efficient text classification. Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL), 2017, 427–431.

[11] Guo X P, Zhao X F, Li Y. Supervised Enhancement Network Incorporating Noise for Small Sample Data Enhancement. Electronic Measurement Technology, 2024, 47 (20): 109-116. DOI: 10.19651/j.cnki-emt.2416647.

[12] Li Z Y. Research on Chinese Named Entity Recognition Method Based on Multi Feature Fusion and Data Enhancement. Beijing University of Chemical Technology, 2024. DOI: 10.26939/d.cnki.g bhgu.2024.000292