Alkadhum Journal of Science, Vol. 1, No. 2 (2023)

Open Access

ISSN:3005-2386 (Online), ISSN: 3005-2378 (Print)



Alkadhum Journal of Science (AKJS)

Journal Homepage: https://alkadhum-col.edu.iq/JKCEAS



Beyond Polarity: The Potential Applications and Impacts of Sentiment Analysis and Emotion Detection

¹Murteza Hanoon Tuama*, ¹Wahhab Muslim Mashloosh, ²Hayder M. Albehadili, ¹Murtadha A. Alazzawi, ³Mahmood A. Al-shareeda

¹Department of Computer Techniques Engineering, Imam Al-Kadhum College (IKC), Iraq ²Missan Oil Company, Maysan, Iraq ³National Advanced IPv6 Centre (NAv6), Universiti Sains Malaysia (USM), Malaysia

Article information

Article history: Received: October, 19, 2023 Accepted: November, 20, 2023 Available online: December, 14, 2023

Keywords: Natural Language Processing, LSTM, RNN, Sentiment Analysis, Contextual Language, Emotion Detection, Deep learning

*Corresponding Author: Murteza Hanoon Tuama Murtezahanoon@gmail.com

DOI: https://doi.org/10.61710/akjs.v1i2.51

This article is licensed under: <u>Creative Commons Attribution 4.0</u> <u>International License</u>.

Abstract

Opinion mining and emotion detection are two important techniques in natural language processing that have gained significant attention in recent years. Opinion mining is the process of identifying and extracting subjective information from text, such as opinions, attitudes, and emotions, while emotion detection is the process of identifying and extracting emotions from text. These techniques have a wide range of applications in various domains, including social media analysis, customer feedback analysis, and product reviews. This paper provides an overview of opinion mining and emotion detection techniques in natural language processing. We discuss the various approaches and methods used in opinion mining and emotion detection, including machine learning, deep learning, and natural language processing techniques. We also explore the challenges and limitations of these techniques, including the subjectivity of language, cultural differences, and the lack of labeled data. Furthermore, we examine the current state of the art in opinion mining and emotion detection, highlighting recent research and developments in these areas. We also discuss the potential applications of these techniques in various domains, including marketing, healthcare, and social media analysis. Overall, this paper provides a comprehensive overview of opinion mining and emotion detection in natural language processing. It provides insights into the methods, challenges, and potential applications of these techniques, and highlights the importance of these techniques in understanding and analyzing subjective information in text.

1. Introduction

Opinion mining and emotion detection are two important techniques in natural language processing that have gained significant attention in recent years. These techniques have a wide range of applications in various domains, including social media analysis, customer feedback analysis, and product reviews. In this paper, we provide an overview of opinion mining and emotion detection techniques in natural language processing, and discuss their applications, challenges, and limitations.

Importance of Opinion Mining and Emotion Detection: Opinion mining and emotion detection have been widely applied in various domains and applications. For example, in social media analysis, these techniques are used to analyze user-generated content such as tweets, Facebook posts, and online reviews to understand the sentiment and emotions expressed by users. In customer feedback analysis, these techniques are used to analyze customer reviews and feedback to improve product recommendations and customer satisfaction. In healthcare, these techniques are used to analyze patient feedback and detect emotional distress in clinical texts.

One important aspect of opinion mining and emotion detection is the use of sentiment lexicons. A sentiment lexicon is a collection of words and phrases that are associated with positive, negative, or neutral sentiment. Researchers have developed various sentiment lexicons, such as SentiWordNet, VADER, and LIWC, that can be used to analyze sentiment and emotions in text. Colin Raffel, et al. [1] used robustness technique called transfer learning with unified text to text transformer.

Another important area of research in opinion mining and emotion detection is the development of domain-specific models. Since the sentiment and emotions expressed in text can vary depending on the domain and context, it is important to develop models that are specific to a particular domain or application. Researchers have explored various techniques for domain adaptation, such as transfer learning and domain adaptation algorithms [1], to develop effective models for sentiment analysis and emotion detection in specific domains. Also rounding vectors are proposed by Franz A. Heinsen [2]. Also most recent state-of-the art work presented by Radu Tudor Ionescu and Andrei M. Butnaru in [3]. Vector of Locally-Aggregated Word Embeddings (VLAWE) were presented in their work and it achieved state of the art work.

