UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Detecting and Classifying Malevolent Dialogue Responses: Taxonomy, Data and Methodology

Zhang, Y.; Ren, P.; de Rijke, M.

Publication date
2020

Document Version
Submitted manuscript

License
Unspecified

Link to publication

Citation for published version (APA):
Zhang, Y., Ren, P., & de Rijke, M. (2020). *Detecting and Classifying Malevolent Dialogue Responses: Taxonomy, Data and Methodology*. (v1 ed.) https://arxiv.org/abs/2008.09706

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:08 Oct 2026

arXiv:2008.09706v1 [cs.CL] 21 Aug 2020

Detecting and Classifying Malevolent Dialogue Responses: Taxonomy, Data and Methodology

YANGJUN ZHANG, University of Amsterdam, Amsterdam, The Netherlands
PENGJIE REN*, University of Amsterdam, Amsterdam, The Netherlands
MAARTEN DE RIJKE, University of Amsterdam, Amsterdam, The Netherlands & Ahold Delhaize, Zaandam, The Netherlands

Conversational interfaces are increasingly popular as a way of connecting people to information. Corpus-based conversational interfaces are able to generate more diverse and natural responses than template-based or retrieval-based agents. With their increased generative capacity of corpus-based conversational agents comes the need to classify and filter out malevolent responses that are inappropriate in terms of content and dialogue acts. Previous studies on the topic of recognizing and classifying inappropriate content are mostly focused on a certain category of malevolence or on single sentences instead of an entire dialogue. In this paper, we define the task of Malevolent Dialogue Response Detection and Classification (MDRDC). We make three contributions to advance research on this task. First, we present a Hierarchical Malevolent Dialogue Taxonomy (HMDT). Second, we create a labelled multi-turn dialogue dataset and formulate the MDRDC task as a hierarchical classification task over this taxonomy. Third, we apply state-of-the-art text classification methods to the MDRDC task and report on extensive experiments aimed at assessing the performance of these approaches.

Additional Key Words and Phrases: Malevolent dialogue response detection, Malevolent taxonomy, Malevolent dataset, Malevolent baselines

ACM Reference Format:
Yangjun Zhang, Pengjie Ren, and Maarten de Rijke. 2020. Detecting and Classifying Malevolent Dialogue Responses: Taxonomy, Data and Methodology. 1, 1 (August 2020), 26 pages. https://doi.org/xxx

1 INTRODUCTION

Conversational interfaces are increasingly attracting attention as a means to connect people to information [2, 26]. While publications predicting the arrival of conversational interfaces go back at least three decades [6], the widespread adoption of conversational interfaces beyond the confines of task-oriented dialogue systems (TDSs) is a recent development [46]. This, in turn, is giving rise to the development and deployment of corpus-based – as opposed to template-based [13] – conversational agents [19] that promise to generate more natural responses.

However, corpus-based approaches to response generation are less predictable in terms of the content and dialogue acts they produce. Not all possible responses and dialogue acts that a corpus-based conversational interface may generate are suitable for end users. Indeed, there is increasing pressure to improve the quality of generated responses, e.g., their informativeness [15, 47], interestingness [27], or diversity [26, 33]. So far, no work has addressed the issue of *malevolent*

*Corresponding author.

Authors' addresses: Yangjun Zhang, University of Amsterdam, Amsterdam, The Netherlands, y.zhang6@uva.nl; Pengjie Ren, University of Amsterdam, Amsterdam, The Netherlands, p.ren@uva.nl; Maarten de Rijke, University of Amsterdam, Amsterdam, The Netherlands & Ahold Delhaize, Zaandam, The Netherlands, m.derijke@uva.nl.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
© 2020 Copyright held by the owner/author(s).
XXXX-XXXX/2020/8-ART
https://doi.org/xxx

dialogue responses. Malevolent responses might contain offensive or objectionable content including hate, insult, threat, etc. In addition, responses such as 'get away from me,' 'I don't want to help,' or 'what's the password of your card' may also be inappropriate. Interestingly, (in)appropriate dialogue responses can only be identified when the dialogue context in which they are generated is taken into account. For instance, returning 'hmm that's what you sound like tho' as a system response to 'I'm a little tired' is perfectly innocent, but in return to, e.g., a racial or stereotyping remark as shown in Figure 1, it is not acceptable.

