An N-Gram Framework for Sentiment and Emotion-Aware Word Association Games

Authors

Rohan Dalal (Pennsylvania State University; rxd5491@psu.edu)
Sanjana Menon (Pennsylvania State University; ssm5808@psu.edu)
Sohan Hajra (University of Illinois at Urbana-Champaign; shajra2@illinois.edu)
Jeremy Blum, D.Sc. (Pennsylvania State University; jjb24@psu.edu)

DOI: 10.1145/3774399.3774407

Abstract

Natural Language Processing (NLP) models have seen impressive advancements in understanding word associations. Still, limited attention has been given to user sentiment and emotions influencing these associations. In this paper, we explore the impact of both sentiment and emotion on the selection of words in an N-gram-based word association game. By integrating GloVe embeddings with SenticNet-based emotion classification and sentiment analysis from VADER, we evaluate how positive and negative sentiments, combined with intense emotions such as delight, ecstasy, terror, and loathing, affect word choices. Our findings suggest that user sentiment and emotion have a significant effect on word selection, with positive players associating positive adjectives with positive nouns, while negative players tend toward the opposite associations. This study highlights the necessity of considering sentiment and emotional intelligence in future NLP systems and presents new applications in areas such as AI-based gaming, behavioral analysis, and human-computer interaction.

Introduction

Word association games are a form of entertainment that challenges players to think creatively while considering the emotional and conceptual context of the words involved. In Apples-to-Apples, players match descriptive adjectives with nouns with the goal of having their selected adjective chosen by a designated judge. The open-ended nature of these games presents a challenge in terms of optimizing gameplay strategy, particularly in how players make decisions based on the emotional and semantic context of words. In particular, the goal is to select words that align with the given target adjective in a way that reflects both semantic similarity and emotional resonance with the judge. This study focuses on the development of an approach that selects words for word association games by simulating a player that can account for sentiment, emotional context, and semantic relationships. This requires not only an understanding of word meanings but also the emotional nuances and subjective preferences of individual players.

The objective of the Refined NGram-based Player (RNP) model is to act as a dynamic player that uses sentiment analysis and emotion classification to align card choices with the given adjective to reflect user biases and emotional profiles. The model uses Valence Aware Dictionary and sEntiment Reasoner (VADER) for sentiment analysis, SenticNet for emotion classification, and GloVe word em beddings to measure semantic similarity between words, enhancing its ability to select more relevant words based on the target adjective. Through preset tests that reflect diverse user profiles, the model's ability to select nouns that are most likely to align with target adjectives is evaluated. The results showed a measurable improvement in card selection alignment.

Related Work

Research in sentiment analysis, semantic similarity, and artificial intelligence has made significant strides in linguistic and emotion analysis. Our work aligns with studies examining how sentiment and emotion interact with AI-driven models, particularly in contexts requiring human- like decision-making.

As Hutto and Gilbert (2014) note, "Sentiment analysis is useful to a wide range of problems that are of interest to human-computer interaction practitioners and researchers, as well as those from fields such as sociology, marketing and advertising, psychology, economics, and political science" (p. 216). Traditional approaches have predominantly relied on rule-based models and machine learning classifiers. For instance, Zainudin et al. (2019) highlighted the potential of K Nearest Neighbors classifiers (KNN) to out perform Support Vector Machines (SVMs) in precision and recall metrics for sentiment classification tasks. Pang and Lee (2008) explored opinion mining, focusing on challenges like domain dependency, sentiment polarity, and subjectivity detection. Their work emphasized the inherent complexities of sentiment analysis, especially in disambiguating nuanced expressions across diverse contests. These studies focused on relatively static con texts, such as reviews or social media posts. Their integration into dynamic, emotion-driven systems like word association games remains underexplored.

