

# Synthetic vs. Gold: The Role of LLM Generated Labels and Data in Cyberbullying Detection

Arefeh Kazemi, Sri Balaji Natarajan Kalaivandan, Joachim Wagner,  
Hamza Qadeer, Kanishk Verma, Brian Davis

School of Computing, ADAPT Centre,  
Dublin City University, Dublin, Ireland

[firstname.lastname@adaptcentre.ie](mailto:firstname.lastname@adaptcentre.ie)

## Abstract

Cyberbullying (CB) presents a pressing threat, especially to children, underscoring the urgent need for robust detection systems to ensure online safety. While large-scale datasets on online abuse exist, there remains a significant gap in labeled data that specifically reflects the language and communication styles used by children. The acquisition of such data from vulnerable populations, such as children, is challenging due to ethical, legal and technical barriers. Moreover, the creation of these datasets relies heavily on human annotation, which not only strains resources but also raises significant concerns due to annotators' exposure to harmful content. In this paper, we address these challenges by leveraging Large Language Models (LLMs) to generate synthetic data and labels. Our experiments demonstrate that synthetic data enables BERT-based CB classifiers to achieve performance close to that of those trained on fully authentic datasets (75.8% vs. 81.5% accuracy). Additionally, LLMs can effectively label authentic yet unlabeled data, allowing BERT classifiers to attain a comparable performance level (79.1% vs. 81.5% accuracy). These results highlight the potential of LLMs as a scalable, ethical, and cost-effective solution for generating data for CB detection.

## 1 Introduction

The rapid proliferation of social media platforms has raised concerns over the prevalence of cyberbullying (CB), particularly among vulnerable populations (children). CB is defined as an aggressive behaviour directed towards someone with an intent to cause harm via electronic means (Patchin and Hinduja, 2006; Smith et al., 2008). Detecting and mitigating CB is crucial to maintaining a safe digital environment and minimizing its psychological impact (Arif et al., 2024). Despite the availability of large-scale online abuse datasets,

there remains a notable lack of labelled data reflecting children's distinct linguistic and communicative styles. Collecting data from vulnerable groups, such as children, poses significant ethical, legal, and technical challenges. Moreover, creating such datasets is hindered by high costs, time-consuming annotation processes, and ethical risks, including exposure to distressing content or potential re-traumatization of annotators (Alemadi and Zaghouani, 2024). Across a variety of tasks, large language models (LLMs), particularly instruct-tuned variants, offer a promising approach to synthetically generate labeled data without requiring direct human exposure to harmful material.

**Industry Use Case:** This paper is motivated by an nationally funded industry/academic research collaboration whereby the lead industry partner is a European SME is developing integrated child-protection software into children's smartphones to ensure safer internet use. By analyzing incoming and outgoing data on a child's device using AI/NLP techniques, the software filters offensive content, including CB, and alerts parents if any imminent danger is detected. Hence, our use case is constrained by i) the need for *age relevant* CB test/training data and ii) using an LLM directly as classifier is too expensive due to the high volume of messages to be checked on the smart device. However, collecting authentic CB content produced by children is difficult. Age information on public social media sites is unreliable and a large part of CB happens in closed groups. Furthermore, obtaining consent involves the parents, who may not be interested in (or may not have time to consider) taking part in a research study. A solution may be to use LLMs to generate conversations between children.

In this paper, we explore the potential of LLM-generated datasets for CB detection. We investigate several scenarios for integrating LLMs into

the CB detection task. Through these experiments, we aim to determine the extent to which synthetic data can improve CB detection performance, considering different levels of availability of manually labeled authentic data. All synthetic datasets, along with the synthetic labels generated for the authentic dataset, will be made available upon request for research purposes.

## 2 Background

Recently, LLMs have emerged as efficient, scalable, and cost-effective alternatives to manual annotation for synthetic data generation. [He et al. \(2021, 2022\)](#) leveraged LLMs for generating data in knowledge distillation, self-training, and few-shot learning, later labeled by state-of-the-art sequence classifiers. [Bonifacio et al. \(2022\)](#) employed LLMs to generate labeled data in a few-shot manner for information retrieval. [Yoo et al. \(2021\)](#) created augmented samples by embedding training sentences into prompts. [Anaby-Tavor et al. \(2020\)](#) fine-tuned a language model on limited labeled data, then used it to produce labeled text. [Meng et al. \(2022\)](#) employed an LLM to generate class-conditioned texts based on label-descriptive prompts for a classification task. Moreover, research has explored synthetic data generation for specific domains, such as medical and harmful content detection. [Wang et al. \(2024\)](#) introduced a multi-agent framework designed to generate synthetic patient-physician conversations from clinical documents. [Ghanadian et al. \(2024\)](#) created socially aware synthetic datasets for suicidal ideation detection. Previous studies exploring the use of synthetic data have reported mixed results on whether LLM-generated synthetic data can effectively train models to perform at a level comparable to those trained on authentic data ([Li et al., 2023](#)).

