

RESEARCH ARTICLE

Socially Aware Synthetic Data Generation for Suicidal Ideation Detection Using Large Language Models

HAMIDEH GHANADIAN¹, ISAR NEJADGHOLI², AND HUSSEIN AL OSMAN¹, (Member, IEEE)

¹Department of Electrical Engineering and Computer Science, University of Ottawa, Ottawa, ON K1N 6N5, Canada

²National Research Council Canada, Ottawa, ON K1V 1J8, Canada

Corresponding author: Hamideh Ghanadian (Hghan053@uOttawa.ca)

ABSTRACT Suicidal ideation detection is a vital research area that holds great potential for improving mental health support systems. However, the sensitivity surrounding suicide-related data poses challenges in accessing large-scale, annotated datasets necessary for training effective machine learning models. To address this limitation, we introduce an innovative strategy that leverages the capabilities of generative AI models, such as ChatGPT, Flan-T5, and Llama, to create synthetic data for suicidal ideation detection. Our data generation approach is grounded in social factors extracted from psychology literature and aims to ensure coverage of essential information related to suicidal ideation. In our study, we benchmarked against state-of-the-art NLP classification models, specifically, those centered around the BERT family structures. When trained on the real-world dataset, UMD, these conventional models tend to yield F1-scores ranging from 0.75 to 0.87. Our synthetic data-driven method, informed by social factors, offers consistent F1-scores of 0.82 for both models, suggesting that the richness of topics in synthetic data can bridge the performance gap across different model complexities. Most impressively, when we combined a mere 30% of the UMD dataset with our synthetic data, we witnessed a substantial increase in performance, achieving an F1-score of 0.88 on the UMD test set. Such results underscore the cost-effectiveness and potential of our approach in confronting major challenges in the field, such as data scarcity and the quest for diversity in data representation.

INDEX TERMS Artificial intelligence, deep learning, large language models, suicide detection, synthetic data generation, transformer based models.

I. INTRODUCTION

According to the World Health Organization¹ more than 700,000 people die due to suicide every year. Suicide remains a global health crisis, accounting for a significant proportion of mortality rates across various age groups. Suicidal ideation, often a precursor to actual suicide attempts, involves the presence of persistent thoughts, contemplation, or planning related to self-harm or death. Early identification of suicide ideation and intervention to protect individuals at risk of suicide are crucial steps in reducing suicide rates and

The associate editor coordinating the review of this manuscript and approving it for publication was Vivek Kumar Sehgal¹.

¹The World Health Organization (WHO).

providing appropriate mental health support. Early detection of suicidal ideation is a complex task, as it requires the integration of various factors, including psychological, social, and environmental variables [1].

In recent years, the proliferation of digital platforms and social media has provided an unprecedented opportunity to capture and analyze large-scale data related to mental health [2], [3]. Machine learning and Natural Language Processing (NLP) techniques have shown promise in detecting linguistic patterns and indicators of suicidal ideation in diverse text-based data sources, such as social media posts, online forums, and electronic health records [4], [5], [6], [7].

However, the use of machine learning technologies requires high volumes of data. Data collection and annotation

processes are time-consuming and impose significant financial costs [8]. Specifically, obtaining a substantial amount of labeled data related to suicide can be challenging and limited due to several factors inherent to the nature of suicide research. The sensitive and stigmatized nature of suicide often presents barriers to data collection. Individuals and organizations may be rightfully reluctant to share personal or confidential information related to suicide, fearing potential negative consequences.

Synthetic data generation offers a viable solution to mitigate the data availability limitation by creating artificially generated data that closely resembles real-world data. Synthetic data generation can be instrumental in machine learning applications as it addresses many challenges of real data collection and annotation. Here we review a list of common challenges in data collection that can be managed through synthetic data generation.

A. DATA SCARCITY

In many NLP tasks, such as mental health-related applications, there may be limited availability of relevant data due to privacy concerns or the complexity and cost associated with manual annotation. Synthetic data generation allows researchers and practitioners to overcome data scarcity issues and augment the limited amount of publically available data [9].

B. DATA DIVERSITY

NLP models trained on limited data may suffer from poor generalization and performance when exposed to diverse and previously unseen examples. Moreover, in real data, certain topics can be undermined or overlooked due to being less discussed. This can happen for several reasons. For example, certain topics may be stigmatized and considered too sensitive or taboo, making people hesitant to openly discuss them. This could include subjects related to mental health, addiction, discrimination, or social issues that carry societal stigmas. Additionally, topics relevant to marginalized or minority communities may receive less discussion due to systemic biases, unequal representation, or limited platforms for their voices to be heard. Also, some topics may be highly specialized or complex, requiring specific expertise or background knowledge to engage in meaningful discussions. Encouraging diverse perspectives and actively seeking out less-discussed topics can contribute to a more comprehensive and nuanced understanding of real-world issues. Synthetic data generation can help enrich the training data by introducing a wider variety of linguistic patterns, sentence structures, vocabulary and topics. This, in turn, improves the model's ability to handle variations in natural language and increases its robustness [10].

C. PRIVACY PRESERVATION

Suicide detection tasks often involve sensitive information. Generating synthetic data allows researchers to create

representative samples that preserve the privacy of individuals while maintaining the statistical properties and distribution of the original data [11].

D. ANNOTATION COST

Suicide detection is a complex task, and high-quality annotation can only be performed by experts and trained annotators, which can be costly [12]. Synthetic data generation addresses the data annotation issue by targeted data generation so that each generated example is pre-labeled with a specific category. Although these labels might be noisy to some extent, they might be preferable in some settings as they come at no additional cost.

To investigate the feasibility and effectiveness of synthetic data generation in the task of suicide ideation generation, we use Generative Large Language Models (GLLMs) for data synthesis and use the generated data to train/test text classifiers. To train classifiers, we fine-tune pretrained BERT-like Large Language Models (LLMs) as state-of-the-art text classifiers.

To enhance the quality of the generated data, we benefit from domain knowledge from psychology. Previous research highlights the importance of incorporating social factors in the design process of NLP systems [13]. Specifically, when generating data with LLMs, external sources of domain knowledge can be leveraged to guide the data generation process [14]. For the task of suicidal ideation detection, such knowledge can be drawn from a vast body of research in psychology devoted to gaining an understanding of the social factors associated with suicidal ideation and behavior. In this work, we review the psychology literature to extract the social factors tightly tied to suicidal ideation and use this knowledge for more effective prompt engineering when generating data with GLLMs. Guiding the data generation with these factors enables the creation of diverse and representative examples of suicidal ideation.

The main contributions of this study are as follows:

- We extracted the relevant social factors associated with suicidal ideation through a comprehensive review of existing literature, research papers and clinical studies to identify key themes related to suicidal ideation. These themes encompass a wide range of factors, including risk factors, common triggers and mental health indicators. Leveraging a socially aware data synthesis approach, we pave the way for more accurate and reliable suicidal ideation detection systems.
- Our study examines three GLLMs' performance in producing synthetic datasets with Zero-Shot and Few-Shot learning techniques. Utilizing the ChatGPT, Flan-T5, and Llama 2 models and leveraging the extracted social factors from the psychology literature, we generated nine datasets with diverse characteristics.
- We trained classifiers by fine-tuning two pre-trained language models, ALBERT and DistilBERT, using the generated datasets. We tested these models on two test sets, the University of Maryland Suicidality dataset (UMD)

and a human-annotated synthetic dataset presented in this paper. Our findings indicate that the GLLMs have significant potential for generating a suicide-related dataset comparable with real available datasets such as UMD. More significantly, the integration of social knowledge may significantly enhance the quality of the generated datasets and lead to more robust classifiers.

- We augmented our best-performing synthetic dataset using subsets of the UMD dataset to evaluate the efficacy of data augmentation in suicidal detection applications. Our results show that models trained with synthetic data augmented with a small set of real-world data can outperform models trained by large annotated real-world datasets.

This paper is organized as follows: In Section II, we review the literature and related background. In Section III, we explain our methodology for generating and evaluating the proposed synthetic data generation, specifics of the classifiers and datasets we use in this work. Section IV presents our results, and Section V discusses these results in detail. Additionally, the conclusion and possible future works are discussed in Section VI. We complete the article by including an ethical statement in Section VII, which delves into the ethical aspects and considerations associated with our work.

II. BACKGROUND AND RELATED WORK

In this section, we review the related work in suicidal ideation detection in psychology as well as NLP research that addresses the task of suicide detection. We also review the previous works that focused on generating synthetic data for a variety of NLP tasks.

A. SUICIDAL IDEATION AND RELATED SOCIAL FACTORS

Suicidal ideation has been a subject of extensive research within the field of psychology. Understanding the underlying elements and risk factors related to suicidal thoughts and behaviors is crucial for developing effective prevention and intervention strategies.