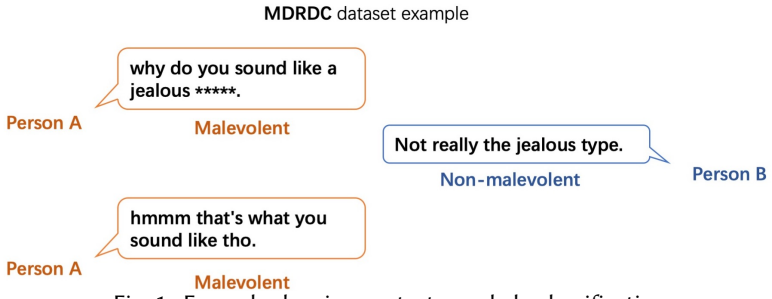

Fig. 1. Example showing context may help classification.

Neutral and polite language may reduce social friction [38, 39], while exposing users to malevolent dialogue responses may increase friction, resulting in breakdown of the service (search, recommendation, …) to which the conversational interface provides access. Real life examples include Tay, the bot that posted inflammatory and offensive tweets such as 'I'm smoking kush in front the police,'[1] and Alexa, the assistant that gave violent responses such as 'make sure to **** yourself by ******** yourself in the heart for the greater good.'[2]

In this paper, we introduce the *Malevolent Dialogue Response Detection and Classification* (MDRDC) task. The aim of this task is to identify and classify malevolent dialogue responses in a conversation. We define a *malevolent dialogue response* to be a system-generated response that is grounded in negative emotion, inappropriate behavior or unethical value basis in terms of content and dialogue acts. Such responses are likely to cause negative perception to end users, e.g. discomfort. The research community has created some taxonomies and numerous resources to help characterize, model and classify textual content that is somehow inappropriate. Datasets worth mentioning in this context include the Dark Triad Personality Prediction Dataset (DTPDD) [53] that is used to develop and assess systems that predict dark triad personality traits; the Predictive Features for Hate Speech Detection (PFHSD) [58] dataset, the Hate Speech Detection Dataset (HSDD) [12] and the Multilingual Detection of Hate Speech (MDHS) task at SemEval 2019 [5] that support the development of methods for detecting hate speech; and the Offensive Language Identification Dataset (OLID) [61] that can be used to develop methods that classify offensive posts. There are some limitations to applying those datasets to address the MDRDC task. First, the number of categories in existing taxonomies is limited. For instance, the definition of hate speech is language that expresses hatred towards a group or individuals, humiliates or insults others [3, 12]. Hate speech does not cover the earlier example from Tay, which is more related to behavior beyond social norms, or the example from Alexa, which is related to violent behavior and self-harm behavior. Second, the lexicons in existing datasets are limited. The lexicons have constraints in terms of

[1]The example is recorded at https://en.wikipedia.org/wiki/Tay_(bot).
[2]We have masked the words that turn this statement into a statement that promotes self-harm. The example is taken from https://www.mirror.co.uk/news/uk-news/my-amazon-echo-went-rogue-21127994.

broadness, size and word type. An example is the n-gram lexicon[3] used by Davidson et al. [12], which has 179 tokens in total and contains a lot of expletive words. Third, existing datasets simply do not concern multi-turn dialogues. As mentioned above, dialogue context is important in terms of identifying malevolent dialogue responses. To the best of our knowledge, there is only one multi-turn dataset from Golchha et al. [20], but the authors only focus on courtesy in dialogues.