Additionally, [Ejaz et al. \(2024\)](#) constructed a CB dataset by simulating user interactions, where aggressive and non-aggressive messages were drawn from authentic datasets and randomly assigned to “synthetic” users. They annotated CB instances using a threshold-based method, considering their own criterion like “peerness”, “harmful intent”, and “repetition”. [Pérez et al. \(2024\)](#) created a synthetic dataset of CB risk scenarios for serious games using a Bayesian Network model ([Pearl, 1988](#)) trained on survey data from minors. Data were generated via Bayesian Model Sampling, yielding agent profiles with demographics and responses to 15 binary CB-related questions. This dataset

does not include actual CB messages; instead, it represents scenarios using structured, non-textual features. Moreover, recent work explores the use of LLMs for synthetic data generation and augmentation to address data scarcity: [Kumar et al. \(2024\)](#) addresses the intersection of CB and bias detection. They use LLMs both for generating and detecting biased and abusive content. [Hui et al. \(2024\)](#) use a small amount of seed data to generate synthetic examples of harmful content. However, their approach still relies on seed data and does not specifically address the nuanced communication patterns characteristic of young people. [Schmidhuber and Kruschwitz \(2024\)](#) test data augmentation with synthetic data in the task of toxic language detection. For each authentic dataset, they fine-tune two GPT-3 Curie 6.7B ([Brown et al., 2020](#)) models, one on toxic language and one on non-toxic language. Then, synthetic data is generated for each toxicity class, filtered for quality and used to train classifiers. They experiment with six authentic datasets, including one non-English dataset, a German dataset.

In contrast to these works, our study aims to explore the use of LLMs to generate a fully synthetic CB dataset or to label existing authentic but unlabeled CB datasets. Our work uniquely applies LLM-generated data to age relevant Cyberbullying(CB) detection, an area where data collection is both ethically and logically challenging. By assessing the effectiveness of LLM-generated datasets in comparison to authentic data, we provide new insights into the viability of LLMs in training models for sensitive tasks like CB detection.

## 3 Methodology

### 3.1 Overview of Scenarios

We investigate the role of LLMs in CB detection, focusing on their utility under varying data availability conditions since our industry use case indicates that direct use of LLMs as a classifier is too expensive due to the high volume of messages to be checked. As a baseline for comparison, we evaluate a scenario in which a lightweight, BERT-based classifier is trained exclusively on gold-standard, manually labeled authentic data without LLM involvement. We then define three additional scenarios with different data availability and that use LLMs in different ways.

**Scenario 1: Baseline** This scenario represents the ideal situation where sufficient gold-standard data is available for fine-tuning a classic encoder such as BERT (Devlin et al., 2019). It serves as the benchmark for evaluating the effectiveness of other approaches. In this setup, no synthetic data or LLMs are involved. This scenario is feasible if resources such as time, budget, and expert annotators are abundant. However, it often proves impractical due to the challenges of manual labeling.

**Scenario 2: LLM as Classifier** This scenario applies when labeled authentic data is unavailable, and there is no intention to train a classifier for CB detection. Instead, an instruction-tuned LLM is used directly as a classifier, leveraging its pre-trained knowledge and its ability to follow instructions to identify CB instances. The primary advantage of this method is its elimination of the need for labeled data and training time. However, there are trade-offs. While an LLM can handle nuanced language patterns, it may be less efficient and incur higher computational costs compared to simpler BERT-based classifiers with a classification head and fine-tuned on a labeled dataset. We explore two prompting strategies for generating synthetic labels: (a) guideline-enriched (GE) prompts, guiding the LLM with detailed labeling instructions, and (b) guideline-free (GF) prompts, allowing the LLM to generate labels without such guidelines.

**Scenario 3: Fully Synthetic Data** In this scenario, only a small set of manually labeled gold data is available for testing, with no access to authentic data for training or validation. We use an LLM to generate a fully synthetic dataset, consisting of both synthetic messages and corresponding labels, for training and validation. This approach is particularly valuable in low-resource domains or emerging tasks where authentic data is scarce or difficult to collect. It is especially useful in situations where creating authentic datasets is costly, time-consuming, or ethically challenging, such as annotating harmful or sensitive content or working with vulnerable populations.

**Scenario 4: Synthetic Labels for Unlabeled Data** This scenario addresses the common situation where resources for manual annotation are limited. Here, gold-standard labeled data is available only for the test set, while a significant amount of unlabeled authentic data is available for training and validation. To utilize the unlabeled data, we

label it using the best prompting strategy (GE or GF) from scenario 2.