One important area of investigation is the identification of risk factors associated with suicidal ideation. Numerous studies have examined the impact of psychological factors such as depression, anxiety, hopelessness, and feelings of worthlessness on the development of suicidal thoughts [15], [16]. These studies investigate the strong association between suicidal thoughts and conditions like depression [17], [18], bipolar disorder [19], borderline personality disorder [20], and substance abuse [21]. By examining the interplay between these conditions, researchers aim to develop targeted interventions to address the unique challenges faced by individuals struggling with suicidal ideation [22], [23]. Additionally, environmental factors such as a history of trauma, social isolation, and access to lethal means have been identified as potential risk factors [24], [25], [26].

Psychology offers valuable insights into the diverse processes and factors that contribute to suicide risk. Psychological theories and frameworks such as the

interpersonal theory of suicide [27], the cognitive model of suicidal behavior [28], and the social-ecological model [29] provide a theoretical foundation for understanding the complex interplay between individual vulnerabilities and environmental factors.

The extensive research conducted on suicidal ideation and associated topics in psychology has significantly contributed to the understanding of the complex factors involved. By unraveling the causes, risk factors, and protective factors associated with suicidal thoughts, researchers aim to develop effective prevention strategies, enhance mental health interventions, and ultimately reduce the global burden of suicide. In Section III-A, we enumerate the social factors that are discussed in the literature as relevant topics to suicidal ideation.

B. SUICIDAL IDEATION DETECTION USING NLP

In recent years, there has been a growing interest in using NLP techniques for suicide prevention [30], [31]. Researchers have developed suicide detection systems to analyze and interpret social media data, including text data. By detecting linguistic markers of distress and other risk factors, these systems can help identify individuals with a risk of suicidality and provide early interventions to prevent such incidents [32].

Several studies indicated the impact of social network reciprocal connectivity on users' suicidal ideation. Hsiung [33] analyzed the changes in user behavior following a suicide case that occurred within a social media group. Jashinsky et al. [34] highlighted the geographic correlation between suicide mortality rates and the occurrence of risk factors in tweets. Colombo et al. [35] focused on analyzing tweets that contained suicidal ideation, with a particular emphasis on the users' behavior within social network interactions that resulted in strong and reciprocal connectivity, leading to strengthened bonds between users. NLP techniques, therefore, offer a promising avenue for suicide prevention efforts, enabling more proactive and effective interventions to support those in need.

1) GENERATIVE LANGUAGE MODELS

Ghanadian et al. [36] utilized ChatGPT for assessing suicidality from social media posts. They performed Zero-Shot and Few-Shot experiments and extensive performance comparison between ChatGPT and two fine-tuned transformer-based models. They also investigated the impact of different temperature parameters on ChatGPT's response generation. The findings of this paper suggest that ChatGPT achieves notable accuracy in the suicidal risk assessment task; however, transformer-based pre-trained models fine-tuned on human-annotated datasets exhibit superior performance. Furthermore, the analysis provides insights into adjusting ChatGPT's hyperparameters to enhance its effectiveness in assisting mental health professionals with this critical task.

Yang et al. [37] conducted a comprehensive evaluation of ChatGPT's mental health analysis and emotional reasoning

ability across five tasks. They also investigated the impact of different emotion-based prompting strategies. Additionally, they explored the use of generative models to generate explanations for the decisions made by ChatGPT, aiming for interpretable mental health analysis. The experimental results revealed that ChatGPT performed better than traditional neural network-based methods such as Convolutional Neural Network (CNN) and Gated Recurrent Unit (GRU) in mental health analysis but still lagged behind advanced task-specific methods.

2) AVAILABLE DATASET

Several datasets have been collected from social media platforms to serve as a resource for creating suicidal ideation detection systems. These datasets encompass a wide range of information collected from various social media sources, including *Twitter*, *Reddit* and other user-generated content. Sinha et al. [38] created a manually annotated dataset from *Twitter* using a lexicon of suicidal phrases and a lexicon along with the social engagement data associated with real-time and historical tweets. The resulting dataset consists of 34,306 tweets with two labels, *Suicidal* and *Non-Suicidal*. Gaur et al. [39] collected and annotated a 5-label Suicide Risk Severity Assessment dataset from *Reddit*, which includes *Suicidal Ideation (ID)*, *Suicidal Behavior (BR)*, *Actual Attempt (AT)*, *Suicide Indicator (IN)* and *Supportive (SU)* categories. This dataset is extracted from *SuicideWatch*² subreddit and has been annotated by four practicing clinical psychiatrists, ensuring the accuracy and reliability of the annotations. The dataset comprises a total of 500 posts, which have been carefully selected to represent a diverse range of content related to suicidal ideation.

Another widely referenced dataset in the field of suicidal ideation detection is the University of Maryland *Reddit Suicidality Dataset (UMD)* [40], [41]. The UMD dataset is a collection of *Reddit* posts and comments created by individuals who expressed suicidal thoughts or behaviors. The dataset contains over 100,000 posts and comments collected from various subreddits, including those related to mental health and suicide prevention, such as *Depression*³ and *SuicideWatch* subreddits. The data was collected over a period of several years and includes the content of the posts and comments, as well as the location and timing of the posts.

C. SYNTHETIC DATA COLLECTION

To overcome the limitations of real-world data availability, NLP researchers have explored the use of synthetic datasets for several applications. For example, He et al. [42] utilized language models to generate synthetic unlabeled text. They introduced the Generate, Annotate, and Learn (GAL) framework that leverages synthetic text for knowledge distillation, self-training, and few-shot learning purposes. To generate the data, they fine-tune pre-trained language

²SuicideWatch subreddit.

³Depression subreddit.

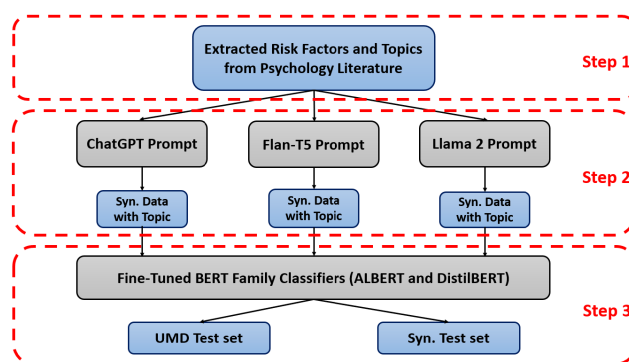


FIGURE 1. Workflow of the proposed methodology.

models on relevant datasets with limited examples. The synthetic text is then annotated with soft pseudo labels using the best available classifier for knowledge distillation and self-training. This paper achieves state-of-the-art results for knowledge distillation with 6-layer transformers on the GLUE leaderboard [43].

Bonifacio et al. [44] presents an effective approach to leverage LLMs in retrieval tasks, resulting in significant improvements across various datasets. Instead of directly utilizing LLMs during the retrieval process, they harness the LLMs' capabilities to generate labeled data using a few-shot learning approach. Subsequently, they fine-tune smaller retrieval models on this synthetic dataset and employ them to re-rank the search results obtained from a primary retrieval system. They provide a novel method to adapt LLMs for Information Retrieval (IR) tasks that were previously deemed infeasible due to their demanding computational requirements. By shifting the computational burden from the retrieval stage to the generation of synthetic data for training, they make it feasible to exploit the power of LLMs without compromising efficiency. In an unsupervised setting, their approach significantly outperforms recently proposed methods, highlighting its superiority in terms of retrieval performance and scalability.

III. METHODOLOGY

In this section, we elaborate on our proposed methodology. Figure 1 shows the workflow we use to generate synthetic datasets and our testing process. As shown in this figure, our method has three steps:

- **STEP 1- Domain knowledge Extraction:** Extract relevant social factors from the psychology literature for an informed prompting of GLLMs in data synthesis.
- **STEP 2- Synthetic Data Generation:** Use three GLLMs to generate socially aware synthetic data, that is, data that covers a wide range of suicide-related topics.
- **STEP 3- Evaluate the effectiveness of Synthetic data:** Train state-of-the-art classifiers with real-world, synthetic, and augmented datasets and test those classifiers on real-world as well as synthetic test sets.

In the following, we explain the three steps described above. The complete implementation of our project, including Zero-Shot Learning and Few-Shot Learning of GLLMs, as well as the fine-tuned classifiers, is available on GitHub.⁴

A. SUICIDE RELATED TOPICS IN PSYCHOLOGY

We conducted a comprehensive search across various academic databases, including PsycINFO, PubMed, and Google Scholar, with keywords and combinations such as “suicide”, “suicidal ideation” and “psychology” to identify relevant articles, research papers, and review papers. Thematic analysis was employed to identify the most recurring and significant topics across the included studies. Repeated topics that demonstrate relevance to suicide in psychology were considered the most related topics.