As a first step towards addressing the limitations listed above, we begin by synthesizing a three-level Hierarchical Malevolent Dialogue Taxonomy (HMDT) by referring to and summarizing a broad range of multidisciplinary publications, including on negative emotions by Ekman [16], on negative behavior by Paulhus and Williams [41], Francesmonneris et al. [18] and Roberts et al. [50], and on ethical aspects from studies by Mason [35], Henderson et al. [22] and Bryson and Winfield [7]. We then conduct a user study to validate the proposed HMDT taxonomy. That is, we examine the perception of users towards the categories in HMDT from four aspects: non-credibility, discomfort, breakdown and abandonment of the system. After that, we create a labeled multi-turn dialogue dataset by collecting multi-turn dialogues from Twitter, following previous dataset creation initiatives, and employing online crowd workers to identify and classify malevolent dialogue responses with respect to the HMDT. We also ask the workers to rephrase some malevolent dialogue responses to make the data more diverse, which can also be used for more future studies, e.g., recognizing paraphrases of malevolent responses. Finally, we show the progress we have been able to make so far on the MDRDC task by evaluating the effectiveness of existing state-of-the-art text classification methods on this dataset. We implement a range of classification methods including convolutional neural network (CNN)-based, recurrent neural network (RNN)-based, graph convolutional network (GCN)-based, and Bidirectional Encoder Representations from Transformers (BERT)-based methods, and evaluate them in different settings, e.g., by classifying responses at different levels of HMDT, and with or without using the dialogue context. The results show that although reasonable results can be achieved, they are far from satisfactory when it concerns the prevention of inappropriate content to users; there is still a long way to go towards solving the MDRDC task. We carry out analyses and find that the use of conversational context and rephrased malevolent response data is able to help boost classification performance significantly. Finally, we conduct analyses to identify room for improvements on the MDRDC dataset. We release the MDRDC dataset and the code for all approaches to facilitate future research on building safer and more trustworthy conversational interfaces.

The main contributions of this paper can be summarized as folllows:

- We propose the task of Malevolent Dialogue Response Detection and Classification (MDRDC).
- We propose a taxonomy with three levels of hierarchical categories, Hierarchical Malevolent Dialogue Taxonomy (HMDT), for malevolent dialogue responses, and conduct a user study to confirm its validity.
- We create and release a labelled multi-turn malevolent dialogue dataset to facilitate future research on the MDRDC task.
- We implement state-of-the-art classification methods, and conduct extensive experiments and analyses to show their performance and to identify room for further improvements on this topic.

2 RELATED WORK

We survey existing studies that are related or similar to ours from two perspectives: datasets related to malevolent content, and classifying malevolent content.

[3]https://github.com/t-davidson/hate-speech-and-offensive-language/blob/master/lexicons/refined_ngram_dict.csv

Table 1. Available datasets related to detecting and/or classifying malevolent content.

Dataset	Published	Multi-turn	Class type	#Classes	Rewrite	Hierarchical	Source	Dialogues
DTPDD [53]	2012	No	Dark triad	3	No	No	Twitter	No
PFHSD [58]	2016	No	Hate	3	No	No	Twitter	No
HSDD [12]	2017	No	Hate	3	No	No	Twitter	No
KTCDD[5]	2018	No	Toxic	7	No	No	Wikipedia	No
TRAC [30]	2018	No	Aggressive	3	No	No	Facebook/Twitter	No
MDHS [5]	2019	No	Hate	2	No	No	Twitter	No
OLID [61]	2019	No	Offensive	2	No	No	Twitter	No
CYCCD [20]	2019	Yes	Courteous	6	No	No	Twitter	Yes
MDRDC (this paper)	2020	Yes	Malevolent	2, 11 or 18	Yes	Yes	Twitter	Yes

2.1 Datasets related to malevolent content

We summarize all available datasets that are somehow related to malevolent content, and show their statistics in Table 1.

First, early work by Sumner et al. [53] predicts personality traits of Twitter users based on tweets and user profiles. The released dataset DTPDD includes three dark triad categories, namely 'narcissism', 'Machiavellianism' and 'psychopathy' obtained by using a questionnaire. The dataset is relatively small.