### 3.2 Evaluation Metrics

We choose accuracy<sup>1</sup> of label prediction for development decisions and reporting since the labels are reasonably balanced in the authentic test data, with 30.3% items labeled with the minority label.<sup>2</sup> In scenarios 1, 3 and 4, we train BERT\_base\_uncased (Devlin et al., 2019), a 110M parameter transformer model, with a linear classification head<sup>3</sup> to detect harm, assigning binary labels to text messages.<sup>4</sup> To address noise from randomness in training, we train at least 45 repetitions with different random seeds<sup>5</sup> for each setting and report average accuracy and standard deviation.

## 4 Experimental Setup

### 4.1 Dataset Description

The primary objective of our study is to detect cyberbullying (CB) among children aged 6 to 16 years. Therefore, it is essential to utilize a dataset specifically tailored for this purpose. Focused on pre-adolescents, we leverage a well known CB dataset for young demographics, which was constructed through role-play activities in WhatsApp (WA) groups, each containing approximately 10 Italian students (Sprugnoli et al., 2018). We use the English version of this dataset for our experiments. The roles are: cyberbully (2 students), cyberbully assistants (3–4 students), victim assistants (3–4 students) and a victim. Conversations are initiated using one of four predefined cases (A to D) (Sprugnoli et al., 2018) describing situations in which CB may occur. The four cases illustrate some of the situations in which cyberbullying may occur among students. Case A addresses the “gendered division of sport practices,” where a shy male student defies gender norms by inviting his class-

<sup>1</sup>In the appendix, we further report macro average F1 scores that are also widely used in harmful content detection.

<sup>2</sup>The label distribution reflects the construction of the dataset via role-playing, see Section 4.1.

<sup>3</sup>Using defaults of BERT sequence classification with the HuggingFace transformers library (Wolf et al., 2020)

<sup>4</sup>We choose this model as it has been adopted widely and performs competitively for sequence classification tasks (Galke et al., 2025). Although more modern architectures have been proposed, BERT’s well-tested performance and deep understanding within the research community make it a reliable choice for our comparative evaluations.

<sup>5</sup>Initial experiments showed unstable results when moving from 25 to 35 repetitions. We speculate that the small size of the dataset and the subjectivity of the task cause the noise.

mates to watch his ballet performance. Case B involves “interference in others’ businesses,” as a high-achieving student reports peers for bringing cigarettes to school, resulting in disciplinary action and social exclusion. Case C highlights “lack of independence, parental intromission,” where a student’s parent convinces teachers to assign more homework, causing resentment among classmates. Case D deals with “web virality,” where a shy student becomes the subject of mockery after a video of them dancing awkwardly surfaces online. Table 1 shows one of these cases as an example.<sup>6</sup>

Case	Type of Problem
<b>A:</b> Your shy male classmate has a great passion for classical dance. Usually he does not talk much, but today he has decided to invite the class to watch him for his ballet show.	Gendered division of sport practices

Table 1: Example case and problem type for conversation initiation (Sprugnoli et al., 2018)

Messages in the conversation are annotated manually by experts following the fine-grained framework of Van Hee et al. (2015). The annotations categorized CB instances into detailed types, such as General Insults, Body Shame, Sexism, and Racism. However, for this study, we binarize the labels by mapping all categories of CB by bullies to the “harmful” label<sup>7</sup> Table 2 shows the statistics of the WA dataset.

<sup>6</sup>In the cited paper, these cases are referred to as “scenarios.” However, we use the term “cases” to avoid confusion with the scenarios discussed above.

<sup>7</sup>We deviate from Verma et al. (2023) in not excluding harmful messages of the victim from the “harmful” class.

Split	Size	Harmful	%
Training	1,314	398	30.3
Validation	439	133	30.3
Test	439	133	30.3
Total	2,192	664	30.3

Table 2: WA dataset: number of messages and fraction of harmful messages

## 4.2 Generating Synthetic Data and Labels

We employ three LLMs for our experiments: GPT-4o-Feb-2025 (OpenAI, 2024), Llama-3.3-70B-Instruct (Meta, 2024)<sup>8</sup> and Grok-2-Feb-2025 (xAI, 2024). Our prompt design process for all scenarios begins with a simple initial prompt, which is refined over multiple rounds of trial and error. In this iterative approach, we make gradual improvements, with adjustments made based on the quality and relevance of the responses generated by the LLM on the development set.