This work also builds on previous work by incorporating VADER, a rule-based sentiment analysis model optimized for short text, such as tweets. VADER has been shown to out perform traditional models, including Linguistic Inquiry and Word Count, in microblogging environments due to its nuanced handling of intensity modifiers, negations, and emoticons (Hutto & Gilbert, 2014). Integrating VADER enables the proposed RNP model to capture sentiment polarity (positive, negative, neutral) with high accuracy.

To further enhance the emotional understanding of words, the RNP model integrates Sen ticNet (Cambria et al., 2016), a concept level sentiment analysis tool that links natural language concepts to fine-grained emotional states such as delight, serenity, and terror. Unlike traditional sentiment analysis frameworks that rely solely on word-level analysis, SenticNet combines cognitive and affective in formation, leveraging commonsense reasoning and psychology to understand both the semantics and connotations of multi- word expressions. By dynamically extracting and applying emotional profiles, SenticNet enables our model to adjust card selection based on the player's emotional context, improving personalization and fostering human-like decision-making.

Word embeddings allow for the quantification of semantic relationships between words. GloVe (Global Vectors for Word Representation) captures co-occurrence statistics, enabling it to model both local and global word contexts (Pennington et al., 2014). Unlike traditional frequency-based methods, GloVe creates dense vector representations that facilitate fine-grained similarity assessments. This is particularly relevant for our study, as semantic alignment between adjectives and nouns is a critical factor in word association games. Despite other alternatives like Word2Vec and FastText offering similar capabilities, GloVe was chosen for its effectiveness in smaller datasets and its pre-trained vectors' ability to generalize across linguistic domains. By integrating GloVe embeddings, our model ensures that selected word pairs maintain both semantic coherence and emotional relevance.

Finally, studies on bigrams provide relevant context for our work. For example, Nguyen et al. (2016) explored bigram entropy analysis to detect significant events on social media, demonstrating how word pair distributions reveal insights into shifts in meaning and social dynamics. Similarly, our model uses bigrams to assess the alignment between adjectives and nouns, ensuring that player choices are both semantically and emotionally consistent with the target words.’

Model Details

The Refined NGram-based Player (RNP) model enhances word association games by integrating semantic similarity, sentiment analysis, and emotion classification to simulate human-like decision-making. The model selects noun cards that align with a given target adjective while accounting for both the semantic meaning and the emotional preferences of the user.

The development of the model relied on two key datasets: an adjective-noun dataset, compiled from open-source linguistic resources, and a user tweets dataset, curated to capture diverse user sentiment profiles categorized as positive, neutral, and negative.

Datasets

Our research involved two key datasets. An adjective-noun dataset was used as a basis for determining the semantic similarity between words and a curated user tweets dataset which was used for analyzing potential judge sentiments. We compiled CSV files of 4,844 adjectives (Siem, 2019) and 6,800 nouns (Leite, 2018) from Kaggle to be used in the word association game. These datasets were preprocessed to remove duplicates and ensure variety.

To create a robust and reproducible dataset for sentiment and emotion analysis, we curated tweets from publicly accessible celebrity accounts due to their expressive language and accessibility. We began by selecting 12 users to represent three sentiment profiles: positive, negative, and neutral. Sentiments were formally defined using VADER sentiment scores, where tweets with a compound score greater than 0.05 were labeled positive, between - 0.05 and 0.05 as neutral, and less than - 0.05 as negative. For each user, we initially collected 10 tweets and analyzed their sentiment scores to select the 5 best candidates for evaluation and testing purposes—2 users with predominantly positive sentiment, 2 with negative sentiment, and 1 with a neutral profile. To refine each sentiment profile, we manually gathered 20 recent tweets (from the past 5 years) per user, ensuring they were in English for uniformity. From these, we randomly sampled 10 tweets per user to build a balanced and representative sentiment profile.

Card Selection Process

The card selection process, depicted in Figure 1, is a multi- step approach that integrates semantic similarity, sentiment analysis, emotion classification, and bigram validation to make intelligent and user-aligned decisions. The objective is to select the noun card that most closely aligns with the target adjective while accounting for the user’s sentiment and emotional profile. This process enables the model to simulate human-like decision-making in word association games.

