WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

“Sentiment Analysis using GANs"

Prof. Nandini .C[1], Prof. Priyanka R[2], Adrian Patrick[3] ,Dinesha N [4], Hitan K[5] , Ajay Krishna[6]

Dayananda Sagar Academy of Technology and Management, Bengaluru
(1DT22CA003)
(1DT23CA400)
(1DT22CA023)
(1DT22CA004)

ABSTRACT :

In emotionally intelligent systems, generating text that accurately reflects human sentiment is a critical challenge. Traditional language models focus on grammar and coherence but often fail to express nuanced emotional tones. This project proposes a *Generative Adversarial Network (GAN)* framework — called *EmotionGAN (EGAN)* — combined with *Transformer-based architectures* like T5 to synthesize high-quality, sentiment-labeled text. The system aims to improve natural interaction in AI-driven communication platforms such as chatbots, voice assistants, and human-robot interaction (HRI) modules. The evaluation includes both automated (BLEU, ROUGE, Self-BLEU) and human-centric metrics (empathy, realism).

Introduction

In the age of artificial intelligence, enabling machines to understand and replicate human emotions through natural language has become a critical aspect of human-computer interaction. One of the most prominent areas contributing to this goal is Sentiment Analysis, which involves identifying and interpreting emotional context embedded in textual data.

Traditional models for text generation, such as rule-based systems or statistical NLP approaches, often generate grammatically correct sentences but lack emotional depth and contextual relevance. This limitation significantly affects the development of intelligent agents like chatbots, virtual assistants, and social robots that are expected to engage with users in a more human-like and empathetic manner.

Recent advancements in deep learning, particularly the use of Generative Adversarial Networks (GANs) and Transformer-based architectures such as T5, have introduced new possibilities for generating synthetically rich and emotionally aligned text. GANs are known for their ability to generate human-like outputs by learning adversarial patterns, while Transformers excel at maintaining context and coherence across longer sequences.

This project explores the generation of sentiment-aware text using GANs, specifically focusing on architectures like EGAN (EmotionGAN)—a combination of LSTM-based generators and CNN-based discriminators. Additionally, we investigate the use of T5 Transformers optimized via Neural Architecture Search (NAS) to improve sentiment control and text diversity.

The overall objective is to create a system that can generate text responses reflecting predefined emotions (joy, anger, sadness, etc.) with high coherence, emotional accuracy, and natural flow. Such a system holds great potential in various real-world applications such as emotion-aware chatbots, human-robot interaction, and automated sentiment-based data augmentation for training classifiers.

By integrating emotion conditioning, adversarial training, and transformer-based optimization, this work aims to push the boundaries of emotionally intelligent AI.

Problem Statement

Lack of Emotional Control in Text Generation:

Traditional NLP models (like LSTMs and even some Transformers) often fail to generate text with controlled or desired emotional tone, limiting their effectiveness in applications like emotionally aware chatbots and human-robot interaction.

Data Scarcity and Imbalance:

High-quality, labeled sentiment datasets are limited, costly to create, and often imbalanced, which impacts model performance and generalization across diverse emotional categories.

GAN Training Instability:

Standard GANs suffer from issues such as training instability, mode collapse, and difficulty in generating coherent long-form text, making them challenging to apply directly for sentiment-aware generation.

Inadequate Sentiment Fidelity in Generated Text:

Existing models lack the precision to generate responses that consistently reflect the desired sentiment category (e.g., joy, anger, sadness), leading to outputs that may be emotionally incorrect or inappropriate.

Limited Diversity and Realism:

Sentiment-controlled text generation must ensure both emotional accuracy and linguistic diversity, which current models often fail to balance effectively.

Need for Real-World Integration:

There is a gap between research-level sentiment generation and practical deployment in systems like chatbots or robots, which need real-time, emotionally coherent, and contextually aware text generation capabilities

Objectives

To develop a GAN-based model that generates emotionally expressive text conditioned on specific sentiment labels (e.g., joy, anger, sadness) to enhance human-computer interaction.