Finally, the literature on opinion mining and emotion detection highlights the ethical and social implications of these techniques. For example, these techniques can be used for surveillance and monitoring of user-generated content, which can raise privacy and ethical concerns. In addition, the accuracy and reliability of these techniques can be affected by biases and cultural differences, which can have social implications.

2. Literature review

Opinion mining and emotion detection are two important techniques in natural language processing that have been widely studied in the literature. Researchers have explored various approaches and methods for analyzing opinions and emotions in text, including machine learning, deep learning, and natural language processing techniques. It is worth to mention that the technique descripted in [4] is one of the efficient works by presenting XLNet. Furthermore, presented FNet in [5] is considered the most competitive and robustness method by providing transformer encoder architectures which sped up by replacing the self-attentions sublayers. However, Amirhossein Abaskohi et al. [6] presented most recent work by proposing LM-CPPF. They mainly recruited generative language models in their work to leverage prompt-based few-shot paraphrasing.

One of the earliest works in opinion mining was the study by Pang and Lee [7], who introduced a supervised learning approach to classify movie reviews as positive or negative. Since then, many researchers have explored various aspects of opinion mining, including feature extraction, sentiment lexicons, and domain adaptation.

Similarly, emotion detection has also been widely studied in the literature. Researchers have explored various methods for identifying emotions in text, including lexicon-based approaches, machine learning, and deep learning techniques. One notable work in this area is the study by Mohammad and Turney [8], who introduced a lexicon-based approach to classify emotions in tweets.

Recently, there has been a growing interest in the use of deep learning techniques for opinion mining and emotion detection. For example, Li et al. [9] proposed a deep learning approach for sentiment analysis that uses attention mechanisms to identify important features in text. In addition, various neural network architectures, such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs), have been explored for emotion detection.

Despite the progress made in opinion mining and emotion detection, there are still several challenges and limitations that need to be addressed. For example, the subjectivity of language and cultural differences can make

it difficult to develop effective models for sentiment analysis and emotion detection. In addition, the lack of labeled data and the need for domain-specific models can also pose challenges in these areas.

Overall, the literature on opinion mining and emotion detection highlights the importance of these techniques in natural language processing and their potential applications in various domains. Future research should focus on developing more effective models that can address the challenges and limitations of these techniques.

3. Background

Opinion mining and emotion detection are two important techniques in natural language processing that have been widely studied in the literature. Researchers have explored various approaches and methods for analyzing opinions and emotions in text, including machine learning, deep learning, and natural language processing techniques. One of the earliest works in opinion mining was the study by Pang and Lee [7], who introduced a supervised learning approach to classify movie reviews as positive or negative. Since then, many researchers have explored various aspects of opinion mining, including feature extraction, sentiment lexicons, and domain adaptation.

In addition, emotion detection has also been widely studied in the literature. Researchers have explored various methods for identifying emotions in text, including lexicon-based approaches, machine learning, and deep learning techniques. The study by Mohammad and Turney [8] introduced a lexicon-based approach to classify emotions in tweets. Similarly, the work by Schuller et al. [13] explores various methods for analyzing emotion and personality in speech and language processing.

Recently, there has been a growing interest in the use of deep learning techniques for opinion mining and emotion detection. Li et al. [9] proposed a deep learning approach for sentiment analysis that uses attention mechanisms to identify important features in text. Wang et al. [11] proposed an attention-based LSTM for aspect-level sentiment classification, while Tang et al. [12] developed a sentiment-specific word embedding for Twitter sentiment classification.

Moreover, the literature on opinion mining and emotion detection highlights the importance of sentiment lexicons in analyzing sentiment and emotions in text. Various sentiment lexicons, such as SentiWordNet, VADER, and LIWC, have been developed to analyze the sentiment and emotions expressed in text Cambria & Hussain [10]. These lexicons can be used in various approaches, such as lexicon-based approaches, to classify the sentiment and emotions expressed in text.