Based on our analysis of the literature, the following social and psychological factors were consistently reported in relation to suicidal ideation in psychology. These topics are not listed in a specific order of importance but represent the consistently reported themes in the literature reviewed:

1) DEPRESSION

Depression emerged as a frequently reported topic, highlighting its strong association with suicidal ideation. Numerous studies have explored the relationship between depressive symptoms, including sadness, loss of interest, feelings of worthlessness, and the increased risk of suicidal thoughts [45], [46], [47].

2) ANXIETY

Anxiety disorders were also commonly associated with suicidal ideation. Research has emphasized the link between excessive worry, fear, and agitation and the presence of suicidal thoughts and behaviors [48].

3) UNEMPLOYMENT

The experience of unemployment has been consistently identified as a topic closely related to suicidal ideation. Studies have examined the psychological distress and negative impact on self-esteem and social support that can arise from unemployment, contributing to suicidal ideation [49].

4) HOPELESSNESS

Hopelessness, characterized by a lack of optimism and a perceived absence of future prospects, has been consistently linked to suicidal ideation. Studies have demonstrated the significant role of hopelessness as a predictor of suicidal thoughts [45], [46], [49].

5) ANGER

The expression and experience of anger have been reported as influential factors in suicidal ideation. Unresolved anger,

hostility, and intense emotional distress have been associated with an increased risk of suicidal thoughts [45].

6) PERFECTIONISM

Perfectionism, marked by excessively high standards and self-criticism, has been identified as a psychological factor related to suicidal ideation. Research has explored the relationship between perfectionistic tendencies and the development of suicidal thoughts and behaviors [45].

7) FAMILY ISSUES

Family-related issues, such as conflict, dysfunctional dynamics, and poor communication, have consistently emerged as topics associated with suicidal ideation. These factors can contribute to a sense of isolation, distress, and feelings of being a burden, increasing the risk of suicidal thoughts [47], [49].

8) RELATIONSHIP PROBLEMS

Difficulties in intimate relationships, including conflicts, breakups, and marital dissatisfaction, have been reported as significant topics in relation to suicidal ideation. Relationship problems can contribute to emotional distress and feelings of hopelessness, leading to thoughts of suicide [49].

9) FINANCIAL CRISIS

Financial difficulties and crises have been consistently linked to suicidal ideation. Economic stressors, such as debt, unemployment, and financial insecurity, can contribute to psychological distress and an increased risk of suicidal thoughts [48].

10) EDUCATION

Issues related to educational pressures, academic stress, and performance expectations have been reported as topics associated with suicidal ideation. Research has highlighted the impact of academic-related stressors on mental well-being and the risk of suicidal thoughts among students [50].

11) BULLYING

Bullying, including physical, verbal, or cyber-bullying, has consistently emerged as a significant topic related to suicidal ideation. The experience of bullying can lead to social isolation, low self-esteem, and emotional distress, contributing to the development of suicidal thoughts [48].

12) DEATH OF LOVED ONES

The loss of close family members or friends through death has been reported as a topic associated with suicidal ideation. Grief, feelings of loneliness, and a sense of being unable to cope with the loss can increase the risk of suicidal thoughts [51].

13) IMMIGRATION

Issues related to immigration, discrimination, and racism have been identified as topics linked to suicidal ideation.

⁴https://github.com/Hamideh-ghanadian/Synthetic_Data_Generation_using_Generative_LLMs

Experiences of marginalization, social exclusion, and acculturative stress can contribute to psychological distress and suicidal thoughts among individuals facing these challenges [52], [53].

14) RACISM

Studies have consistently highlighted the significant impact of racial discrimination on suicidal ideation. Experiencing racism and racial prejudice can increase the risk of suicidal thoughts [54].

B. SYNTHETIC DATA GENERATION

We utilized three generative language models to generate a synthetic dataset related to suicidal ideation. GLLM's foundation is constructed with transformers. Transformers are a class of deep learning models, first introduced by Vaswani et al. [55] in 2017. Researchers build state-of-the-art NLP models using transformer-based architectures because they can be quickly trained on large datasets, and studies have shown that they are better at modeling long-term dependencies in natural language text [56]. GLLMs, including ChatGPT, FlanT5, and Llama are designed with the primary purpose of generating coherent and contextually relevant text. They excel at tasks such as text generation [57], completion [58], and dialogue generation [59]. These models are typically based on decoder transformer architectures and focus on the generative aspect of language which involves the auto-regressive generation, where the models predict the next word based on the preceding context. Generative models are trained on a vast corpus of text, however, their main strength lies in their ability to generate text that flows naturally and contextually appropriate.

We aim to build a diverse dataset in order to train a generalizable and robust model in suicidal ideation detection. In total, nine different datasets are generated with different specifications and models.

1) ChatGPT

The language model utilized by ChatGPT is *gpt-3.5-turbo*,⁵ which is one of the most advanced language models developed by OpenAI. ChatGPT accept a sequence of messages as an input and produce a message generated by the model as an output. Although the chat format is primarily intended for conversations spanning multiple turns, it is also equally useful for single-turn tasks that do not involve any conversations. We used the *OpenAI Python library*⁶ to access the *ChatCompletion* functionality of the *gpt-3.5-turbo* model through its API.

In this project, we evaluate the capability of ChatGPT in Zero-Shot Learning and Few-Shot Learning settings to generate a diverse suicidality dataset. However, we are primarily focused on Zero-Shot Learning methods as ChatGPT has exhibited superior performance in this setting

compared to Few-Shot Learning for a suicidal ideation detection task. Ghanadian et al. [36] conducted an extensive comparison of the Zero-Shot and Few-Shot approaches using ChatGPT. According to their findings, fine-tuning a model on Few-Shot setting might yield poorer performance compared to Zero-Shot in various scenarios. In Few-Shot, with few examples available for fine-tuning, the risk of overfitting increases. The model might learn specific nuances or noise within the limited few-shot data, leading to poor generalization on unseen examples. Moreover, Few-shot learning relies on a small subset of labeled examples, which might not adequately represent the entire diversity of the dataset. The model might fail to capture the complexity and variability present in the broader dataset during fine-tuning.

The temperature hyperparameter in ChatGPT is a crucial parameter that influences the generated output. A higher temperature value, such as 1.0, increases the randomness and produces more varied responses. Conversely, a lower temperature value, such as 0.1, reduces randomness and generates more focused and deterministic responses. Ghanadian et al. [36] investigated the effect of the temperature parameter on the generated output of ChatGPT for suicide risk assessment. Furthermore, the authors introduced a parameter known as the "Inconclusiveness Rate," which indicates the proportion of test cases that do not produce a definitive or conclusive result. According to their paper, this parameter decreases as the temperature parameter is increased. As such, for this project, we have configured the temperature parameter of ChatGPT to be 1.

We generated five datasets using ChatGPT. Four of these datasets are informed by 14 main suicide-related topics in psychology, while one dataset is generated without providing any specific topics. For incorporating suicide-related topics in data generation, we utilized prompt engineering techniques. Prompt engineering involves carefully crafting and designing the prompts provided to the model to elicit desired responses. By employing prompt engineering strategies, we aimed to enhance the quality and relevance of the generated dataset. This methodology allowed us to tailor the dataset generation process to align with our specific objectives and requirements. For prompt engineering, we drew inspiration from a short course on ChatGPT Prompt Engineering.⁷

In both Zero-Shot and Few-Shot Learning, we initiated the prompt construction process with a simple initial prompt and iteratively refined it through multiple rounds of trial and error. This iterative approach allowed us to gradually evolve the prompt, making necessary adjustments based on the observed outcomes. The following example shows a prompt we used to generate a binary dataset (*Suicide/non-suicide*), guided by the topics introduced in Section III-A, in Zero-shot Learning setting.

⁵<https://platform.openai.com/docs/models/gpt-3-5>

⁶<https://github.com/openai/openai-python>

⁷ChatGPT Prompt Engineering for Developers.

```
Criteria="" These are the criteria of
different suicide risk level:
Risk Level=Non Suicidal: I do not see
evidence that this person is at risk
for suicide
Risk Level=Suicidal: I believe this
person is at high risk of attempting
suicide in the near future. ""
```

```
prompt = f""
Your task is to generate a suicidal text for
each of the following "topics" with
different Risk levels.
```

```
1-Depression
2-Anxiety
3-Hopelessness
4-Anger
5-Perfectionism
6-Family issues
7-Relationship problems
8-Unemployment
9-Financial Crisis
10-Education
11-Being Bullied
12-Death of closed one
13-Immigration
14-Racism
```

```
Provide the answers in JSON format with the
following columns: text, topic, risk
level.
```

```
Risk level criteria: ``{ Criteria }``
```

```
""
```

In Few-Shot Learning, the prompt is structured to include two examples for each category (8 in total) from the training set of UMD Dataset, followed by a text generation question. This approach enables the model to learn from a limited set of labeled examples before generating a dataset. Moreover, by combining the Few-Shot Learning methodology with the inclusion of psychology topics in the prompt, we aim to enhance the model's ability to generate meaningful and contextually relevant responses when dealing with suicide-related discussions.