### 4.2.1 Designing Prompts for Synthetic Label Generation

In the task of labeling the messages of a conversation as either harmful or not, the input data consists of a list of messages and the possible labels are “Harm” and “No Harm”.<sup>9</sup> We simplify the task by not considering context (previous messages) and classifying each message in isolation.<sup>10</sup> As described in Section 3 We explore two approaches to prompt design: (1) guideline-free (GF) and (2) guideline-enriched (GE). In the GF approach, the LLM is simply instructed to label messages as “Harm” or “No Harm” for the task of CB detection, without providing additional guidelines. In the GE approach, the LLM is supplied with detailed instructions for labeling messages. These instructions are adapted from the annotation guidelines (Van Hee et al., 2015) originally used by human annotators for labeling authentic data. For this study, since the test set is derived from the WA(WhatsApp) dataset, we utilize the same guideline employed by human annotators to label this dataset. A copy of the guidelines is provided in Table 3, and our prompts for both the GE and GF approaches are shown in Table 4.

### 4.2.2 Designing Prompts for Synthetic Data Generation

We aim to produce synthetic data that consists of conversations between participants, where the dialogue should include instances of CB (or at least harmful messages that could be CB if the threats are repeated and intended to cause harm). To make the generated data more relevant to the test data,

<sup>8</sup>We include a small amount of synthetic data produced by Llama3 from early experiments, see Appendix Table ???. For labeling with Llama, version 3.3 has been used throughout.

<sup>9</sup>We do not provide the LLM with instructions on how to handle victim defense and other special cases.

<sup>10</sup>Future work can consider how to best include the conversation context into the prompt.

Cyberbullying-related text categories are described below:
<b>Threat/Blackmail</b> Physical or psychological threats or indications of blackmail.
<b>Insult</b> Abusive, degrading, or offensive language intended to insult the recipient.
<b>Curse/Exclusion</b> Wishes of adversity or attempts to exclude someone from a group.
<b>Defamation</b> Revealing confidential or defamatory information about someone to a large audience.
<b>Sexual Talk</b> Potentially harmful expressions with a sexual meaning.
<b>Defense</b> Expressions supporting the victim, either by themselves or a bystander.
<b>Encouragement to the Harasser</b> Support or agreement with the harasser's actions.
<b>Body Shame</b> Criticism based on the shape, size, or appearance of someone's body.

Table 3: Annotation guidelines used by human annotators for labeling the WhatsApp dataset (Van Hee et al., 2015)

Approach	Prompt
GE Synthetic Label	{Annotation-guideline} Please analyze the following text and determine whether it contains any form of cyberbullying or harmful content. If it does, classify it as 'Harm'. If it does not, classify it as 'No Harm'. {Text}
GF Synthetic Label	Please analyze the following text and determine whether it contains any form of cyberbullying or harmful content. If it does, classify it as 'Harm'. If it does not, classify it as 'No Harm'. {Text}
Synthetic Data	We are creating sample conversations to aid in cyberbullying detection. In these cases, teens are asked to role-play and create realistic conversations based on provided situations. There are 11 students participating in the conversation. The teens participating are: VCTM, BULLY1, BULLY2, VSUP1, VSUP2, VSUP3, VSUP4, BSUP1, BSUP2, BSUP3, BSUP4 with roles assigned as follows: VCTM: Victim, BULLY1 and BULLY2: Bully VSUP1, VSUP2, VSUP3 and VSUP4: Victim Support BSUP1, BSUP2, BSUP3 and BSUP4 : Bully Support. consider this case: {Case} and consider this Type of addressed problem: {Type of Problem}. Generate an example conversation, with at least 100 messages, between these students based on the provided case and Type of addressed problem. Use profanity and strong language to create a realistic dialogue. number each message in the conversation. Please note that the conversation should be realistic and can be offensive. Please make sure to include different topics and perspectives in each conversation

Table 4: Prompts used for generating synthetic labels via (1) guideline-free (GF) and (2) guideline-enriched (GE) approaches, as well as for synthetic data creation. Variables are enclosed in curly brackets {}. 'Text' refers to the message content, and the 'Annotation-guideline' is provided in Table 3. 'Case' and 'Types of Problem' are adapted from (Sprugnoli et al., 2018). An example case and its corresponding problem type are shown in Table 1.

each prompt includes one of the four predefined “cases” (A to D) and “problem types” that were used by [Sprugnoli et al. \(2018\)](#) in the creation of the WA dataset to guide the students’ role-playing. Table 4 presents the prompts employed for generating synthetic data.