Despite the progress made in opinion mining and emotion detection, there are still several challenges and limitations that need to be addressed. For example, the subjectivity of language and cultural differences can make it difficult to develop effective models for sentiment analysis and emotion detection. In addition, the lack of labeled data and the need for domain-specific models can also pose challenges in these areas. Medhat et al. [15] provide a survey of various sentiment analysis algorithms and applications and highlight the challenges and limitations of these techniques. Furthermore, Narsimha Chilkuri et al. [16] presented parallelizing linear time – invariant(LTI) which can be parallelized during training leading to avoid shortcoming of training RNNs on GPUs.

4. Motivations

Motivations for Studying Opinion Mining and Emotion Detection: Opinion mining and emotion detection techniques have gained significant attention in recent years due to their potential applications and impact on society. One of the key motivations for studying these techniques is their ability to analyze and understand subjective information in text, such as sentiment and emotions. This is particularly important in the age of social media and online communication, where people express their opinions and emotions in various forms of digital communication. Also recent works show the capability of processing sequential inputs [17] which can leverage studying of long time dependencies. Thus, it can be recruited to be adapted for opinion mining and emotion detection.

Another motivation for studying these techniques is their potential applications in various domains and applications. For example, in social media analysis, these techniques can be used to analyze the sentiment and emotions expressed in user-generated content, such as tweets and Facebook posts, to understand public opinion

and behavior. In customer feedback analysis, these techniques can be used to analyze customer reviews and feedback to improve product recommendations and customer satisfaction. In healthcare, these techniques can be used to analyze patient feedback and detect emotional distress in clinical texts, which can help healthcare professionals to provide better care and support.

Moreover, the impact of these techniques extends beyond specific applications, as they have the potential to transform the way we analyze and understand subjective information in text. By providing insights into the sentiment and emotions expressed in text, these techniques can help us to understand human behavior and decision-making processes, and to develop more effective strategies for communication and persuasion.

In addition, these techniques raise important ethical and social implications, such as issues related to privacy, bias, and cultural differences. By studying and developing these techniques, we can address these challenges and ensure that they are used in a responsible and ethical manner.

Overall, the motivations for studying opinion mining and emotion detection techniques are multifaceted and extend beyond specific applications [18]. These techniques have the potential to transform the way we analyze and understand subjective information in text, and to provide insights into human behavior and decision-making processes. By studying and developing these techniques, we can ensure that they are used in a responsible and ethical manner, and contribute to the development of a more informed and communicative society.

5. Challenges and Limitations

Despite the progress made in opinion mining and emotion detection techniques, there are still several challenges and limitations that need to be addressed. One of the main challenges is the subjectivity of language, which can make it difficult to develop effective models for sentiment analysis and emotion detection. Language is highly contextual and can vary depending on the culture, age, gender, and other factors of the speaker or writer. Therefore, it can be challenging to develop models that can accurately capture and analyze the sentiment and emotions expressed in text [19].

Another challenge is the lack of labeled data, which is necessary for training and evaluating sentiment analysis and emotion detection models. The availability of labeled data varies across domains, and it can be costly and time-consuming to collect and annotate data for specific domains or applications. Moreover, the need for domain-specific models can also pose challenges, as the same models may not be effective across different domains or applications.

In addition, there are ethical and social implications associated with the use of these techniques, such as issues related to privacy, bias, and cultural differences. For example, the use of sentiment analysis and emotion detection techniques can raise concerns about the privacy and security of personal information, as well as the potential for biased or discriminatory outcomes based on cultural or linguistic differences.

Need for Further Research: To address these challenges and limitations, there is a need for further research in opinion mining and emotion detection. This research can focus on developing more effective models for sentiment analysis and emotion detection, as well as exploring new methods and techniques for analyzing subjective information in text. In addition, research can also focus on addressing ethical and social implications associated with the use of these techniques, and developing guidelines and best practices for using these techniques in a responsible and ethical manner.

Moreover, there is a need for more comprehensive and diverse datasets for training and evaluating sentiment analysis and emotion detection models. This can involve collecting and annotating data across various domains and applications, and addressing issues related to bias and cultural differences.

6. Sentiment lexicons

Sentiment lexicons are dictionaries that contain words and phrases with positive or negative sentiment. These lexicons are commonly used in sentiment analysis tasks, where the goal is to determine the sentiment of a given

text. Sentiment lexicons can also be used in emotion detection tasks, where the goal is to detect the emotional content of a text.