2) FLAN-T5

FLAN-T5 models are instruction fine-tuned across a diverse set of tasks, aiming to enhance their zero-shot performance on various tasks. During instruction tuning, pretrained models undergo fine-tuning using drafts of instructions that guide them on how to perform a specific task. These instructions can include real-time feedback to assist the model in learning from its mistakes and improving at a faster rate. By providing explicit guidance and incorporating feedback mechanisms, the instruction-tuning process enables the model to refine its performance and enhance its ability to accurately execute the given task. This iterative approach of incorporating instructions and feedback facilitates the model's learning process, allowing it to adapt and improve its performance based on the provided guidance.

In this project, we utilized *Flan-T5-XXL*⁸ presented by Google Research [60] in a Zero-Shot setting. Two datasets are generated using Flan-T5, one with topics and another without topics. Moreover, similar to ChatGPT, the temperature value is set to 1, and the same prompt structure is utilized.

3) LLAMA 2

LLaMA (Large Language Model Meta AI) is an auto-regressive language model constructed based on transformer architecture. Similar to other generative models, LLaMA operates by taking a sequence of words as its input and making predictions about the subsequent word, iteratively producing text in a recursive manner. It is a collection of state-of-the-art foundational language models, with parameter counts ranging from 7 billion to 65 billion. The foundation models were trained on large unlabeled datasets, making them ideal for fine-tuning on a variety of tasks. The newest version of this model, Llama 2, expanded its pre-training corpus size, allowing the model to learn from a more extensive and diverse set of publicly available data. Additionally, the context length of Llama 2 has been doubled, enabling the model to consider a more extensive context when generating responses, leading to improved output quality and accuracy [61]. In this paper, we used Llama 2-13B, presented by Meta in the Zero-Shot setting. In total, we generated two datasets with Llama2, one with topics and another without topics. These datasets were created using the temperature of 1 and maintained the same prompt structure as ChatGPT.

C. EVALUATION OF SYNTHETIC DATASET

To evaluate the utility and effectiveness of synthetic datasets, we fine-tuned pre-trained transformer-based language models, ALBERT and DistilBERT, to train classifiers with each set of the generated synthetic data. We compared the trained classifiers with classifiers with similar structures fine-tuned with real-world data as the benchmark model.

ALBERT and DistilBERT are two pre-trained language models from the BERT family of LMs. The BERT model was initially proposed by Devlin et al. [62] as a bidirectional language model pretrained on a large corpus comprising the Toronto Book Corpus and Wikipedia. The model is named bidirectional because it can simultaneously gather the context of a word from either direction. Unlike the generative models such as ChatGPT, FlanT5 or Llama, which include a decoder structure, the BERT family of language models are encoder models and can be fine-tuned for specific tasks such as classification tasks.

The ALBERT model was proposed by Lan et al. [63] to reduce memory consumption and increase the training speed compared to BERT. In other words, ALBERT is a more lightweight version of BERT that maintains its high level of accuracy, making it a powerful tool for various NLP applications. The DistilBERT model was proposed by Sanh et al. [64]. The authors reported it has 40% fewer

⁸<https://huggingface.co/google/flan-t5-xxl>

parameters than BERT and runs 60% faster while preserving over 95% of BERT's performances as measured on the GLUE language understanding benchmark. Both models are designed as lightweight alternatives to BERT, with ALBERT emphasizing parameter efficiency and DistilBERT focusing on knowledge transfer through distillation. Overall, ALBERT, with a smaller number of parameters, shows more efficient performance compared to DistilBERT.

To fine-tune these models, we utilized the Huggingface library [65]. The Huggingface is an open-source library and data science platform that provides tools to build, train and deploy ML models. We compare our classification results with baseline ALBERT⁹ and DistilBERT¹⁰ models fine-tuned on the UMD dataset by Ghanadian et al. [36]. We used the Trainer¹¹ class from Huggingface transformers¹² for feature-complete training in PyTorch.

The hyperparameters were selected based on the default values commonly used in similar studies. The final hyperparameters used in our experiments were Learning Rate= $2e^{-5}$, Batch Size = 4, Dropout Rate = 0.1, and Maximum Sequence Length = 512. By comparing the performance of these models on synthetic datasets against the baseline, we can assess the efficiency of using the synthetic datasets and gauge the improvements achieved through our fine-tuning process.

To conduct a comprehensive assessment of the fine-tuned classifiers' performance, we generated two distinct sets of testing subsets. Furthermore, we created an augmented dataset to showcase the application of synthetic data in the suicidal ideation detection domain.

1) TESTING SUBSETS

We selected two test sets for evaluation purposes: The first test dataset is the test subset of the UMD datasets utilized in [36], which is annotated as a 4-class dataset. We employed a 10-20-70 split for validation, test, and training sets, respectively. Out of the entire dataset, 10% was allocated for validation purposes, ensuring the model's hyperparameters and configurations were appropriately set. 20% of the data was set aside as a test set to evaluate the model's performance on unseen data and ensure its generalizability. The remaining 70% formed the training set, where the bulk of the data was utilized to train the model and learn the underlying patterns. This distribution was chosen to provide substantial data for training while reserving enough distinct data for validation and robust performance testing. The detailed description of the Multi-class UMD dataset is presented in Table 1.

Furthermore, to employ binary classification, we binarize the UMD Dataset. Based on the definition of each class, "No Risk" and "Low Risk" classes are considered as Non-Suicidal and "Moderate Risk" and "High Risk" as

TABLE 1. The description of the training and testing subset of UMD dataset used in [36] for multi-class task.

Multi-class Dataset	No Risk	Low Risk	Moderate Risk	High Risk
Training Subset	27.45 %	16.39 %	31.90 %	24.24 %
Number of Users	154	92	179	136
Testing Subset	24.41 %	11.62 %	26.74 %	37.20 %
Number of Users	42	20	46	64

TABLE 2. The description of the training and testing subset of UMD dataset for binary task.

Binary Dataset	Non Suicidal	Suicidal
Training Subset	43.84%	56.14%
Number of Users	246	315
Testing Subset	36.3	63.94
Number of Users	62	110

Suicidal. Table 2 presents the description of binarized UMD dataset. The second testing set is composed of 10% of each synthetic dataset generated in our project. This test set is annotated independently by two human annotators. A notable 89% of the labels, initially generated by the generative models, were agreed upon by the human annotators. However, for the remaining 11% of the data, the labels were altered based on the decision of the annotators. In cases where both annotators agreed on a label, that label was retained. Conversely, when disagreements arose, the annotators engaged in discussions to ultimately reach a consensus on the appropriate label. The details of the generated synthetic dataset are presented in Table 3 as the eleventh dataset.

2) AUGMENTED DATASET

Data augmentation involves enriching a dataset by introducing variations to its existing instances or generating entirely new instances. This process is designed to enhance the diversity and quality of the dataset, which, in turn, can lead to improved model performance and generalization. Hence, in this study, we augment the best performing synthetic dataset generated by LLMs with different subset sizes of UMD dataset. Starting with 10% of the UMD training subset, this subset is combined with the selected synthetic dataset. The augmented dataset, which is a mix of synthetic and real instances, is used to fine-tune the pretrained models. We continue this process by increasing the number of real data instances, such as 20% and 30%, until achieving comparable results to those obtained from the model trained on the full UMD dataset.

IV. RESULTS

In this section, first, we present the characteristics of each synthetic dataset. Second, we report a comprehensive comparison of the models fine-tuned with them. Third, we report the data augmentation results. For evaluation, we report two widely-used metrics in this task, accuracy and

⁹ALBERT.

¹⁰DistilBERT.

¹¹Trainer.

¹²Huggingface Transformers.

TABLE 3. Detailed description of generated synthetic datasets.

Dataset #	Model	Learning Method	Topic-Oriented	# of Class	# of Instances
1	Chat GPT	Zero-Shot	Yes	2	549
2	Chat GPT	Zero-Shot	No	2	646
3	Chat GPT	Few-Shot	Yes	2	545
4	Chat GPT	Zero-Shot	Yes	4	492
5	Chat GPT	Few-Shot	Yes	4	594
6	Flan-T5	Zero-Shot	Yes	2	561
7	Flan-T5	Zero-Shot	No	2	502
8	Llama 2	Zero-Shot	Yes	2	395
9	Llama 2	Zero-Shot	No	2	613
10	Mix Dataset	Zero-Shot	Yes	2	1352
11	Synthetic Testing Set	Zero-Shot	N/A	2	318

TABLE 4. Performance evaluation of ALBERT and DistilBERT models on multi-class datasets generated by ChatGPT.

Models	Metrics	Non-Synthetic	ChatGPT	ChatGPT
		UMD Dataset	Zero-Shot	Few-Shot
ALBERT	Accuracy	0.865	0.41	0.36
	F1-Score	0.87	0.43	0.27
DistilBERT	Accuracy	0.77	0.06	0.06
	F1-Score	0.75	0.1	0.12

F-score, to provide a complete and informative evaluation of the performance of the classification models [66].