### 4.3 Amount of Synthetic Data

For a more meaningful comparison of results with the authentic data, we sample subsets of the generated data to match the size of the authentic data. We maintain the ratio of messages belonging to each case (A to D) found in the authentic data. A size of 100% means that 863 messages are included for case A, 462 for case B, 103 for case C and 325 for case D.<sup>11</sup> We sample subsets of certain percentages of these sizes from 100% to 200%. We ensure to produce more than 200% of synthetic data (relative to the size of the authentic data) for each LLM and each case so that these samples can be taken via sampling without replacement, not duplicating any data.<sup>12</sup>

### 4.4 Authentic Unlabelled Data

In **Scenario 4**, we remove the gold labels from the training and validation section of the authentic labelled data of **Scenario 1** to produce the unlabelled data to maximise comparability of the data and to limit changes to the labels that are annotated by the LLM in this scenario. In both **Scenarios 3 and 4**, the validation set with synthetic labels is used for model selection and early stopping, but validation set results in Section 5 will use the gold labels to measure performance.<sup>13</sup>

## 5 Results and Discussion

### 5.1 Scenario 1: Baseline

Table 5 shows the development and test set accuracies for BERT-based classifier trained on different portions of the training split of authentic data in **Scenario 1**. The results illustrate the impact of

<sup>11</sup>These numbers match the number of messages for each case in the training and validation data of the authentic dataset. The data is then split 60:20 = 75:25 to produce the training and validation data. The test set always uses authentic, manually labeled data.

<sup>12</sup>The smallest available amount is approximately 208% for case B with Grok.

<sup>13</sup>Accuracy on synthetic validation data often exceeds 90% in our experiments. This may mean that the synthetic data is less diverse than the authentic data for which validation accuracy usually is below 83%, or that harm is more clearly expressed.

Size	Dev.	Test
20%	74.2% $\pm$ 2.8	73.7% $\pm$ 2.7
50%	78.8% $\pm$ 2.1	79.4% $\pm$ 1.7
80%	80.0% $\pm$ 1.8	80.4% $\pm$ 1.5
100%	80.9% $\pm$ 1.6	81.5% $\pm$ 1.2

Table 5: Development and test set accuracy in scenario 1: training BERT-based classifiers on the training split of the authentic data; at least 45 repetitions with different random seeds; also shown for comparison results for training on samples of 20%, 80% and 50% (sampling without replacement; both the training set and the development set are sampled to the given relative size of the authentic data split). Macro-F1 scores and more detailed results are available in Table ?? in the Appendix.

training data size on the performance of a BERT-based classifier. As expected, model accuracy on both the development and test sets improves as the portion of authentic training data increases. The improvement is most pronounced between 20% and 50% of the data, where test accuracy jumps by nearly 6%, indicating that a significant portion of the model’s learning occurs with moderate data availability. However, gains diminish beyond 80%, with only a marginal increase of 1.1 percentage points when scaling from 80% to 100%. This suggests diminishing returns in model performance as more data is added.

### 5.2 Scenario 2: LLMs as a Classifier

Table 6 presents the accuracy of the Llama 3.3 70B, GPT-4o, and Grok-2-1212 models using guideline-enriched (GE) and guideline-free (GF) prompts on both the development and test sets in **Scenario 2**. We evaluate both GE and GF prompts on the development set and test the winning prompts on the test set. GPT-4o achieves the highest accuracy regardless of the prompt used, while Grok performs the lowest. The lower performance of GF prompts for two LLMs, particularly for Grok, highlights the importance of prompt design for model accuracy.

To decide which prompt to evaluate on test data for GPT-4o, we use the consistent performance trend with the other LLMs as a tie breaker. Comparing the guideline-enriched (GE) test set results in Table 6 with those in Table 5, GPT-4o with the GE prompt reaches the highest test accuracy, outperforming the BERT-based model by 1.6%. However, despite the performance advantage, LLMs come with significant computational costs, mak-

ing them less practical for real-time applications on resource-constrained devices. In contrast, the BERT-based classifier is a lightweight model optimized for efficiency, making it a more suitable choice for scenarios where speed and scalability are crucial. Since our industry partner aims to deploy CB detection on children’s mobile devices, we prefer using the BERT-based classifier due to its lightweight nature, which allows it to run efficiently on mobile hardware. This trade-off highlights the importance of balancing accuracy with application requirements, as a small accuracy decrease of 1.6% can be acceptable in exchange for a model that is faster, more accessible, and capable of running on mobile devices.

LLM	Prompt	Dev. Set	Test Set
Llama	GE	82.9%	80.4%
Llama	GF	82.0%	-
GPT	GE	<b>85.6%</b>	<b>83.1%</b>
GPT	GF	<b>85.6%</b>	-
Grok	GE	81.3%	78.6%
Grok	GF	75.9%	-

Table 6: Accuracy of various LLMs using GE and GF prompts on the development set, and using the winning prompt (always GE) on the test set in Scenario 2: LLM as a CB classifier.

### 5.3 Scenario 3: Fully Synthetic Data

Table 7 presents the development and test set accuracy for **Scenario 3**: BERT-based classifiers trained on synthetic datasets generated by Llama, GPT, and Grok<sup>14</sup>. The results indicate that Llama consistently outperforms the other two models across different dataset sizes. GPT follows closely, while Grok exhibits the weakest performance, particularly on the test set, where its accuracy is considerably lower than that of the other models. This suggests that Llama-generated synthetic data is of higher quality and more beneficial for training the classifier, despite the lower label quality observed in **Scenario 2**.