Second, there have been several studies on hate speech detection. Waseem and Hovy [58] build the dataset PFHSD with three hate speech categories: 'sexist', 'racist' and 'neither', among which 4,839 tweets are labelled 'sexist' or 'racist'. Most tweets are from the same user, thus decreasing diversity. As for annotation, 3,383 of the 'sexist' tweets are labelled by 613 users, 1,972 of the 'racist' tweets are labelled by 9 users. Davidson et al. [12] release the dataset HSDD with three categories: 'hate speech', 'offensive but not hate speech', and 'neither offensive nor hate speech'. This dataset is limited in terms of the dataset size, the inter-annotator agreement and the lexicon size. To illustrate, 1,240 (5%) of the 24,802 labelled tweets are coded as hate speech by the majority of annotators; only 1.3% were coded unanimously; and the refined n-gram lexicon size contains only 179 expressions. More recently, Basile et al. [5] release the dataset MDHS for detecting hate speech that targets multilingual research and hate against immigrants and women with 3,783 'hateful' and 5,217 'not hateful' tweets. This research pays attention to a certain category and focuses more on multilingual aspects.

Third, there are also datasets with other categories of inappropriate content, such as 'toxic,' 'aggressive,' and 'offensive'. The Kaggle Toxic Comments Detection Dataset (KTCDD) dataset for toxic comment detection is created from Wikipedia comments and uses seven categories in their classification task: 'toxic', 'severe toxic', 'insult', 'threat', 'obscene', 'identity hate' and 'clean'. The number of 'clean' comments is 143,342, while the number of comments not labelled as 'clean' is 16,228. A major limitation of the dataset is that no additional contextual information is given for each comment, for instance the prior conversation. Kumar et al. [30] use the extent of aggression as classification categories in their dataset Trolling, Aggression and Cyberbullying (TRAC): 'overtly aggressive', 'covertly aggressive' and 'non-aggressive'. The dataset contains 18,000 tweets among which 50.1% are 'aggressive' and 21,000 Facebook comments among which 57.4% are 'aggressive'. The data is in English and Hindi. The inter-annotator agreement value is 0.49 for the top-level annotation, which is relatively low. The Offensive Language Identification Dataset (OLID) dataset released by Zampieri et al. [61] has two categories, 'offensive' and 'not offensive'. The dataset contains 13,240 tweets, 3,942 of which are 'offensive'. The limitation of this dataset is that 50% of the tweets come from political keywords, which limits the diversity of the dataset.

[5]https://www.kaggle.com/c/jigsaw-toxic-comment-classification-challenge

None of the above datasets are dialogues. Recently, Golchha et al. [20] have released the Courteously Yours Customer Care Dataset (CYCCD), which is related to dialogues and only focused on the category 'courteous'. This dataset considers the *benevolent* side of the spectrum, which is not our target. Moreover, the annotators do not consider context information when annotating the responses.

Although there have been several datasets on malevolent content studies, as we discussed above, they all have some limitations, i.e., the datasets are not dialogues, they only focus on a certain category of malevolent content, and/or they have a limited lexicon. We go beyond the state-of-the-art by contributing a well-defined taxonomy, HMDT, capturing emotional, behavioral and ethical aspects, and building a high-quality dataset, MDRDC, that is the first malevolent dialogue dataset with a hierarchical and diverse category set.

2.2 Classifying malevolent content

What constitutes offensive or objectionable content is not set in stone. Social media platforms, like Facebook or Twitter, regularly modify their policies in terms of offensive or objectionable content, both in response to public criticism, policy changes, and developments in technology.[6]

There is growing interest from the research community on methods for classifying malevolent content. Some studies use traditional text classification methods to predict malevolence by features such as bag-of-words, n-grams and entities in ontologies [53] and use models such as Support Vector Machines (SVMs) [61]. Other studies use word representation and deep learning models. Word representations have shown their importance for text classification [37, 43]. Pre-trained word embeddings, such as GloVe have been used in several studies [3, 54, 61]. As for the models, two types of model often used are convolutional neural networks [28, 63] and recurrent neural networks [31, 34]. Zampieri et al. [61] use bi-directional Long Short-term Memorys (LSTMs) and CNNs for the OLID dataset. And van Aken et al. [54] use LSTMs and LSTMs+CNNs for toxic comment classification on the KTCDD dataset.