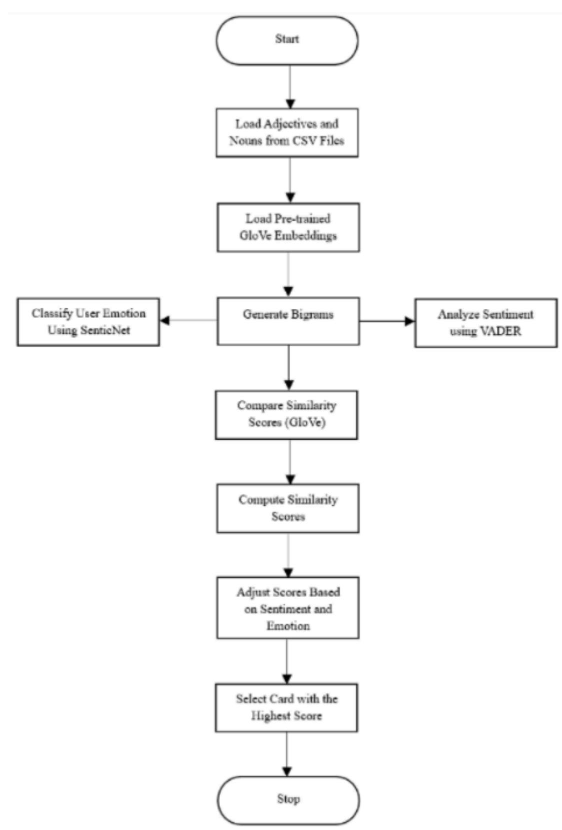

Figure 1: The overall architecture of the Sentiment Analysis Model.

The first step involves computing semantic similarity between the target adjective and each card in the player’s hand using GloVe word embeddings. These pre-trained embeddings map words into a high-dimensional vector space, allowing the model to quantify the semantic closeness between words. The similarity score for each card is calculated using the cosine similarity function, which measures how closely the vector representation of the adjective aligns with the noun. To emphasize the importance of contextual relevance in word associations, this similarity score is amplified by a factor of two—ensuring that cards with higher semantic alignment are prioritized.

Once the semantic similarity scores are established, the model incorporates sentiment analysis using VADER (Valence Aware Dictionary and sEntiment Reasoner). VADER assigns a compound sentiment score to each noun card and the target adjective, classifying the words as positive, neutral, or negative. The user’s sentiment profile, derived earlier through VADER analysis of their tweets, is then compared with the sentiment of the cards. The model dynamically adjusts scores based on this alignment. For instance, if the user has a positive sentiment and the target adjective is also positive, the model assigns a higher score to noun cards carrying positive sentiments. Conversely, when a negative sentiment user interacts with a negative adjective, cards with stronger negative connotations are favored. This step personalizes card selection to resonate with the user’s sentiment bias. The model further refines the selection process by incorporating emotional context using Sentic Net. Each user is assigned a dominant emotional state (e.g., ”delight,” ”terror,” or ”serenity”) that influences the decision-making process. Cards that align with the user’s emotional state are given a boost in score, while cards that conflict with the emotion receive a penalty. Strong emotions, such as ”ecstasy” or ”loathing,” exert a greater influence, amplifying their effect on the scores, whereas milder emotions like ”serenity” lead to more subtle score adjustments. This ensures that the model reflects not only the user’s senti ment but also their emotional intensity.

To ensure linguistic validity, the model verifies whether the selected noun card forms a valid bigram with the target adjective. This validation is performed using Part-of-Speech (POS) tagging to confirm that the first word is an adjective and the second is a noun. Additionally, WordNet synset validation ensures that the chosen noun is contextually relevant to the adjective. If the bigram passes both checks, the score for the card receives a further boost, enhancing its likelihood of being selected.