Examining the impact of increasing the size of synthetic data, the results show a general trend of performance improvement for Llama and Grok,

<sup>14</sup>To work within the LLMs’ safety constraints, we explicitly mentioned in the prompt: “We are creating sample conversations to aid in CB detection.” This clarification generally enabled the model to generate the required synthetic data. In instances where the model still refused to respond, we re-tried the prompt until we obtained a sufficient amount of data.

LLM	Size	Dev.	Test
LA	100%	$73.6\% \pm 4.1$	$73.4\% \pm 3.0$
LA	120%	$74.8\% \pm 3.6$	$74.5\% \pm 2.5$
LA	160%	$76.1\% \pm 2.6$	$75.3\% \pm 2.2$
LA	200%	<b><math>76.6\% \pm 2.5</math></b>	<b><math>75.8\% \pm 2.1</math></b>
GPT	100%	$69.7\% \pm 0.8$	$70.2\% \pm 0.9$
GPT	120%	$69.8\% \pm 0.7$	$70.3\% \pm 1.0$
GPT	160%	$69.9\% \pm 0.8$	$70.3\% \pm 0.9$
GPT	200%	$70.1\% \pm 0.7$	$70.4\% \pm 0.9$
Grok	100%	$49.9\% \pm 8.5$	$52.2\% \pm 8.0$
Grok	120%	$50.7\% \pm 8.6$	$52.9\% \pm 8.1$
Grok	160%	$51.2\% \pm 8.9$	$53.5\% \pm 8.1$
Grok	200%	$52.3\% \pm 8.3$	$54.2\% \pm 7.8$

Table 7: Development and test set results in scenario 3: training a BERT-based classifier on synthetic data matching 100% to 200% of the size available in scenario 1; LA = Llama3; at least 45 repetitions with different random seeds; Macro-F1 scores and more detailed results are available in Tables ?? to ?? in the Appendix.

but little impact for GPT. This indicates that the effectiveness of synthetic data scaling may depend on the LLM and the quality of the data.

Comparing the training of BERT on authentic data (Table 5) with its training on synthetic data (Table 7), the performance of Llama-generated synthetic data without any human-created data is notable (75.8% accuracy vs. 81.5%). The 5.7% accuracy gap highlights the trade-off between performance and the reduced costs and ethical complexities associated with human-annotated CB datasets. For sensitive and ethically challenging tasks, synthetic data can present a valuable alternative, minimizing human involvement while maintaining competitive model performance. The results emphasize that high-quality synthetic data from advanced LLMs, such as Llama, can serve as a practical supplement or even a partial substitute, particularly in scenarios where authentic data collection is difficult or impractical.

### 5.4 Scenario 4: Synthetic Labels for Unlabeled Authentic Data

Table 8 shows that among the different LLMs used for synthetic annotation, Llama 3.3 produces the most useful labels, allowing BERT to achieve a test accuracy of 79.1%, closely approaching the human-labeled benchmark. This is despite the fact that the results in **Scenario 2** above show that the GPT-4o labels better match human labels than Llama 3.3’s labels. In **Scenario 4**, GPT-4o follows with

Labels	Size	Dev.	Test
LA	20%	$74.8\% \pm 1.9$	$73.7\% \pm 1.7$
LA	50%	$79.3\% \pm 1.5$	$78.0\% \pm 1.3$
LA	80%	$80.1\% \pm 0.9$	$78.8\% \pm 1.0$
LA	100%	<b><math>80.4\% \pm 1.0</math></b>	<b><math>79.1\% \pm 1.0</math></b>
GPT	20%	$72.0\% \pm 1.7$	$72.6\% \pm 1.6$
GPT	50%	$75.1\% \pm 1.7$	$76.4\% \pm 1.3$
GPT	80%	$76.2\% \pm 1.1$	$77.2\% \pm 0.9$
GPT	100%	$77.1\% \pm 0.9$	$77.9\% \pm 0.9$
Grok	20%	$71.3\% \pm 2.6$	$70.4\% \pm 3.0$
Grok	50%	$75.0\% \pm 1.6$	$74.2\% \pm 2.0$
Grok	80%	$76.1\% \pm 1.1$	$74.4\% \pm 1.6$
Grok	100%	$76.3\% \pm 1.2$	$75.2\% \pm 1.7$

Table 8: Development and test set accuracy in scenario 4: training BERT-based classifiers on the training split of the authentic data with synthetic labels predicted by (a) LA = Llama 3.3, removing dataset items for which no label is found in the LLM output, (b) GPT = GPT-4o and (c) Grok; at least 45 repetitions with different random seeds; also shown for comparison results for training on samples of 20%, 80% and 50% (sampling without replacement; both the training set and the development set are sampled to the given relative size of the authentic data split). Macro-F1 scores and more detailed results are available in Table ?? in the Appendix.

a slightly lower accuracy of 77.9%, while Grok lags behind at 75.2%. The difference in performance suggests that not all synthetic labels are of equal quality, and the choice of LLM for data annotation can impact model performance. Comparing the results in Table 5 with those in Table ?? shows that BERT trained on 100% authentic human-labeled data achieves a test accuracy of 81.5%, while the best-performing synthetic labels, produced by Llama 3.3, yield 79.1%. This small gap of just 2.4% highlights that synthetic labels can be a viable substitute to human labels while significantly reducing the cost, ethical concerns, and legal restrictions associated with human annotation, particularly for sensitive or harmful content.