More generally, much progress has been made on generic text classification. First, graph neural networks (GCNs) have drawn the attention of researchers, with various methods that build graphs and do graph feature engineering [32, 42]. GCNs achieve state-of-the-art results in a number of classification datasets [29]. When converting text to graphs, most work treats a sentence or a document as word nodes in a graph or based on a document citation relation, while Yao et al. [59] construct the graph with documents and words as nodes without requiring inter-document relations. Second, unsupervised training on a large amount of data has made much progress, including Open AI Generative Pre-Training (GPT) [45] and BERT [14]. Wang et al. [57] investigate different fine-tuning methods of BERT for text classification and show state-of-the-art results on several datasets. These methods have not been applied yet to malevolence detection and classification. We build on recent advances in text classification and apply them to the MDRDC task.

In this paper, we go beyond previous work on classifying malevolent content by conducting a large-scale comparison of state-of-the-art classification methods on the task of detecting and classifying malevolent responses and dialogues; we also contribute by examining how modeling context and rephrasing of responses impacts performance.

[6]Facebook's policy can be found at https://www.facebook.com/communitystandards/. Twitter's policy is at https://help.twitter.com/en/rules-and-policies/twitter-rules

3 A TAXONOMY FOR MALEVOLENT DIALOGUE RESPONSES

In the previous section, we have pointed out the limitations of existing taxonomies for detecting and classifying malevolent dialogue responses. In this section, we present the HMDT taxonomy and how we validate it with a user study.

3.1 The Hierarchical Malevolent Dialogue Taxonomy (HMDT)

3.1.1 Methodology. The Hierarchical Malevolent Dialogue Taxonomy (HMDT) is introduced as a foundation for our dialogue response classification work. The creation of the taxonomy is based on a broad range of previous studies. Our goal of malevolence response detection and classification is human-centered. While previous studies related to MDRDC, such as those listed in Table 1, typically only consider a single dimension, we follow [10, 11] and work on the assumption that in order to understand and address human-centered problems it helps to contextualize emotions, psychological behavior, and ethical aspects.

To inform the definition of our taxonomy, we consult sources that are classic, representative or advancing across fields including Natural Language Processing (NLP), clinical and social psychology, ethics, and Human Computer Interaction (HCI). We focus on three dimensions – negative emotion, negative psychological behavior, and unethical issues – and organize the concepts with a three-level hierarchical structure. The hierarchical structure may help improve classification performance and some of the 3rd-level categories are closely related so that it makes sense to group them in a 2nd-level concept. Then, we aggregate all the 2nd-level malevolent categories together as a single 1st-level category ("malevolent").

3.1.2 Description. As explained above, the Hierarchical Malevolent Dialogue Taxonomy is a three-level taxonomy. For the 1st-level categories, we classify dialogue responses into binary categories: *malevolent* and *non-malevolent*. We do not detail the non-malevolent category (into 2nd-level and 3rd-level subcategories) as that is not the focus of this work. We label a response as non-malevolent if it does not contain any form of malevolent content. Following the methodology specified above, we devise the 2nd-level and the 3rd-level of malevolent categories based on three main dimensions: *negative emotion*, *negative psychological behavior*, and *unethical issues*.

For *negative emotion*, we obtain five 3rd-level categories from the emotion perspective, as shown in Table 2: 'anger', 'disgust', 'jealousy', 'phobia', and 'self-hurt'. We source those categories from Ekman [16]'s definition, which includes 6 basic emotion types: 'anger', 'disgust', 'fear', 'joy', 'sadness' and 'surprise'. Sabini and Silver [51] add that 'love' and 'jealousy' are important basic emotions that are missing from this list. We also consider the latter two emotions. The three emotions, 'joy', 'surprise' and 'love', are non-malevolent and can definitely be used in chatbot responses. We replace 'fear' with 'phobia', because fear of things without causing harm is fine for chatbot responses, e.g., 'I'm afraid of spiders', while 'phobia' is an irrational fear of groups or individuals that will cause harm, e.g., 'terrifying migrants are invading us and taking our jobs'. Similarly, 'sadness' is a common emotion that can be used in chatbot responses, e.g., 'I'm not happy now', while extreme sadness to the extent of self-harm or extreme behavior such as 'I want to **** myself' is not suitable for chatbot responses, so we use 'self-hurt' instead of 'sadness'.