The final step involves aggregating the scores from all the aforementioned components—semantic similarity, sentiment alignment, emotional context, and bigram validation. The noun card with the highest total score is selected as the optimal match for the target adjective. In cases where no card stands out or multiple cards receive identical scores, the model defaults to selecting a random card to ensure smooth and continuous gameplay.

Experimentation and Testing

The experimentation and testing phase of this study was divided into two stages to comprehensively evaluate the effectiveness of the RNP model in a word association game context. The first stage focused on benchmarking the Base RNP Model against traditional approaches for card selection, including models like Word2Vec, Random selection, EditDis tance, and SpaCy. This comparison aimed to assess the Base RNP Model’s ability to integrate semantic similarity, sentiment analysis, and emotional alignment in decision making relative to these established methods. The second stage of testing introduced personalized sentiment profiles—Positive, Negative, and Neutral—by modifying the Base RNP Model to account for user-specific emotional and sentiment contexts. This stage evaluated the model’s ability to adapt its card selection process to mimic user personas and align closely with their emotional profiles. Together, these stages aim to provide a holistic evaluation of the model’s performance, both as a general-purpose AI-driven player and as a personalized, emotion-aware participant in the game.

Base RNP Model

The Base RNP Model in our study serves as a foundational approach to card selection in word association games, relying on a combination of semantic similarity and static rules to make decisions. Unlike models integrated with user data, the Base RNP Model operates without personalization, applying generic sentiment scores derived from VADER and focusing solely on static attributes of the words. While it calculates semantic similarity using GloVe embeddings to ensure contextual relevance between adjectives and nouns, it does not account for user-specific emotional or sentiment-driven nuances. The absence of dynamic scoring and adaptability renders it a more traditional, rule-based solution, making it ideal for baseline comparisons. This approach makes the model a static decision-maker, where its output is determined by the pre-defined logic and the hand of cards, without any variability based on user-specific inputs. The general-purpose framework of this model highlights the limitations of non-personalized models and serves as a benchmark to evaluate the enhanced performance of sentiment and emotion-aware models, which dynamically adjust scoring and adapt to user profiles.

RNP Model with User Sentiment Profiles

The RNP Model with User Sentiment Profiles builds upon the Base RNP Model by integrating user-specific sentiment and emotional data to dynamically tailor its decision- making process. Unlike the static approach of the Base RNP Model, this advanced model personalizes card selection by leveraging user sentiment (positive, negative, or neutral) and dominant emotional states (e.g., delight, terror, serenity). Through dynamic weighting, the model adjusts scoring based on the alignment between the user's sentiment and the target adjective's sentiment. For instance, when the user sentiment is positive, the model prioritizes cards with positive sentiment for positive adjectives while penalizing negative ones. Conversely, for a user with negative sentiment, the model reverses this approach, favoring negative cards for negative adjectives.

Additionally, the model incorporates emotion specific adjustments, where dominant user emotions further influence scoring. High arousal emotions like delight or loathing amplify the selection bias, either favoring or penalizing specific cards depending on the con text, while neutral emotions like serenity have subtler effects. This emotion-aware bias enhances the model's ability to mimic human-like decision-making by considering both the semantic relevance of the word pairings and the user's emotional context. By reweighting semantic similarity, sentiment alignment, and emotional biases, the model creates a personalized and context-sensitive experience. Over multiple runs, the model does not ”train” in the machine learning sense (i.e., it doesn't update its parameters or learn from past decisions). However, the presence of a user specific sentiment profile makes the decision making process appear tailored for the user. As a result, this model behaves as a personalized decision-maker where every user will likely get a different outcome based on their profile, even with the same target adjective and hand of cards.

First Stage of Testing

The first stage of testing evaluates the effectiveness of the Base RNP Model against traditional card selection models by using a curated dataset of adjective-noun pairs. After creating the finalized user tweets dataset, which provided sentiment and emotion profiles for selected users; we manually generated 10 target adjectives and corresponding hands of cards for each of the six emotions: delight, ecstasy, serenity, enthusiasm, terror, and loathing. Each target adjective represented a specific emotion, and the hands of cards consisted of nouns that varied in sentiment and semantic similarity to the target adjective.