One of the main advantages of using sentiment lexicons is that they can be easily integrated into sentiment analysis models. In lexicon-based approaches, the sentiment of a text is determined by the presence of words with positive or negative sentiment. For example, if a text contains more positive words than negative words, it is classified as having a positive sentiment.

Sentiment lexicons can also be used in conjunction with machine learning and deep learning algorithms. In this approach, sentiment lexicons are used as features in sentiment analysis models, along with other linguistic features, such as Part-of-Speech (POS) tags and syntactic features. This approach can improve the accuracy of sentiment analysis models, particularly when labeled data is limited.

However, the use of sentiment lexicons is limited by the coverage and accuracy of the lexicon. Sentiment lexicons may not capture the nuanced meanings and contexts of words and phrases, which can result in inaccurate sentiment analysis. Moreover, the sentiment of a text can also be affected by the ordering and combination of words, which may not be captured by sentiment lexicons.

7. Discussion of Key Findings

The literature review on opinion mining and emotion detection has highlighted several key findings and insights that have important implications for the development and application of these techniques. Some of these implications include:

- 1. The importance of domain-specific models and datasets: The effectiveness of sentiment analysis and emotion detection is highly dependent on the domain and application. Therefore, the development of domain-specific models and datasets is critical for achieving accurate and reliable results. This has implications for researchers and practitioners who need to consider the specific context and domain when applying sentiment analysis and emotion detection techniques.
- 2. The need for effective feature engineering: Feature engineering is an important aspect of sentiment analysis and emotion detection, as the choice of features can significantly impact the accuracy and robustness of the models. Therefore, researchers and practitioners need to carefully consider the selection and combination of features when designing sentiment analysis and emotion detection models.
- 3. The potential for deep learning algorithms: Deep learning algorithms, such as CNNs and RNNs, have shown promising results in sentiment analysis and emotion detection tasks. These algorithms can effectively capture the context and semantics of text, which can improve the accuracy of sentiment analysis and emotion detection models. Therefore, researchers and practitioners may consider incorporating deep learning algorithms into their sentiment analysis and emotion detection models.
- 4. The limitations and challenges associated with sentiment analysis and emotion detection: The literature review has identified several limitations and challenges associated with sentiment analysis and emotion detection, such as the subjectivity of language, the lack of labeled data, and ethical and social implications. These limitations and challenges should be carefully considered when applying sentiment analysis and emotion detection techniques, and further research is needed to address these issues. The key findings and insights from the literature review have important implications for the development and application of sentiment analysis and emotion detection techniques. By carefully considering these implications, researchers and practitioners can effectively apply these techniques in various domains and applications, and contribute to the development of a more communicative and informed society.

8. Discuss the potential applications and impact of these techniques in various domains and applications

Sentiment analysis and emotion detection have numerous potential applications and impacts across various domains and applications. Some of these potential applications and impacts include:

- 1. Social media: Sentiment analysis and emotion detection can be used to analyze the sentiment and emotions expressed in social media data, such as tweets and posts, and provide insights into public opinion on various topics. This can have implications for businesses, governments, and organizations, who can use this information to improve their products, services, and communication strategies.
- 2. Healthcare: Sentiment analysis and emotion detection can be used to analyze patient feedback, medical records, and social media data related to healthcare, and provide insights into patient experiences, attitudes, and emotions. This can have implications for improving patient care, identifying areas of concern, and developing new treatment strategies.
- 3. Business: Sentiment analysis and emotion detection can be used to analyze customer feedback, product reviews, and social media data related to businesses, and provide insights into customer experiences, attitudes, and emotions. This can have implications for improving customer satisfaction, identifying areas of concern, and developing new marketing strategies.
- 4. Politics: Sentiment analysis and emotion detection can be used to analyze political discourse, news articles, and social media data related to politics, and provide insights into public opinion and sentiment towards political issues and candidates. This can have implications for improving political campaigns, communication strategies, and policy decisions.

The potential impact of sentiment analysis and emotion detection on society is significant. These techniques can contribute to the development of a more communicative and informed society, by providing insights into the attitudes and emotions of individuals and groups. However, the ethical and social implications of these techniques, such as issues related to privacy, bias, and cultural differences, need to be carefully considered and addressed. Overall, sentiment analysis and emotion detection have numerous potential applications and impacts across various domains and applications, and can contribute to the development of a more insightful and empathetic society.