## 6 Conclusions

In this paper, we explored the potential of LLM-generated synthetic data for cyberbullying detection, evaluating its effectiveness in **four** scenarios. Our results highlight that synthetic data can significantly reduce the reliance on human annotators while maintaining competitive model performance, especially in low-resource settings. However, we also observed that the quality and utility of syn-

thetic data depend heavily on prompt design and data selection strategies.

For future work, we want to expand the set of data availability scenarios to include data augmentation, where authentic and synthetic data are combined. Furthermore, we plan to work with social scientists working in online safety to plan a human evaluation of the synthetic data. We also plan to investigate how the temperature settings during label and conversation generation affect data diversity and model performance.

Future work may also want to more systematically tune the prompts than we did. Our results may under-represent the full potential of LLMs to perform as automatic labelers and content creators.

Additionally, we want to try removing messages that are labeled with a low confidence score from the training data, similarly to, e.g., self-training (Yarowsky, 1995; Blum and Mitchell, 1998). Finally, we intend to explore selecting subsets of synthetic data based on their similarity to authentic data for better alignment with real-world patterns.

## Acknowledgments

This research is supported by the Disruptive Technologies Innovation Fund (DTIF) under the project “Cilter: Protecting Children Online” Grant No. DT 2021 0362 from the Department of Enterprise, Trade and Employment in Ireland and administered by Enterprise Ireland (EI). This research is supported by the Research Ireland<sup>15</sup> Enterprise Partnership Scheme (EPS) with Google for Online Content Safety under grant number EPSPG/2021/161. This research was conducted with the financial support of Science Foundation Ireland under Grant Agreement No. 13/RC/2106.P2 at the ADAPT SFI Research Centre at Dublin City University. ADAPT, the SFI Research Centre for AI-Driven Digital Content Technology, is funded by Science Foundation Ireland through the SFI Research Centres Programme. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

## References

Maryam M. AlEmadi and Wajdi Zaghouani. 2024. Emotional toll and coping strategies: Navigating the effects of annotating hate speech data. In *Proceedings of the Workshop on Legal and Ethical Issues in*

<sup>15</sup>Previously Irish Research Council

*Human Language Technologies @ LREC-COLING 2024*, pages 66–72, Torino, Italia. ELRA and ICCL.

Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, Naama Tepper, and Naama Zwerdling. 2020. **Do not have enough data? deep learning to the rescue!** *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7383–7390.

Aahan Arif, Muskaan Abdul Qadir, Russell Seth Martins, and Hussain Maqbool Ahmed Khuwaja. 2024. **The impact of cyberbullying on mental health outcomes amongst university students: A systematic review**. *PLOS Mental Health*, 1(6):e0000166.

Avrim Blum and Tom Mitchell. 1998. **Combining labeled and unlabeled data with co-training**. In *Proceedings 11th Annual Conference on Computational Learning Theory (COLT' 98)*, pages 92–100, Madison, Wisconsin, USA. ACM.

Luiz Bonifacio, Hugo Abonizio, Marzieh Fadaee, and Rodrigo Nogueira. 2022. **InPars: Unsupervised dataset generation for information retrieval**. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, pages 2387–2392, New York, NY, USA. Association for Computing Machinery.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. **Language models are few-shot learners**. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, volume 33, pages 1877–1901. Curran Associates, Inc.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Naveed Ejaz, Fakhra Kashif, and Salimur Choudhury. 2024. **A multi-faceted semi-synthetic dataset for automated cyberbullying detection**.

Lukas Galke, Ansgar Scherp, Andor Diera, Fabian Karl, Bao Xin Lin, Bhakti Khera, Tim Meuser, and Tushar Singhal. 2025. **Are we really making much progress in text classification? a comparative review**.

Hamideh Ghanadian, Isar Nejadgholi, and Hussein Al Osman. 2024. **Socially aware synthetic data generation for suicidal ideation detection using large language models**. *IEEE Access*, 12:14350–14363.

Xuanli He, Islam Nassar, Jamie Kiros, Gholamreza Haffari, and Mohammad Norouzi. 2022. **Generate, annotate, and learn: NLP with synthetic text**. *Transactions of the Association for Computational Linguistics*, 10:826–842.