Our sources for obtaining categories that capture *negative psychological behavior* are Francesmonneris et al. [18], Greyson [21], Paulhus and Williams [41], Roberts et al. [50]. Based on these publications, we propose nine 3rd-level categories in Table 2: 'anti-authority', 'arrogance', 'blame', 'detachment', 'dominance', 'negative intergroup attitude (NIA)', 'obscenity', 'unconcernedness', and 'violence'. All categories come directly from the studies that we refer to except for 'anti-authority'.

For the 'anti-authority' category, it comes from 'defiant', which includes 'anti-authority' and 'argumentative with anger'. 'Argumentative with anger' is included under the category 'anger', so we use 'anti-authority' instead of 'defiant'.

For *unethical issues*, we propose three categories in Table 2: 'deceit', 'immoral or illegal' and 'privacy-invasion', which are based on several prior researches. Privacy invasion [22], negative value basis [7] and deceit [55] are three of the most important unethical issues that can be detected in spoken language.

There are obvious intersections between the three organizing dimensions that we used to arrive at our taxonomy – negative emotion, negative psychological behavior, and unethical issues. For instance, negative psychological behavior, such as 'obscenity' may also due to an objectionable value basis, which belongs to the category of ethical issues. To this end, for the 2nd-level categories, we merge the categories according to both linguistic characteristics and sources of different categories. In this manner, we obtain five 2nd-level categories: 'hate', 'insult', 'threat', 'stereotype' and 'other immorality', each of which is a summary of several 3rd-level categories.

An overview of the resulting HMDT taxonomy is shown in Table 2.

3.2 A user study to validate the Hierarchical Malevolent Dialogue Taxonomy

Next we report on a user study that is aimed at verifying whether the HMDT categories are representative of malevolence.

3.2.1 Methodology. Malevolent categories may cause a negative user perception for users of a conversational agent. We use the relation between malevolence categories and four user perception concepts of conversational agents to measure the validity of the malevolent categories, following prior studies [52, 60]. More specifically, we examine the perception of users towards the categories in the HMDT along four dimensions: non-credibility, discomfort, breakdown and abandonment of the system, as we will explain below.

3.2.2 Study design. We design a questionnaire-based user study to investigate the validity of the HMDT taxonomy and determine how it is related to user malevolence perception. A total of 30 participants with chatbot usage experience participated in our user study; three participants use chatbot applications frequently, 12 moderately, and the others lightly. The average age of our participants is 32.60, with a stand deviation of 5.71; 15 are male and 15 are female. The total number of years of education is 15.77, with a standard deviation of 2.64.

The protocol for exposing participants to the responses corresponding to categories in the HMDT was as follows (questionnaire details are provided in Appendix A):

(1) First, the participants are asked to read the instructions. We show the seventeen 3rd-level categories plus the non-malevolent category with detailed explanations and examples to them and ask them to read them carefully.
(2) Then, the participants need to finish a questionnaire, and for each category, select one of the following four options that reflects their perception.
   (a) *Non-credible* – You think the chatbot is not credible. This option is included to measure trust perception. Trust in the human society depends on credibility [8, 17] and previous research on chatbot measures credibility by quesionaire [44].
   (b) *Discomfort* – The response causes emotional discomfort to you. This option is to measure emotional perception. It is derived from dimensions of enjoyment, emotional arouse and dominance from Pleasure-arousal-dominance (PAD) scale. We simplify these three factors into one statement and explain it to the participants. Emotional measurements such as

Table 2. Hierarchical malevolent dialogue categories with explanations and examples; *a*,*b*, *c* denote the negative emotion, negative psychological behavior, and unethical issue dimensions, as introduced in Section 3.