To establish a ground truth for the testing phase, we categorized these six emotions into positive, neutral, and negative sentiment groups based on psychological and linguistic standards. Positive emotions include de light, ecstasy, and enthusiasm, as they convey happiness, excitement, and positivity. Neutral emotions, such as serenity, represent calmness or a lack of strong sentiment. Negative emotions, including terror and loathing, were defined by their association with fear, disgust, or negativity. This classification ensured a clear and consistent framework for evaluating the model's performance, as the ground truth determined whether the selected card aligned with the target adjective's emotional and semantic context. A total of 60 examples (10 per emotion) were compiled into a CSV file, providing a structured and reproducible basis for evaluation.

The goal of the testing against the structured dataset was to assess the extent to which the Base RNP Model could align its selections with the target emotion and sentiment, providing a benchmark for comparing it with other traditional models. This approach highlighted the importance of sentiment and emotional understanding in creating contextually accurate and human-like gameplay

In this stage, we evaluated the performance of the Base RNP Model by comparing it with the following models and each model had its unique approach for selecting cards given a target adjective and hand of cards (nouns):

1. Word2Vec model: Leverages pre-trained Word2Vec embeddings to calculate the semantic similarity between the target adjective and nouns in the hand. The card with the highest similarity score is selected, ensuring contextually relevant choices.
2. SpaCy Model: Utilizes SpaCy's large English language model for word embeddings and integrates Hugging Face's transformers for sentiment analysis. This model incorporates humor-based decision-making for a more dynamic and engaging approach.
3. Random Model: Makes entirely arbitrary selections, offering no semantic or contextual basis for its choices.
4. EditDistancePlayer Model: Chooses the card with the smallest Levenshtein distance to the target, prioritizing lexical closeness.

The testing procedure involved extracting examples from the curated dataset, consisting of predefined emotions, corresponding target adjectives, and their associated noun combinations (hand of cards).

Each model was run across all 60 target adjectives and their respective hand of cards (noun combinations). The chosen card, sentiment score, and emotion for each example were stored, generating a dataset, which was further used for modelling the results.

To evaluate the performance of each model, we utilized Root Mean Squared Error (RMSE) as the primary metric. For this study, we assigned a sentiment value of +1 to positive emotions and -1 to negative emotions. The model- generated sentiment score was then compared against these ground truth values to calculate RMSE.

Model Accuracy

The models were evaluated for accuracy on the test set of curated examples. Figure 2 shows the RMSE values for each model across the six emotions: delight, ecstasy, serenity, enthusiasm, terror, and loathing. The plots illustrate that the RNP model consistently outperformed the others in positive emotions like delight and ecstasy, with notably lower RMSE values. However, the performance dropped for negative emotions such as loathing, especially for negative sentiment users. The RNP model performed best overall, especially in categories with extreme positive emotions.

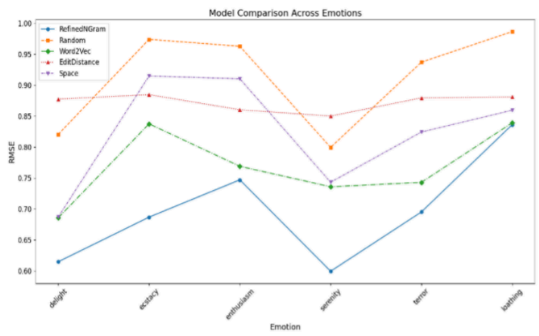

Figure 2: Performance of the models as a function of the underlying emotions.