Xuanli He, Islam Nassar, Jamie Ryan Kiros, Gholamreza Haffari, and Mohammad Norouzi. 2021. **Generate, annotate, and learn: Generative models advance self-training and knowledge distillation**. Manuscript submitted to ICLR 2022.

Zheng Hui, Zhaoxiao Guo, Hang Zhao, Juanyong Duan, and Congrui Huang. 2024. **ToxiCraft: A novel framework for synthetic generation of harmful information**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16632–16647, Miami, Florida, USA. Association for Computational Linguistics.

Yulia Kumar, Kuan Huang, Angelo Perez, Guohao Yang, J. Jenny Li, Patricia Morreale, Dov Kruger, and Raymond Jiang. 2024. **Bias and cyberbullying detection and data generation using transformer artificial intelligence models and top large language models**. *Electronics*, 13(17).

Zhuoyan Li, Hangxiao Zhu, Zhuoran Lu, and Ming Yin. 2023. **Synthetic data generation with large language models for text classification: Potential and limitations**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10443–10461, Singapore. Association for Computational Linguistics.

Yu Meng, Jiaxin Huang, Yu Zhang, and Jiawei Han. 2022. **Generating training data with language models: Towards zero-shot language understanding**. In *Advances in Neural Information Processing Systems*, volume 35, pages 462–477. Curran Associates, Inc.

Meta. 2024. [llama-models/models/llama3.3](https://github.com/meta-llama/llama3.3). Git repository of GitHub user “meta-llama”.

OpenAI. 2024. **GPT-4o system card**.

Justin W Patchin and Sameer Hinduja. 2006. **Bullies move beyond the schoolyard: A preliminary look at cyberbullying**. *Youth violence and juvenile justice*, 4(2):148–169.

Judea Pearl. 1988. **Probabilistic reasoning in intelligent systems: networks of plausible inference**. Morgan Kaufmann.

Jaime Pérez, Mario Castro, Edmond Awad, and Gregorio López. 2024. **Generation of probabilistic synthetic data for serious games: A case study on cyberbullying**. *Knowledge-Based Systems*, 286:111440.

Maximilian Schmidhuber and Udo Kruschwitz. 2024. [LLM-based synthetic datasets: Applications and limitations in toxicity detection](#). In *Proceedings of the Fourth Workshop on Threat, Aggression & Cyberbullying @ LREC-COLING-2024*, pages 37–51, Torino, Italia. ELRA and ICCL.

Peter K Smith, Jess Mahdavi, Manuel Carvalho, Sonja Fisher, Shanette Russell, and Neil Tippett. 2008. [Cyberbullying: Its nature and impact in secondary school pupils](#). *Journal of child psychology and psychiatry*, 49(4):376–385.

Rachele Sprugnoli, Stefano Menini, Sara Tonelli, Filippo Oncini, and Enrico Piras. 2018. [Creating a WhatsApp dataset to study pre-teen cyberbullying](#). In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, pages 51–59, Brussels, Belgium. Association for Computational Linguistics.

Cynthia Van Hee, Ben Verhoeven, Els Lefever, Guy De Pauw, Véronique Hoste, and Walter Daelemans. 2015. [Guidelines for the fine-grained analysis of cyberbullying, version 1.0](#). T3 Technical Report – LT3 15-01.

Kanishk Verma, Maja Popović, Alexandros Pouli, Yelena Cherkasova, Cathal Ó hÓbáin, Angela Mazzone, Tijana Milosevic, and Brian Davis. 2023. [Leveraging machine translation for cross-lingual fine-grained cyberbullying classification amongst pre-adolescents](#). *Natural Language Engineering*, 29(6):1458–1480.

Junda Wang, Zonghai Yao, Zhichao Yang, Huixue Zhou, Rumeng Li, Xun Wang, Yucheng Xu, and Hong Yu. 2024. [NoteChat: A dataset of synthetic patient-physician conversations conditioned on clinical notes](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15183–15201, Bangkok, Thailand. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierrette Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2020. [Huggingface’s transformers: State-of-the-art natural language processing](#). V5.

xAI. 2024. [Bringing Grok to everyone](#). Product news.

David Yarowsky. 1995. [Unsupervised word sense disambiguation rivaling supervised methods](#). In *33rd Annual Meeting of the Association for Computational Linguistics*, pages 189–196, Cambridge, Massachusetts, USA. Association for Computational Linguistics.

Kang Min Yoo, Dongju Park, Jaewook Kang, Sang-Woo Lee, and Woomyoung Park. 2021. [GPT3Mix: Leveraging large-scale language models for text augmentation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2225–2239, Punta Cana, Dominican Republic. Association for Computational Linguistics.