A. SYNTHETIC DATA GENERATION

A total of nine datasets are generated. An extensive description of these datasets, as well as a mixed set and a test subset, is presented in Table 3. As shown in Table 3, we created binary datasets and four-class datasets, each with the option of including or not including the topics. The binary datasets contain two classes, which allows us to evaluate the model's ability to distinguish between suicidal ideation and non-suicidal instances. On the other hand, the four-class datasets involve multiple categories, enabling us to explore more nuanced predictions of suicidal ideation levels, including "No Risk", "Low Risk", "Moderate Risk" and "High Risk" classes. Moreover, the option to include or not include the topics in these datasets allows us to investigate the impact of information provided by topics on the model's performance.

By comparing the results from datasets with and without topics, we can gain insights into how incorporating topic-related data enhances or influences the model's effectiveness in suicidal ideation detection. Furthermore, as explained in Section III-C1, we created a synthetic testing dataset comprising 10% of each dataset which is annotated by human experts.

B. FINE-TUNED CLASSIFIERS

Two models, ALBERT and DistilBERT are fine-tuned with the generated synthetic datasets. Table 4 presents the results of the performance evaluation of models fine-tuned with multi-class synthetic datasets generated by ChatGPT in

Zero-Shot and Few-Shot settings, tested on the multi-class UMD test set. Considering the poor performance of the multi-class synthetic dataset, we have chosen to disregard the multi-class aspect and proceed solely with the binary approach.

Moreover, Table 5 provides the performance of the models trained on the binary synthetic datasets generated by ChatGPT, Flan-T5 and Llama models and tested on the binary UMD testing subset. We compare the results for the synthetic datasets with those of the UMD training set. Table 5 shows that incorporating topic in generating the datasets significantly improves the performance of the models. For instance, for Llama 2, the topic-oriented dataset increased the F1-score and accuracy of the ALBERT model by 10% and 14% points, respectively. We also created a mixed dataset, including all topic-oriented datasets, to further evaluate the effects of topics on the performance of the models. With both ALBERT and DistilBERT, an F1-score of 0.82 is achieved by the mixed dataset, which is significantly higher than the DistilBERT model trained on the UMD dataset and comparable with the performance of the ALBERT model fine-tuned on the UMD dataset with an F1-score of 0.87.

Table 6 presents the results of models included in Table 5 but tested on synthetic testing datasets. Similar to the results of Table 5, all of the topic-oriented datasets show significant improvement compared to the datasets without any topics. ChatGPT-generated training data, with an F1-score of 0.82, exhibits the best performance, while the performances of Flan-T5 and Llama2-generated datasets are acceptable. Moreover, the mixed dataset shows a 0.81 F1-score, which is an 11% improvement compared to the model trained with the UMD dataset.

C. DATA AUGMENTATION

Based on the results presented in Table 5 and Table 6, the datasets generated by ChatGPT in the Zero-Shot setting show the best results compared to the other datasets. As explained in section III-C2, the augmented dataset now contains a mix of synthetic and real data instances. The augmented dataset is used to fine-tune the pretrained models and then evaluated

TABLE 5. Performance evaluation of the ALBERT and DistilBERT models fine-tuned with binary datasets and tested on UMD testing subset.

Models	Metrics	Non-Synthetic	ChatGPT			Flan-T5		Llama 2		Mix Dataset
		UMD Dataset	With Topic Few-Shot	Without Topic Zero-Shot	With Topic Zero-Shot	Without Topic	With Topic	Without Topic	With Topic	With Topic
ALBERT	Accuracy	0.87	0.67	0.70	0.71	0.48	0.62	0.33	0.75	0.77
	F1-Score	0.87	0.66	0.79	0.79	0.54	0.64	0.49	0.78	0.82
DistilBERT	Accuracy	0.77	0.61	0.63	0.64	0.59	0.77	0.32	0.75	0.76
	F1-Score	0.75	0.59	0.69	0.71	0.61	0.84	0.15	0.77	0.82

TABLE 6. Performance evaluation of the ALBERT and DistilBERT models fine-tuned with binary datasets and tested on synthetic testing subset.

Models	Metrics	Non-Synthetic	ChatGPT			Flan-T5		Llama 2		Mix Dataset
		UMD Dataset	With Topic Few-Shot	Without Topic Zero-Shot	With Topic Zero-Shot	Without Topic	With Topic	Without Topic	With Topic	With Topic
ALBERT	Accuracy	0.67	0.71	0.81	0.81	0.34	0.63	0.48	0.70	0.83
	F1-Score	0.70	0.69	0.78	0.82	0.41	0.69	0.24	0.73	0.81
DistilBERT	Accuracy	0.40	0.65	0.83	0.85	0.63	0.86	0.49	0.63	0.78
	F1-Score	0.61	0.61	0.81	0.81	0.69	0.84	0.12	0.69	0.73

TABLE 7. Performance evaluation of the ALBERT model fine-tuned with the augmented dataset (synthetic data + a subset of the UMD train set) and tested on UMD and synthetic testing subsets.

Test Set	Metric	UMD Dataset	10% (Avg. of 3 Folds)*	20% (Avg. of 3 Folds)*	30% (Avg. of 3 Folds)*
UMD Testing Set	Accuracy	0.87	0.75	0.81	0.83
	F1-Score	0.87	0.79	0.84	0.88
Synthetic Testing Set	Accuracy	0.67	0.87	0.87	0.90
	F1-Score	0.70	0.83	0.86	0.88

* Standard Deviation < 2%

on two separate testing sets. In each iteration, three folds, each comprising 10% of non-overlapping random samples from the UMD dataset, are added to the synthetic data. Subsequently, the average¹³ of the accuracy and F1-score are calculated and reported in Table 7. If the model’s performance with the augmented dataset is less than the model trained with the UMD dataset, additional real-world data is gradually incorporated. For instance, the percentage of real data can be increased to 20% in the next iteration, and the training and evaluation process is repeated.

Throughout the iterations, the model’s performance is closely monitored and compared to the baseline model trained solely on the UMD dataset. The aim is to identify the point at which the augmented dataset starts producing results comparable to or even surpassing those of the baseline model. The process continues until an optimal percentage of real data is found, where the model achieves similar results as the baseline. This ratio indicates the ideal balance between synthetic and real data for achieving high model performance and generalization. Table 7 shows the results of each augmentation process until we achieved the F1-score of 0.87 on the UMD testing subset at 30% augmentation rate and F1-score of 0.85 on the synthetic testing subset at 10% augmentation rate.

V. DISCUSSIONS

This study focuses on the generation of synthetic datasets using generative models and subsequently assessing the

¹³we also calculated the standard deviation of the metrics which were always <0.02.

performance of models fine-tuned with these datasets. Our synthetic data generation framework addresses two limitations of real-world data collection and annotation. First, we address the data scarcity and annotation cost by generating micropost-like suicidal/non-suicidal text. Second, we address the lack of diversity in real-world data by forcing the generative models to create a balanced number of examples related to each of the psychological and social factors impacting suicidality. Integrating insights from psychology into the NLP pipeline in this context can illuminate previously unexplored facets of suicide and mental health detection in social media.

We created several datasets, including binary and multi-class, in Zero-Shot and Few-Shot settings, topic-oriented and non-topic-oriented, with three different generative LLMs. Early in our experiments (Table 4), we observed that ChatGPT is not able to produce high-quality multi-class datasets in either the Zero-Shot or the Few-Shot settings. Generating multi-class datasets using LLMs such as ChatGPT is more complex and challenging task due to the inherent complexities involved in distinguishing between multiple and fine-grained, classes. Even with the availability of a high-quality dataset, one should anticipate lower accuracies in multi-class scenarios. This is largely attributed to the ambiguous boundaries that exist between these classes, creating a complex landscape that proves difficult for any classifier to navigate successfully. Moreover, the creation of such datasets necessitates not only a detailed prompt but also specific instructions that outline the multi-class scenarios. This process demands a nuanced understanding and a level of specificity that often poses a considerable challenge to ChatGPT. Longer [67].

As a result, we opted to exclusively create binary datasets and focus our investigation on how topics impact the overall generalizability of the fine-tuned models.

Our results show the critical role of incorporating domain knowledge in synthetic data generation. We extracted the relevant social topics from the Psychology literature and used that to create more focused prompts for data generation. Table 8 displays a selection of binary samples generated by

TABLE 8. Generated synthetic samples using extracted social and psychological topics by ChatGPT.