To overcome data limitations by creating synthetic, sentiment-labeled datasets using GANs or Transformer-based models optimized through Neural Architecture Search (NAS).

To ensure emotional consistency and linguistic quality in generated text by evaluating models using metrics like BLEU, ROUGE, Jaccard Similarity, and Sentiment Consistency Score.

To compare the performance of GANs and Transformers in sentiment-aware text generation in terms of coherence, diversity, sentiment accuracy, and applicability to real-world scenarios.

To integrate the generated text into applications such as emotion-aware chatbots, educational assistants, and human-robot interaction systems for improved user engagement and experience.

Methodology

The goal of this research is to build a sentiment-aware text generation system using *Generative Adversarial Networks (GANs)* to enable *emotionally intelligent human-robot interaction*. The system should produce grammatically correct, contextually relevant, and sentiment-aligned textual responses.

Dataset Selection and Preprocessing

1. The proposed system requires datasets that contain text annotated with emotional or sentiment labels. The following datasets are considered
2. *IMDB Reviews*: Contains long movie reviews labeled as positive or negative, useful for binary sentiment generation.
3. *Twitter Sentiment140*: A dataset of short texts (tweets) labeled as positive, negative, or neutral, ideal for modeling noisy, real-world sentiment.
4. *DailyDialog*: A high-quality dialogue dataset that includes multiple turns in a conversation and emotional tags such as joy, anger,sadness, etc.

For all selected datasets, the following preprocessing steps are applied:

- *Text Cleaning*: Removal of emojis, HTML tags, stop words, and special characters to normalize the dataset.
- *Tokenization*: Splitting of text into tokens (words or subwords) using standard NLP libraries.
- *Vectorization*: Text is converted to numerical format using pre-trained word embeddings such as *Word2Vec*, *GloVe*, or contextual encoders like *BERT*.
- *Sentiment Tagging*: Labels are preserved or inferred using external sentiment classification tools if necessary.
- *Sequence Padding*: Uniform sequence lengths are ensured using zero-padding to support mini-batch training.

GAN Architecture

The model design follows the principles of *Conditional GANs (cGANs)*. The architecture consists of two neural networks trained adversarially:

3Generator

1. The generator is responsible for producing sentiment-aligned synthetic text. It receives:
2. A *noise vector* (z) for randomness.
3. A *sentiment label* (y) indicating the target emotional tone (e.g., joy, sadness).
4. Optionally, an *input prompt* or dialogue context.
5. The generator may be implemented using:
6. *LSTM/GRU layers* to capture sequential dependencies in text.
7. *Transformer layers* (e.g., T5, GPT) to improve fluency and context retention.

Discriminator

The discriminator evaluates the authenticity and emotional alignment of the text. It is trained to distinguish:

- Real vs. generated text samples.
- Correct vs. incorrect sentiment match.

It may include a *classification head* to predict both realism and emotional correctness.

Conditional Learning

Both G and D are conditioned on sentiment embeddings, which allows the system to control the emotional tone of output text.

Model Training and Optimization

Loss Functions:

- *Adversarial Loss*: Binary Cross Entropy between real/fake labels.
- *Sentiment Consistency Loss*: Penalty for sentiment mismatch using a pre-trained sentiment classifier.
- *Optional*: BLEU or ROUGE loss to encourage coherence.

Training Strategy:

- Generator and Discriminator are updated alternately.
- Techniques like *label smoothing*, *gradient clipping*, and *adaptive learning rates* are used to prevent mode collapse and instability.

Epochs and Batch Size:

- Hyperparameters are determined empirically via validation performance.

System Workflow

The following pipeline is implemented to ensure real-time usability and modularity:

1. The system receives a *user input text*.
2. A *sentiment detection module* classifies the emotional tone of the input.
3. The sentiment label and input prompt are passed to the *GAN Generator*.
4. The Generator outputs a synthetic response aligned with the sentiment.
5. The Discriminator validates the realism and sentiment alignment of the output.