9. Conclusion

1. The literature review on opinion mining and emotion detection has highlighted several key findings and insights. Firstly, sentiment analysis and emotion detection are critical techniques for analyzing the attitudes and emotions of individuals and groups, and have numerous potential applications and impacts across various domains and applications. Secondly, a variety of approaches and methods can be used in sentiment analysis and emotion detection, including machine learning algorithms, deep learning algorithms, and natural language processing techniques. Each approach and method has its own strengths and weaknesses, and researchers and practitioners should carefully consider the specific domain and application when selecting and applying these techniques. Thirdly, the literature review has identified several limitations and challenges associated with sentiment analysis and emotion detection, such as the subjectivity of language, the lack of labeled data, and ethical and social implications. Further research is needed to address these issues and improve the effectiveness and robustness of sentiment analysis and emotion detection have significant potential for contributing to the development of a more communicative and informed society, but careful consideration of the specific context and limitations is necessary to achieve accurate and reliable results

The findings from the literature review have several implications for future research and development in sentiment analysis and emotion detection. Firstly, there is a need for further research on domain-specific models and datasets, as the effectiveness of sentiment analysis and emotion detection is highly dependent on the domain and application. Therefore, more efforts should be devoted to developing domain-specific models and datasets to achieve accurate and reliable results in various domains.

Secondly, there is a need for further research on effective feature engineering and model tuning. Different features can significantly impact the accuracy and robustness of sentiment analysis and emotion detection models. Therefore, more research is needed to explore different feature selection and combination methods, and to optimize the models for different applications.

Thirdly, there is a need for further research on the ethical and social implications of sentiment analysis and emotion detection. These techniques have significant potential for contributing to the development of a more communicative and informed society, but there are also concerns related to privacy, bias, and cultural differences. Therefore, more research is needed to explore the ethical and social implications of these techniques and to develop methods for addressing these concerns.

Lastly, there is a need for more research on the combination of different approaches and methods in sentiment analysis and emotion detection. The literature review has highlighted the strengths and weaknesses of different approaches and methods, and combining these methods may lead to more effective and robust sentiment analysis and emotion detection models.

Based on the literature review, the following recommendations for future research in sentiment analysis and emotion detection can be made:

- 2. Develop more accurate and robust sentiment analysis and emotion detection models: There is a need for further research on developing more accurate and robust sentiment analysis and emotion detection models that can effectively capture the nuances and complexities of language. This can be achieved through the development of more sophisticated machine learning and deep learning algorithms, as well as through the use of more domain-specific models and datasets.
- 3. Address the ethical and social implications of sentiment analysis and emotion detection: There is a need for further research on the ethical and social implications of sentiment analysis and emotion detection, particularly in terms of privacy, bias, and cultural differences. This can be achieved through the development of more transparent and interpretable models, as well as through the implementation of ethical and social guidelines and standards.
- 4. Combine different approaches and methods in sentiment analysis and emotion detection: There is a need for more research on the combination of different approaches and methods in sentiment analysis and emotion detection, particularly in terms of combining machine learning, deep learning, and natural language processing techniques. This can lead to more effective and robust sentiment analysis and emotion detection models.
- 5. Develop more effective feature engineering and model tuning methods: There is a need for more research on effective feature engineering and model tuning methods in sentiment analysis and emotion detection, particularly in terms of selecting and combining different features to improve the accuracy and robustness of the models.
- 6. Explore new domains and applications of sentiment analysis and emotion detection: There is a need for more research on the application of sentiment analysis and emotion detection in new domains and applications, such as in healthcare, finance, and education. This can lead to new insights and applications of sentiment analysis and emotion detection, as well as contribute to the development of a more communicative and informed society.

10. Summary

In conclusion, sentiment analysis and emotion detection are critical techniques for analyzing the attitudes and emotions of individuals and groups, and have numerous potential applications and impacts across various domains and applications. The literature review has identified several key findings and insights, including the different approaches and methods used in sentiment analysis and emotion detection, the importance of domain-specific models and datasets, the need for effective feature engineering and model tuning, and the limitations and challenges associated with these techniques.