Topics	Labels	Sample
Depression	1	Growing up, I always felt like nobody wanted me. No one ever seemed to care about my thoughts, feelings, or emotions. As I got older, these feelings intensified, and I started to feel like I was a burden on everyone around me. I never felt like I belonged anywhere and I constantly struggled with my depression. The constant emptiness and loneliness eventually became too much to bear, and I began to contemplate ending my life.
	0	I've been struggling with mental health problems for as long as I can remember. Depression have been my constant companions. But I've never let it stop me from living my life. I've learned to manage my symptoms with the help of therapy and medication. Some days are harder than others, but I keep pushing forward and focusing on the good in my life. I've come to accept that mental illness is a part of me, but it doesn't define me.
Bullying	1	I felt like I was constantly being attacked by my classmates. They called me names, pushed me around, and even vandalized my belongings. The worst part was feeling so alone and misunderstood. I didn't want to burden my parents with my problems so I kept it all inside. The idea of ending it all seemed like the only way to escape the pain and torment.
	0	When I was in middle school, I was constantly bullied for being overweight. It was tough going to school every day and constantly being made fun of by my peers. I felt so alone and helpless. But I didn't let the bullies control my life. I joined a support group and found friends who lifted me up. I also started exercising and eating healthier, not to please anyone else, but to feel better about myself. It wasn't an easy journey, but it was worth it. I'm happy to say that today I am confident and proud of who I am.
Unemployment	1	After losing my job, I stayed unemployed for months. My savings ran out and with no source of income, my bills piled up. The constant fear of not being able to provide for myself and my family drove me to the brink of despair. I feel worthless and like a burden on everyone around me. The future seems bleak and hopeless, and I wonder if it's worth it to keep going.
	0	After graduating from college, I struggled to find a job in my field for a few months. It was frustrating and disheartening, but I kept applying and networking. Eventually, I landed a job in a related field that I enjoy. It wasn't my dream job, but it paid the bills and gave me experience. I'm still looking for my dream job, but I'm grateful for what I have and optimistic about my future prospects.
Relationship problems	1	I thought I had found the one but it seems like I was wrong because he left me for someone else. I don't know how to deal with this pain. I can't sleep, I can't eat, and I just want to disappear. Maybe everything would be easier if I just ended it all.
	0	My relationship with my partner hasn't been going well lately. We have been arguing over small things, and it's affecting our mental health. We decided to go for couples therapy, and it's been a turning point for us. We learned to communicate better and understand each other's perspective. Now we are in a better place and happier than ever before.

ChatGPT within the synthetic dataset using social topics. The table provides an illustration of specific examples generated by this GLLM.

Table 5 presents the results of fine-tuned models on synthetic datasets and tested on the UMD dataset. Comparison between topic-oriented datasets and no topic-oriented datasets shows the significant effects of including the topics on the performance of the generated datasets. Informing the data generation with topics in Flan-T5 and Llama2 increased the F1-Score of the ALBERT model by 10% and 29% points, respectively. Fine-tuning models on topic-oriented synthetic datasets allows them to gain diverse domain-specific knowledge and patterns. Moreover, non-topic-oriented synthetic datasets might lack specificity, leading to noise and irrelevant content. In contrast, topic-oriented datasets are curated to focus on a specific domain, reducing the chances of generating irrelevant or out-of-context text.

We showed that the BERT family fine-tuned with real-world data can achieve an F1 score ranging from 0.75 to 0.87, depending on the complexity of their structure. Specifically, DistilBERT, a less efficient model from the BERT Family, achieves an F1-score of 0.75, while ALBERT, a more optimized model designed for speed and accuracy, attains an F1-score of 0.87. In contrast, both DistilBERT and ALBERT achieve a consistent F1-score of 0.82 when trained on purely synthetic data and tested on real-world data.

With this, we demonstrate that the diversity of synthetic data compensates for model complexity irrespective of its architecture. This not only underscores the considerable potential of synthetic data but also suggests that it can mitigate the limitations of real-world data in capturing diverse topics. Most notably, our results emphasize an optimal strategy that involves augmenting synthetic data with real data. This innovative method achieves performances comparable to the ALBERT model, even when relying on merely 30% of the manually annotated dataset. This solidifies our proposed method as a cost-effective alternative, addressing the challenges of data scarcity and diversity more effectively than the current benchmarks. However, Synthetic datasets often exhibit a distributional shift from real-world data. This shift arises due to the inherent differences in the data generation processes between synthetic and real domains. As a result, models trained solely on synthetic datasets may not be applicable in real-world situations, leading to a lack of robustness and adaptability. Therefore, exploring hybrid approaches that combine synthetic and real-world data for training can offer a more comprehensive solution. Leveraging both sources allows models to learn from the strengths of synthetic data while adapting to the intricacies of real-world environments.

As presented in Table 5 and Table 7, our study's central objective was to investigate the potential of synthetic and

augmented data in training models to perform effectively on real-world data. Given this setup, the chance of overfitting is inherently reduced since the training (synthetic) and testing (real-world) datasets are obtained from distinct distributions. Moreover, to better understand the performance, robustness, and limitations of the fine-tuned classifiers, we curated an additional test set by manual annotation of a subset of the synthetic data. Additional tests ensure that the models do not overfit a particular dataset and can handle a variety of data distributions and scenarios. Table 6 presents the performance results of the fine-tuned models evaluated on the human-annotated synthetic dataset. Notably, the topic-oriented ChatGPT dataset stands out with an F1-score of 0.82, demonstrating its superior performance compared to the other datasets. Specifically, the model trained with the UMD dataset falls short in handling the synthetic test set, presumably because of its less diverse topics.

VI. CONCLUSION AND FUTURE WORKS

The accurate identification of suicidal ideation from textual data holds paramount importance for early intervention and prevention efforts. Natural Language Processing (NLP) techniques have shown promise in this domain, but the scarcity and sensitivity of real suicide-related data pose significant challenges. Gathering and annotating real suicide-related data is a resource-intensive and ethically sensitive process. Synthetic data generation methods, such as text generation models and data augmentation techniques, offer a more cost-effective way to supplement real data. Also, our synthetic datasets offer a potential solution by providing additional social and psychological context in training instances for models to address the limitations of the existing real data.

Our data augmentation results show that incorporating synthetic data into the training pipeline helps diversify the dataset and enhance model generalization. Real data is often limited in size, leading to over-fitting and reduced model performance. However, by carefully blending synthetic data with real data, we can bolster the model's performance while maintaining a balance between practicality and sensitivity.

Moving forward, exploring the diversity of Language Models (LLMs) stands as an intriguing avenue for future research. Investigating and quantifying the extent of diversity within LLMs across various domains, languages, and training methodologies could offer valuable insights. Future works could delve into developing robust metrics or methodologies specifically tailored to assess and measure diversity within these models.

This paper has effectively highlighted the advantages of using synthetic data generation techniques in detecting suicidal ideation. However, the field still presents numerous opportunities for further research and refinement. Future initiatives could focus on the adaptation of models to various linguistic and cultural environments, acknowledging the diverse ways people express suicidal thoughts across different languages and cultures. Furthermore, a holistic approach

that integrates multiple data modalities, such as images, audio, or behavioral data, alongside textual information could enhance the detection process. It's also crucial to set up a framework that allows for the continuous evaluation and optimization of models, given the ever-changing nature of online communication patterns and user behaviors.

VII. ETHICAL CONSIDERATIONS

For this research, we obtained ethics approval from the research ethics board at the University of Ottawa. Moreover, the UMD dataset was used with authorization from its creators, and we adhered to the terms of use and ethical standards¹⁴ provided by them.

The use of LLMs for suicide-related synthetic datasets raises several ethical considerations. Firstly, synthetic datasets should be generated in a way that avoids perpetuating or amplifying biases present in the original data. It is important to carefully examine the underlying data and the algorithms used in generating synthetic datasets to ensure fairness and mitigate potential biases.

Secondly, the process of generating synthetic datasets should be transparent and well-documented. It is essential to provide clear information about the methods used, assumptions made, and limitations of the synthetic data. This enables others to assess and evaluate the validity and appropriateness of using synthetic datasets.

Thirdly, to use synthetic datasets in sensitive applications or decision-making processes, accountability and liability should be considered. Care should be taken to understand the potential impact and consequences of decisions or actions based on synthetic data and establish mechanisms for addressing any negative outcomes or biases that may arise.