Second Stage of Testing

The second stage of testing focused on assessing the ability of the RNP with user sentiment profiles to adapt its card selection process to personalized user sentiment profiles. This stage directly addressed our research question: can the model mimic user sentiment in word association games to provide more human-like and engaging gameplay? To achieve this, we introduced five personalized sentiment profiles derived from a curated user tweet dataset (refer to Section 3.1.2), representing three distinct sentiment orientations:

- Positive sentiment profiles: Two users with consistently positive sentiment scores, e.g. Taylor Swift.
- Negative sentiment profiles: Two users with consistently negative sentiment scores, e.g. Piers Morgan.
- Neutral sentiment profile: One user with neutral sentiment scores, e.g. Rihanna.

The model was modified to integrate these sentiment profiles by dynamically adjusting scoring based on user sentiment and emotion. The personalized RNP Model with user sentiment profiles were then compared to the base RNP Model to evaluate the impact of sentiment personalization on decision-making.

Using the same dataset of target adjectives and noun combinations generated in Stage 1, we calculated the sentiment scores for the personalized models and the Base RNP Model across all six emotions: delight, ecstasy, serenity, enthusiasm, terror, and loathing. RMSE values were computed to measure the accuracy of each model's sentiment alignment relative to the ground truth. The inclusion of user-specific sentiment profiles allowed us to evaluate how effectively the RNP Model could refine its selections to align with user sentiment, emotional context, and the semantic relevance of target adjectives.

The results from this stage formed the basis for further analysis, providing insights into the benefits of integrating user sentiment data. By comparing the personalized models against the non-personalized Base RNP Model, we demonstrated the model's potential to enhance user engagement and decision-making in word association games through sentiment aware personalization. This stage highlighted the adaptability and context-sensitivity of the RNP Model with User Sentiment profiles, showcasing its ability to emulate human-like decision-making in emotionally dynamic scenarios.

Evaluation of Model Variability

The RNP model exhibits more consistency in its selections. The box plot of RMSE distribution in Figure 3 highlights the variability in each model's performance, showing which models are consistent and which exhibit greater fluctuation.

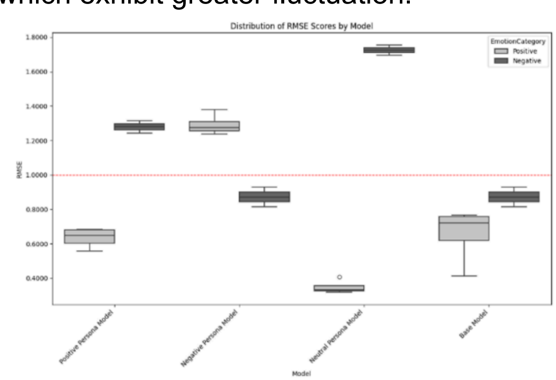

Figure 3: Root Mean Squared Error Distribution across all models and emotions.

As seen in Figure 3, the RNP model has a relatively narrow RMSE distribution, reflecting its consistency in predicting across different emotional categories. In contrast, the Random and EditDistance models exhibit much wider distributions, indicating that their predictions are less reliable. This reinforces earlier observations about the need for models to handle sentiment and semantic context better, particularly in high-variance categories.

Figure 3 illustrates the distribution of RMSE scores for the RNP Model with User Senti ment profiles and Base RNP Model across both positive, neutral, and negative emotional categories during the second stage of testing. The RNP Model with User Sentiment profiles displays a relatively narrow range of RMSE values for both positive and negative emotions, indicating its robustness and consistency in adapting to user sentiment profiles. This aligns with the model's design to account for semantic, sentiment, and emotional factors during decision-making.

The red dashed line in the plot serves as a reference threshold for acceptable RMSE values. While the RNP Model with User Sentiment profiles consistently performs below this threshold, indicating higher accuracy, the Base RNP Model frequently exceeds it, highlighting their limitations in predicting sentiment-aligned outcomes. The stark contrast between the models underscores the importance of incorporating sentiment and emotion-aware mechanisms in achieving reliable and human-like decision- making in word association games.