- The final response is forwarded to the *robot's dialogue interface* or a chatbot front-end.
- The system can be optionally extended to *multi-modal inputs* such as camera-based emotion detection or speech tone analysis.

Evaluation Metrics

To assess the model's performance, both quantitative and qualitative metrics are employed:

1. *BLEU* and *ROUGE*: Evaluate grammatical fluency and textual similarity.
2. *Self-BLEU*: Measures the diversity of generated outputs.
3. *Jaccard Similarity*: Compares overlap of sentiment-specific tokens between reference and generated texts.
4. *Sentiment Consistency Score*: Uses an independent classifier to verify that the generated output reflects the intended sentiment.
5. *Human Evaluation (optional)*: Responses are rated by humans on scales such as empathy, fluency, and emotional correctness.

Deployment Considerations

- Initial deployment focuses on *text-only input/output*, suitable for conversational agents and chatbots.
- Planned extensions include integration with *robotic speech interfaces*, *voice assistants*, and *camera-based facial emotion detection* for richer interaction.
- The system is designed to be *modular*, allowing future upgrades with multimodal data streams or new GAN variants.

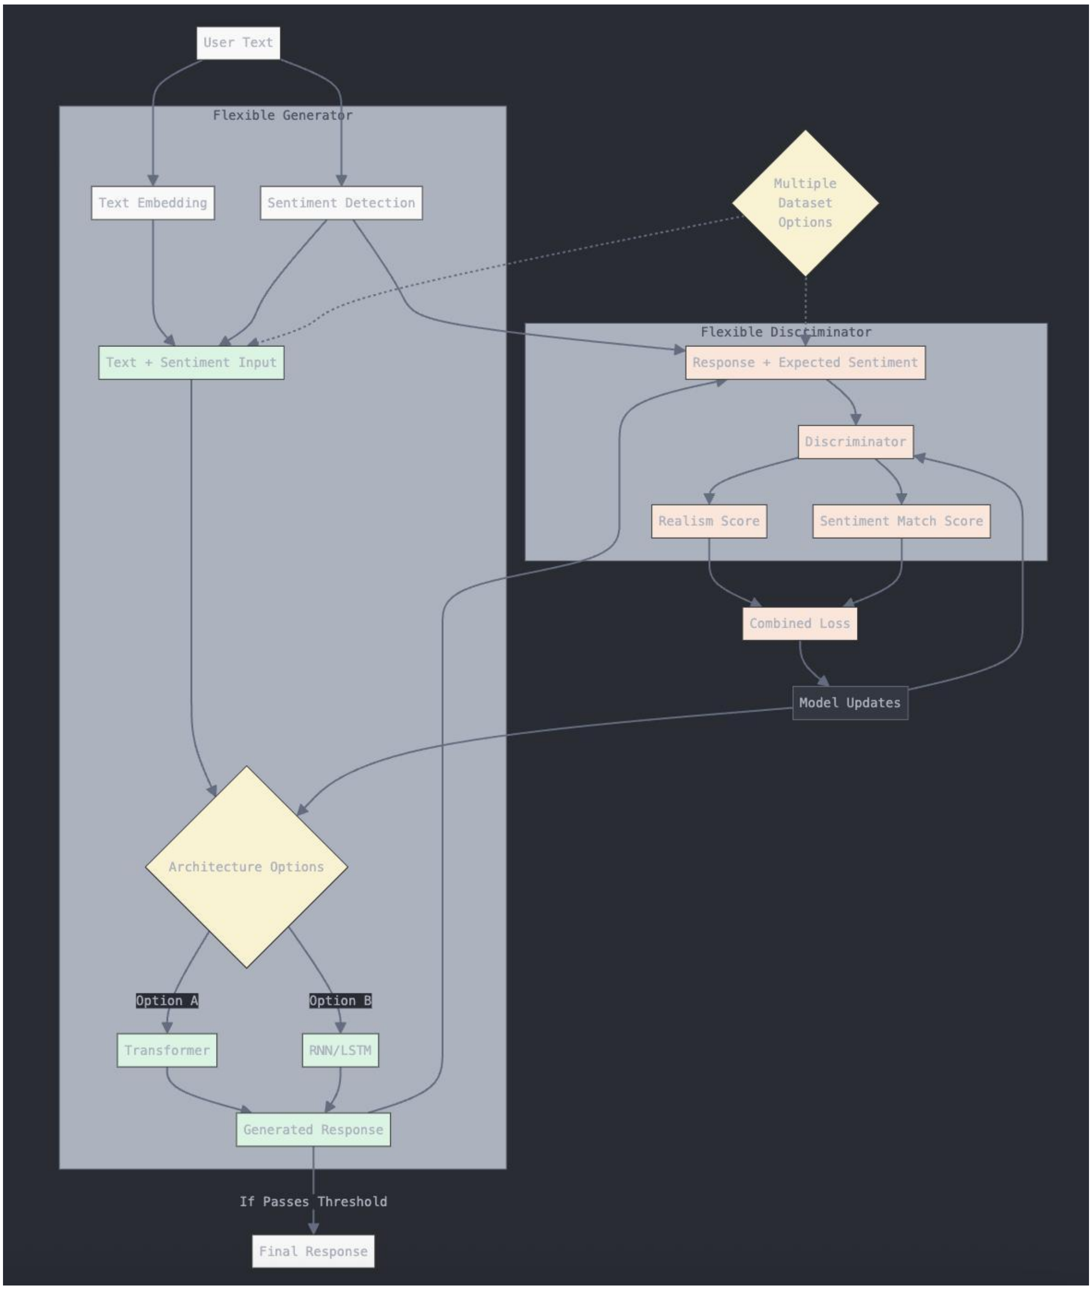

Results:

Anticipated Outcomes:

- Generated responses will align with the detected sentiment (positive, negative, neutral).Text will be contextually relevant and grammatically correct

Evaluation Metrics (Planned):

- BLEU / ROUGE: To assess fluency and coherence.
- Sentiment Consistency Score: Measures how closely the generated responsereflectsintended emotion.
- Human Evaluation (optional later): Response quality, empathy, and relevance.

Expected Impact on Human-Robot Interaction:

- Robots will respond more naturally and empathetically.
- Emotional alignment enhances trust and engagement.

Potential Application Areas:

- Healthcare: Emotional support robots for mental health.
- Education: Encouraging learning assistants.

- Customer Service: Emotion-aware chatbot agents.

Conclusion

This research presents a novel approach to sentiment-aware text generation using Generative Adversarial Networks (GANs), specifically Conditional GANs (cGANs), to address the limitations of traditional natural language generation models in capturing emotional context. The proposed system demonstrates the effectiveness of GAN-based architectures in generating emotionally coherent and contextually appropriate text, paving the way for emotionally intelligent human-computer interaction. The flexibility of the architecture allows seamless integration with future enhancements, including multimodal emotion detection through text, voice, and visual inputs. Moreover, this work lays a strong foundation for future advancements in affective computing and emotionally responsive dialogue systems, with promising applications in healthcare, education, and customer service domains.

REFERENCES

1. Hempel, A., Smith, J., & Lee, T. (2023). Sentiment-based engagement models for human-robot interaction. *Journal of Human-Robot Interaction, 12*(1), 45–62
2. ScienceDirect. (2022). Generative adversarial networks for sentiment text generation: A review. *ScienceDirect Article*.
3. ScienceDirect. (2024). Conditional GANs for sentiment classification on Twitter data. *ScienceDirect Article*.
4. Lu, Y., Chen, H., & Zhang, L. (2023). Emotion and topic-aware dialogue generation for natural human-robot interaction. *Transactions on Affective Computing*, 14(2), 310–323.