REFERENCES

- [1] D. De Berardis, G. Martinotti, and M. Di Giannantonio, "Understanding the complex phenomenon of suicide: From research to clinical practice," *Frontiers Psychiatry*, vol. 9, p. 61, 2018.
- [2] E. R. Kumar and N. Venkatram, "Predicting and analyzing suicidal risk behavior using rule-based approach in Twitter data," *Soft Comput.*, pp. 1–9, 2023.
- [3] A. Raza, F. Rustam, H. U. R. Siddiqui, I. D. L. T. Diez, B. Garcia-Zapirain, E. Lee, and I. Ashraf, "Predicting genetic disorder and types of disorder using chain classifier approach," *Genes*, vol. 14, no. 1, p. 71, Dec. 2022.
- [4] A. Abdulsalam and A. Alhothali, "Suicidal ideation detection on social media: A review of machine learning methods," 2022, *arXiv:2201.10515*.
- [5] Z. Li, J. Zhou, Z. An, W. Cheng, and B. Hu, "Deep hierarchical ensemble model for suicide detection on imbalanced social media data," *Entropy*, vol. 24, no. 4, p. 442, Mar. 2022.
- [6] D. Kodati and R. Tene, "Identifying suicidal emotions on social media through transformer-based deep learning," *Appl. Intell.*, vol. 53, no. 10, pp. 11885–11917, 2023.
- [7] M. M. S. Fareed, A. Raza, N. Zhao, A. Tariq, F. Younas, G. Ahmed, S. Ullah, S. F. Jillani, I. Abbas, and M. Aslam, "Predicting divorce prospect using ensemble learning: Support vector machine, linear model, and neural network," *Comput. Intell. Neurosci.*, vol. 2022, pp. 1–15, Jul. 2022.
- [8] Q. Wei, A. Franklin, T. Cohen, and H. Xu, "Clinical text annotation—What factors are associated with the cost of time?" in *Proc. AMIA Annu. Symp.* Bethesda, MD, USA: American Medical Informatics Association, Dec. 2018, pp. 1552–1560.

¹⁴The University of Maryland Reddit Suicidality Dataset.

- [9] R. Babbar and B. Schölkopf, "Data scarcity, robustness and extreme multi-label classification," *Mach. Learn.*, vol. 108, nos. 8–9, pp. 1329–1351, Sep. 2019.
- [10] S. I. Nikolenko, *Synthetic Data for Deep Learning*, vol. 174. Cham, Switzerland: Springer, 2021.
- [11] Y. Lu, M. Shen, H. Wang, X. Wang, C. van Rechem, and W. Wei, "Machine learning for synthetic data generation: A review," 2023, *arXiv:2302.04062*.
- [12] H. Chau, S. Balaneshin, K. Liu, and O. Linda, "Understanding the tradeoff between cost and quality of expert annotations for keyphrase extraction," in *Proc. 14th Linguistic Annotation Workshop*, 2020, pp. 74–86.
- [13] D. Hovy and D. Yang, "The importance of modeling social factors of language: Theory and practice," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Human Lang. Technol.*, 2021, pp. 588–602.
- [14] Z. Hu, Y. Xu, W. Yu, S. Wang, Z. Yang, C. Zhu, K.-W. Chang, and Y. Sun, "Empowering language models with knowledge graph reasoning for open-domain question answering," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, Abu Dhabi, United Arab Emirates, 2022, pp. 9562–9581.
- [15] J. C. Franklin, J. D. Ribeiro, K. R. Fox, K. H. Bentley, E. M. Kleiman, X. Huang, K. M. Musacchio, A. C. Jaroszewski, B. P. Chang, and M. K. Nock, "Risk factors for suicidal thoughts and behaviors: A meta-analysis of 50 years of research," *Psychol. Bull.*, vol. 143, no. 2, pp. 187–232, 2017.
- [16] K. H. Bentley, J. C. Franklin, J. D. Ribeiro, E. M. Kleiman, K. R. Fox, and M. K. Nock, "Anxiety and its disorders as risk factors for suicidal thoughts and behaviors: A meta-analytic review," *Clin. Psychol. Rev.*, vol. 43, pp. 30–46, Feb. 2016.
- [17] L. Orsolini, R. Latini, M. Pompili, G. Serafini, U. Volpe, F. Vellante, M. Fornaro, A. Valchera, C. Tomasetti, S. Fraticelli, M. Alessandrini, R. La Rovere, S. Trotta, G. Martinotti, M. Di Giannantonio, and D. De Berardis, "Understanding the complex of suicide in depression: From research to clinics," *Psychiatry Invest.*, vol. 17, no. 3, pp. 207–221, Mar. 2020.
- [18] N. H. Kalin, "Insights into suicide and depression," *Amer. J. Psychiatry*, vol. 177, no. 10, pp. 877–880, Oct. 2020.
- [19] L. da Silva Costa, Á. P. Alencar, P. N. Neto, M. do Socorro Vieira dos Santos, C. G. L. da Silva, S. de França Lacerda Pinheiro, R. T. Silveira, B. A. V. Bianco, R. F. F. P. Júnior, M. A. P. de Lima, A. O. A. Reis, and M. L. R. Neto, "Risk factors for suicide in bipolar disorder: A systematic review," *J. Affect. Disorders*, vol. 170, pp. 237–254, Jan. 2015.
- [20] J. Paris, "Suicidality in borderline personality disorder," *Medicina*, vol. 55, no. 6, p. 223, May 2019.
- [21] K. H. Lee, J. S. Jun, Y. J. Kim, S. Roh, S. S. Moon, N. Bukonda, and L. Hines, "Mental health, substance abuse, and suicide among homeless adults," *J. Evidence-Informed Social Work*, vol. 14, no. 4, pp. 229–242, Jul. 2017.
- [22] C. Okolie, M. Dennis, E. S. Thomas, and A. John, "A systematic review of interventions to prevent suicidal behaviors and reduce suicidal ideation in older people," *Int. Psychogeriatrics*, vol. 29, no. 11, pp. 1801–1824, Nov. 2017.
- [23] E. M. Kleiman, B. J. Turner, S. Fedor, E. E. Beale, J. C. Huffman, and M. K. Nock, "Examination of real-time fluctuations in suicidal ideation and its risk factors: Results from two ecological momentary assessment studies," *J. Abnormal Psychol.*, vol. 126, no. 6, pp. 726–738, 2017.
- [24] N. Leigh-Hunt, D. Bagguley, K. Bash, V. Turner, S. Turnbull, N. Valtorta, and W. Caan, "An overview of systematic reviews on the public health consequences of social isolation and loneliness," *Public Health*, vol. 152, pp. 157–171, Nov. 2017.
- [25] J. Holt-Lunstad, T. B. Smith, M. Baker, T. Harris, and D. Stephenson, "Loneliness and social isolation as risk factors for mortality: A meta-analytic review," *Perspect. Psychol. Sci.*, vol. 10, no. 2, pp. 227–237, Mar. 2015.
- [26] A. Allchin, V. Chaplin, and J. Horwitz, "Limiting access to lethal means: Applying the social ecological model for firearm suicide prevention," *Injury Prevention*, vol. 25, no. 1, pp. i44–i48, Sep. 2019.
- [27] K. A. V. Orden, T. K. Witte, K. C. Cukrowicz, S. R. Braithwaite, E. A. Selby, and T. E. Joiner Jr., "The interpersonal theory of suicide," *Psychol. Rev.*, vol. 117, no. 2, pp. 575–600, 2010.
- [28] A. Wenzel and A. T. Beck, "A cognitive model of suicidal behavior: Theory and treatment," *Appl. Preventive Psychol.*, vol. 12, no. 4, pp. 189–201, Oct. 2008.
- [29] R. J. Cramer and N. D. Kapusta, "A social-ecological framework of theory, assessment, and prevention of suicide," *Frontiers Psychol.*, vol. 8, p. 1756, Oct. 2017.
- [30] A. C. Fernandes, R. Dutta, S. Velupillai, J. Sanyal, R. Stewart, and D. Chandran, "Identifying suicide ideation and suicidal attempts in a psychiatric clinical research database using natural language processing," *Sci. Rep.*, vol. 8, no. 1, p. 7426, May 2018.
- [31] C. A. Bejan, M. Ripperger, D. Wilimitis, R. Ahmed, J. Kang, K. Robinson, T. J. Morley, D. M. Ruderfer, and C. G. Walsh, "Improving ascertainment of suicidal ideation and suicide attempt with natural language processing," *Sci. Rep.*, vol. 12, no. 1, p. 15146, Sep. 2022.
- [32] M. J. Vioules, B. Moulahi, J. Aze, and S. Bringay, "Detection of suicide-related posts in Twitter data streams," *IBM J. Res. Develop.*, vol. 62, no. 1, pp. 7:1–7:12, Jan. 2018.
- [33] R. C. Hsiung, "A suicide in an online mental health support group: Reactions of the group members, administrative responses, and recommendations," *CyberPsychol. Behav.*, vol. 10, no. 4, pp. 495–500, Aug. 2007.
- [34] J. Jashinsky, S. H. Burton, C. L. Hanson, J. West, C. Giraud-Carrier, M. D. Barnes, and T. Argyle, "Tracking suicide risk factors through Twitter in the U.S.," *Crisis*, vol. 35, no. 1, pp. 51–59, Jan. 2014.
- [35] G. B. Colombo, P. Burnap, A. Hodorog, and J. Scourfield, "Analysing the connectivity and communication of suicidal users on Twitter," *Comput. Commun.*, vol. 73, pp. 291–300, Jan. 2016.
- [36] H. Ghanadian, I. Nejadgholi, and H. Al Osman, "ChatGPT for suicide risk assessment on social media: Quantitative evaluation of model performance, potentials and limitations," 2023, *arXiv:2306.09390*.
- [37] K. Yang, S. Ji, T. Zhang, Q. Xie, Z. Kuang, and S. Ananiadou, "Towards interpretable mental health analysis with large language models," 2023, *arXiv:2304.03347*.
- [38] P. P. Sinha, R. Mishra, R. Sawhney, D. Mahata, R. R. Shah, and H. Liu, "#suicidal—A multipronged approach to identify and explore suicidal ideation in Twitter," in *Proc. 28th ACM Int. Conf. Inf. Knowl. Manage.*, Nov. 2019, pp. 941–950.
- [39] M. Gaur et al., "Knowledge-aware assessment of severity of suicide risk for early intervention," in *Proc. World Wide Web Conf.*, 2019, pp. 514–525.
- [40] A. Zirikly, P. Resnik, Ö. Uzuner, and K. Hollingshead, "CLPsych 2019 shared task: Predicting the degree of suicide risk in Reddit posts," in *Proc. 6th Workshop Comput. Linguistics Clin. Psychol.*, Jun. 2019, pp. 24–33.
- [41] H.-C. Shing, S. Nair, A. Zirikly, M. Friedenber, H. Daumé III, and P. Resnik, "Expert, crowdsourced, and machine assessment of suicide risk via online postings," in *Proc. 5th Workshop Comput. Linguistics Clin. Psychol., From Keyboard Clinic*, 2018, pp. 25–36.
- [42] X. He, I. Nassar, J. Kiros, G. Haffari, and M. Norouzi, "Generate, annotate, and learn: NLP with synthetic text," *Trans. Assoc. Comput. Linguistics*, vol. 10, pp. 826–842, Aug. 2022.
- [43] A. Rashid, V. Lioutas, and M. Rezagholizadeh, "MATE-KD: Masked adversarial text, a companion to knowledge distillation," 2021, *arXiv:2105.05912*.
- [44] L. Bonifacio, H. Abonizio, M. Fadaee, and R. Nogueira, "InPars: Unsupervised dataset generation for information retrieval," in *Proc. 45th Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, Jul. 2022, pp. 2387–2392.
- [45] J. Boergers, A. Spirito, and D. Donaldson, "Reasons for adolescent suicide attempts: Associations with psychological functioning," *J. Amer. Acad. Child Adolescent Psychiatry*, vol. 37, no. 12, pp. 1287–1293, Dec. 1998.
- [46] E. D. Klonsky, A. M. May, and B. Y. Saffer, "Suicide, suicide attempts, and suicidal ideation," *Annu. Rev. Clin. Psychol.*, vol. 12, pp. 307–330, Jan. 2016.
- [47] R. Vilhjálmsdóttir, E. Sveinbjarnardóttir, and G. Kristjansdóttir, "Factors associated with suicide ideation in adults," *Social Psychiatry Psychiatric Epidemiol.*, vol. 33, no. 3, pp. 97–103, Feb. 1998.
- [48] C. Lázaro-Pérez, P. M. Gómez, J. Á. Martínez-López, and J. Gómez-Galán, "Predictive factors of suicidal ideation in Spanish university students: A health, preventive, social, and cultural approach," *J. Clin. Med.*, vol. 12, no. 3, p. 1207, Feb. 2023.
- [49] J.-I. Lee, M.-B. Lee, S.-C. Liao, C.-M. Chang, S.-C. Sung, H.-C. Chiang, and C.-W. Tai, "Prevalence of suicidal ideation and associated risk factors in the general population," *J. Formosan Med. Assoc.*, vol. 109, no. 2, pp. 138–147, Feb. 2010.
- [50] A. Farabaugh, S. Bitran, M. Nyer, D. J. Holt, P. Pedrelli, I. Shyu, S. D. Hollon, S. Zisook, L. Baer, W. Busse, T. J. Petersen, M. Pender, D. D. Tucker, and M. Fava, "Depression and suicidal ideation in college students," *Psychopathology*, vol. 45, no. 4, pp. 228–234, 2012.