Limitations

While the RNP model demonstrated significant strengths in handling positive emotions, its performance in dealing with negative emotions such as loathing and terror was noticeably weaker. This suggests that the model's sensitivity to negative sentiment could be enhanced, particularly in distinguishing between nuanced negative emotions. The difficulty in managing these emotions limits the model's overall applicability in scenarios where users may exhibit more complex negative emotional states.

Additionally, the relatively small dataset, 5 users with 10 tweets each, may limit the model's generalizability. A larger and more diverse dataset would likely offer better insights into the full range of emotional expressions and improve the robustness of the model in real-world applications. The limited dataset makes it challenging to fully assess the model's scalability and adaptability to broader, more diverse user groups.

Finally, the use of static GloVe embeddings for measuring semantic similarity may not fully capture the dynamic contextual nuances present in more sophisticated word relationships. Dynamic embeddings could potentially provide a richer, more accurate understanding of both sentiment and semantics, leading to improved performance in the word association game.

Conclusion and Future Work

This study presented a novel approach to incorporating user sentiments into word selection for word association games, leveraging the RNP model to enhance gameplay personalization and alignment with emotional contexts. The experimentation and testing phases highlighted the model's ability to integrate semantic similarity, sentiment analysis, and emotional profiling into decision-making, offering both strengths and areas for improvement.

The RNP model demonstrated strong performance in handling positive emotions such as delight and ecstasy, consistently aligning with user sentiment and enhancing semantic coherence. However, the model struggled with negative emotions like loathing and terror, revealing limitations in adapting to nuanced negative sentiment profiles. This inconsistency underscores the need for further refinement in handling complex emotional states. While the inclusion of user personas highlighted the model's potential to personalize gameplay, it also exposed gaps in performance balance across different emotional spectrums.

The results suggest that sentiment-aware AI models like the RNP model have the potential to enhance user engagement by tailoring gameplay to emotional contexts. However, user engagement was not directly evaluated in this study. While our results indicate potential for improving interactivity, future studies should explicitly measure user engagement through qualitative or quantitative evaluations, such as user surveys or gameplay metrics, to validate this claim. Additionally, the dataset size was a limitation of this study. A larger and more varied dataset, including real- world user input, would improve the model's generalizability and robustness. Advanced transformer-based architectures like BERT (Devlin, Chang, Lee, & Toutanova, 2018) and GPT (Radford & Narasimhan, 2018) could further enhance performance by capturing dynamic contextual relationships and enabling richer sentiment and semantic analysis. Future work could incorporate these models alongside multi-modal data sources such as images or videos to expand the emotional scope of gameplay. Finally, integrating real-time sentiment and emotion analysis from platforms like Twitter into word association games presents an exciting avenue for research. This could enable more responsive and emotionally intelligent game play, enhancing personalization and engagement. Expanding the applications of such models to education or mental health contexts could also offer significant societal benefits, demonstrating the versatility and impact of emotion-aware AI.

References

Kaur, Sumandeep, Geeta Sikka, and Lalit Kumar Awasthi. 2018. "Sentiment Analysis Approach Based on N-Gram and KNN Classifier." In 2018 First International Conference on Secure Cyber Computing and Communication (ICSCCC), 38:1–4. IEEE.

Hutto, C., and Eric Gilbert. 2014. "VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text." Proceedings of the International AAAI Conference on Web and Social Media 8 (1): 216–25. https://doi.org/10.1609/icwsm.v8i1.14550

Wilson, Theresa, and Janyce Wiebe. n.d. "Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis." Aclan thology.org. Accessed September 16, 2024. https://aclanthology.org/H05-1044.pdf.

Erik Cambria, Daniel Olsher, Dheeraj Rajagopal. 2014. "SenticNet 3: A Common and Common-Sense Knowledge Base for Cognition-Driven Sentiment Analysis." Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, pages 1515-1520.

Kanna Akella, N. Venkatachalam, K. Gokul, Keunho Choi and Ramachandraprabhu Tyakal. 2017. "Gain Customer Insights Using NLP Techniques." SAE International Journal of Materials and Manufacturing, pages 333-337.