The implications of these findings have important implications for future research and development in sentiment analysis and emotion detection. Further research is needed to develop more accurate and robust models, address

the ethical and social implications of these techniques, combine different approaches and methods, develop more effective feature engineering and model tuning methods, and explore new domains and applications.

Sentiment analysis and emotion detection have significant potential for contributing to the development of a more communicative and informed society. By carefully considering the specific context and limitations of these techniques, researchers and practitioners can effectively apply these techniques in various domains and applications, and contribute to the development of a more insightful and empathetic society.

Conflict of Interest: The authors declare that there are no conflicts of interest associated with this research project. We have no financial or personal relationships that could potentially bias our work or influence the interpretation of the results.

References

- [1]C. Raffe and N. Shazeer, et al., "Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer", *Journal of Machine Learning Research*, vol. 21, pp 1-67, 2020.
- [2] H. Franz A, "An algorithm for routing vectors in sequences." arXiv preprint arXiv:2211.11754, 2022.
- [3] R. T. Ionescu and A. M. Butnaru, "Vector of locally-aggregated word embeddings (vlawe): A novel documentlevel representation", arXiv preprint arXiv:1902.08850, 2019.
- [4] Y. Zhilin, et al. "Xlnet: Generalized autoregressive pretraining for language understanding.", Advances in neural information processing systems 32, 2019.
- [5] L.T. James, et al. "Fnet: Mixing tokens with fourier transforms." arXiv preprint arXiv:2105.03824, 2021.
- [6] A.Amirhossein, S. Rothe, and Y. Y.zadeh. "LM-CPPF: Paraphrasing-Guided Data Augmentation for Contrastive Prompt-Based Few-Shot Fine-Tuning." *arXiv preprint arXiv:2305.18169*, 2023.
- [7] P. Bo and L. Lee, "Foundations and Trends® in information retrieval." *Foundations and Trends*® *in Information Retrieval*, vol. 2, no. 1-2, pp 1-135, 2008.
- [8] M. Saif, and P. Turney, "Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon." *Proceedings of the NAACL HLT 2010 workshop on computational approaches to analysis* and generation of emotion in text, pp 26-34, 2010.
- [9] Y.Li and X.Xu, "Attention-based convolutional neural network for sentiment analysis". *IEEE Access*, vol. 7, pp 17544-17551, 2019.
- [10] E. Cambria and A. Hussain, "Opinion mining and sentiment analysis". *IEEE Intelligent Systems*, vol. 27, no. 6, pp 46-59, 2012.
- [11] Y.Wang , M. Huang, X. Zhu, L. Zhao, "Attention-based lstm for aspect-level sentiment classification", Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pp 606-615, 2017.
- [12] D. Tang, F.Wei, N.Yang, M. Zhou, T. Liu, B. Qin, "Learning sentiment-specific word embedding for twitter sentiment classification", *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pp 1555-1565, 2014.
- [13] B. Schuller, A. Batliner, and F.Burkhardt, "Computational paralinguistics: Emotion, affect and personality in speech and language processing", *John Wiley & Sons*, 2013.
- [14] B.Liu, "Sentiment analysis and opinion mining", *Synthesis Lectures on Human Language Technologies*, vol. 5, no. 1, pp 1-167, 2012.
- [15] W. Medhat , A.Hassan, and H. Korashy , "Sentiment analysis algorithms and applications: A survey", *Ain Shams Engineering Journal*, vol. 5, no. 4, pp 1093-1113, 2014.
- [16] Chilkuri, N. Reddy, and C. Eliasmith, "Parallelizing legendre memory unit training." *International Conference on Machine Learning. PMLR*, pp 1898-1907, 2021.
- [17] Rusch, T. Konstantin, and Siddhartha Mishra, "Unicornn: A recurrent model for learning very long time dependencies." *International Conference on Machine Learning. PMLR*, pp 9168-9178, 2021.
- [18] Rusch, T. Konstantin, and S. Mishra, "Coupled Oscillatory Recurrent Neural Network (coRNN): An accurate and (gradient) stable architecture for learning long time dependencies." *arXiv preprint arXiv:2010.00951*, 2020.
- [19] Li, Xianming, and L. Jing, "Angle-optimized text embeddings." arXiv preprint arXiv:2309.12871, 2023.