- [51] J. R. Peteet, G. Maytal, and H. Rokni, "Unimaginable loss: Contingent suicidal ideation in family members of oncology patients," *Psychosomatics*, vol. 51, no. 2, pp. 166–170, Mar. 2010.
- [52] K. A. Ratkowska and D. De Leo, "Suicide in immigrants: An overview," *Open J. Med. Psychol.*, vol. 2, no. 3, pp. 124–133, Jul. 2013.
- [53] J. D. Hovey, "Acculturative stress, depression, and suicidal ideation in Mexican immigrants," *Cultural Diversity Ethnic Minority Psychol.*, vol. 6, no. 2, pp. 134–151, 2000.
- [54] B. T. Keum, M. J. Wong, and R. Salim-Eissa, "Gendered racial microaggressions, internalized racism, and suicidal ideation among emerging adult Asian American women," *Int. J. Social Psychiatry*, vol. 69, no. 2, pp. 342–350, Mar. 2023.
- [55] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 5–8.
- [56] T. Wolf et al., "Transformers: State-of-the-art natural language processing," in *Proc. Conf. Empirical Methods Natural Lang. Syst. Demonstrations*, 2020, pp. 38–45.
- [57] R. Luo, L. Sun, Y. Xia, T. Qin, S. Zhang, H. Poon, and T.-Y. Liu, "BioGPT: Generative pre-trained transformer for biomedical text generation and mining," *Briefings Bioinf.*, vol. 23, no. 6, Nov. 2022, Art. no. bbac409.
- [58] X. Xie, N. Zhang, Z. Li, S. Deng, H. Chen, F. Xiong, M. Chen, and H. Chen, "From discrimination to generation: Knowledge graph completion with generative transformer," in *Proc. Companion Web Conf.*, Apr. 2022, pp. 162–165.
- [59] F. Mi, Y. Li, Y. Zeng, J. Zhou, Y. Wang, C. Xu, L. Shang, X. Jiang, S. Zhao, and Q. Liu, "PanGu-bot: Efficient generative dialogue pre-training from pre-trained language model," 2022, *arXiv:2203.17090*.
- [60] H. Won Chung et al., "Scaling instruction-finetuned language models," 2022, *arXiv:2210.11416*.
- [61] H. Touvron et al., "Llama 2: Open foundation and fine-tuned chat models," 2023, *arXiv:2307.09288*.
- [62] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," 2018, *arXiv:1810.04805*.
- [63] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "ALBERT: A lite BERT for self-supervised learning of language representations," 2019, *arXiv:1909.11942*.
- [64] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter," 2019, *arXiv:1910.01108*.
- [65] T. Wolf et al., "HuggingFace's transformers: State-of-the-art natural language processing," 2019, *arXiv:1910.03771*.
- [66] M. Sokolova and G. Lapalme, "A systematic analysis of performance measures for classification tasks," *Inf. Process. Manage.*, vol. 45, no. 4, pp. 427–437, Jul. 2009.
- [67] Y. Bang, S. Cahyawijaya, N. Lee, W. Dai, D. Su, B. Wilie, H. Lovenia, Z. Ji, T. Yu, W. Chung, Q. V. Do, Y. Xu, and P. Fung, "A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity," 2023, *arXiv:2302.04023*.



HAMIDEH GHANADIAN received the M.A.Sc. degree in electrical engineering from the University of Ottawa, in 2018, where she is currently pursuing the Ph.D. degree in electrical engineering. Her work particularly focuses on the application of natural language processing techniques on suicide and mental health detection on social media platforms and exploring the ways in which NLP can be used to better understand human psychology. Her research interests include natural language processing, applied machine learning, social media processing, and the explainability of AI systems.



ISAR NEJADGHOLI received the Ph.D. degree in artificial intelligence from the Amirkabir University of Technology, Iran, and the Ph.D. degree from the University of Ottawa, Canada, in 2016. She is currently a Senior Research Scientist with the National Research Council Canada and an Adjunct Professor with the University of Ottawa. Her work focuses on responsible AI, specifically on evaluating and improving the transparency and fairness of natural language processing systems. Her research interests include machine learning applications, particularly natural language processing, social media data analysis, and medical text processing.



HUSSEIN AL OSMAN (Member, IEEE) received the Ph.D. degree in electrical and computer engineering from the University of Ottawa, in 2014. He is currently an Associate Professor with the School of Electrical Engineering and Computer Science, University of Ottawa. He is a member with the Multimedia Computing and the Distributed and Collaborative Virtual Environments Research Laboratories. He leads the Multimedia Processing and Interaction Group. His research interests include the application of artificial intelligence in affective computing and biomedical engineering. In particular, he is interested in the development of multi-modal affect recognition methods using deep artificial neural networks to estimate facial expressions and speech sentiment. He studies remote physiological signal measurement using video signals and applies this technology to biomedical and human-computer interaction (HCI) applications. He conducts research in HCI, especially the development of serious games intended for physical rehabilitation and education.

• • •