Rahim Dehkharghani, Yucel Saygin, Berrin Yanikoglu and Kemal Oflazer. 2015. "SentiTurkNet: a Turkish polarity lexicon for sentiment analysis." Language Resources and Evaluation, pages 667-685.

Preslav Nakov, Sara Rosenthal, Svetlana Kiritchenko, Saif M. Mohammad, Zornitsa Kozareva, Alan Ritter, Veselin Stoyanov and Xiaodan Zhu. 2016. "Developing a successful SemEval task in sentiment analysis of Twitter and other social media texts." Language Resources and Evaluation, pages 35-65.

Vishakha Joseph, Chandra Prakash Lora, Narmadha T. 2024. "Exploring the Application of Natural Language Processing for Social Media Sentiment Analysis." 2024 3rd International Conference for Innovation in Technology (INOCON), pages 1-6.

Erik Cambria, Soujanya Poria, Rajiv Bajpai, Bjoern Schuller. 2016. SenticNet 4: A Semantic Resource for Sentiment Analysis Based on Conceptual Primitives. The COLING 2016 Organizing Committee, pages 2666–2677.

Siem, J. (2019). Adjectives List [Data set]. www.kaggle.com/datasets/jordansiem/ adjectives-list

Leite, M. (2019). List of Nouns [Data set]. www.kaggle.com/datasets/leite0407/list-of nouns

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional Transformers for lan guage understanding. In arXiv [cs.CL]. https://doi.org/10.48550/ARXIV.1810.04805

Radford, A., & Narasimhan, K. (2018). Improving Language Understanding by Generative Pre-Training.

Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. Foundations and Trends® in Information Retrieval, 2(1–2), 1–135. https://doi.org/10.1561/1500000011

Cambria, E., Zhang, X., Mao, R., Kwok, K., & Chen, M. (2024). SenticNet 8: Fusing emotion AI and commonsense AI for interpretable, trustworthy, and explainable affective computing. In Proceedings of the 2024 Conference on Affective Computing and Intelligent Interaction (ACII). College of Computing and Data Science, Nanyang Technological University.

Rohan Dalal is a student at the Pennsylvania State University majoring in Applied Data Sciences. His academic interests include artificial intelligence, machine learning, and natural language processing. He enjoys exploring the intersection of technology and human behavior, with a strong focus on data driven insights and socially responsible AI applications.

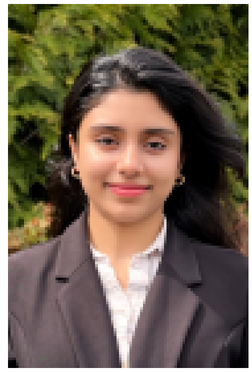

Sanjana Menon is a student at the Pennsylvania State University, majoring in Computational Data Science. Her research interests include natural language processing, sentiment analysis, and computational approaches to human decision-making. She is also interested in the ethical implications of AI, particularly in socially sensitive applications such as evaluation systems.

Sohan Hajra is a student at the University of Illinois at Urbana-Champaign, majoring in Mathematics and Computer Science. His research interests span the development of intelligent financial technologies, including algorithmic trading strategies, fraud detection models, and decentralized finance (DeFi) protocols. He is also passionate about transforming extended reality (XR) through advanced spatial computing and interactive perception algorithms, as well as advancing natural language processing (NLP) for real-world applications in information retrieval, conversational AI, and semantic understanding.

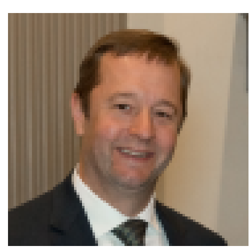

Dr. Jeremy Blum is an associate professor of Computer Science and Chair of the Mathematics and Computer Science programs at the Pennsylvania State University, Harrisburg. Dr. Blum received a D.Sc. in Computer Science and an M.S. in Computational Sciences, both from the George Washington University, as well as a B.A. in Economics from Washington University. His research interests include computer science education and transportation safety.