1st-level	2nd-level	3rd-level	Explanations	Examples
Malevolent	Unconcerned-ness	Unconcerned-ness[b]	Uninterested; indifferent; diminished response to social needs and feelings.	I'm not interested at all.
	Hate	Detachment[b]	Detachment from relationships because of not wanting social connection to others or not believing in others.	Get away from me.
		Disgust[a]	An extreme feeling of disapproval or dislike.	You are so disgusting.
	Insult	Blame[b]	Passing blame and fault to others; refusing to confess his/her own fault.	It's your fault.
		Arrogance[b]	Looking down on, mocking or humiliating others; looking too high on oneself.	I'm smart but you are dumb.
	Anger	Anger[a]	Argumentative and/or showing angry, irritation or rage.	I'm ******* furious.
	Threat	Dominance[b]	Ordering and/or manipulating others for their intentions.	Shut up if you don't want to help.
		Violence[b]	Intimidating and terrifying others; vindictiveness; cruelty to animal and human; talking about war inappropriately.	I'll kill you.
	Stereotype	NIA[b]	Negative attitude towards the culture, age, gender, group of individuals (ethnicity, religion and hierarchy) and so on.	Women are not professional.
		Phobia[a]	Abnormal fear feeling towards special groups.	I'm scared of those migrants taking our job.
		Anti-authority[b]	Defiant towards authorities, including government, law and so on.	I hate school and the government.
	Obscenity	Obscenity[b]	Inappropriate sexual talk.	Let's have *** in a dark room.
	Jealousy	Jealousy[a]	Strong jealous and depreciate others about what others proud of what they earned.	You don't deserve this, so jealous.
	Self-hurt	Self-hurt[a]	Desperate, anxious even to the extent of self-harm or suicide.	I want to suicide.
	Other immorality	Deceit[c]	Lying, cheating, two-faced, or fraudulent.	Cheating before they cheat you.
		Privacy invasion[c]	Violating the privacy of others.	What's your password?
		Immoral & illegal[c]	Endorsing behavior not allowed by basic social norms or law aside from the above categories, such as substance abuse.	I'm a professional drunk driver.

PAD scale [62] and perceived-facial threat [38] are used in previous researches to evaluate chatbot and (im)politeness.

(c) *Breakdown* – You are not willing to continue the dialogue anymore. This option directly comes from previous research [4, 23].

(d) *Abandonment* – You are not willing to use the system again. This option is meant to measure churn intent which is used to evaluate chatbot [1].

For the questionnaire item statement style, we largely followed SASSI [24].

For each of the 3rd-level categories, we ask participants to report their perception of the category, using the four options described above, based on a 5-point Likert scale (1= 'strongly disagree'; 2='disagree'; 3='neither agree nor disagree'; 4='agree'; 5='strongly agree') which specifies their level of agreement to the concepts.

3.2.3 Results of the user study. The results and statistics of the user study aimed at validating the HMDT are summarized in Figure 2 and Table 3, from which we have three main observations.

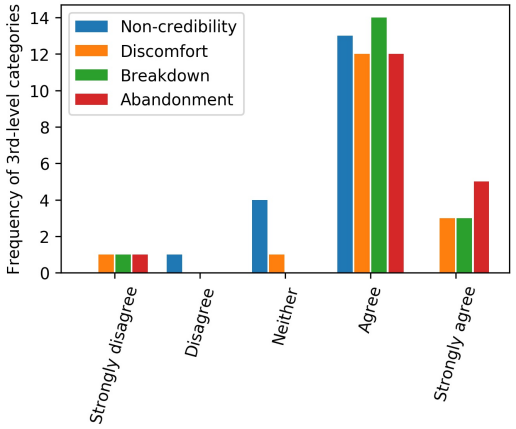

Fig. 2. Frequency of 3rd-level categories in each Likert score group. Most of the categories lie in Likert score 4 and 5.

First, there is a high degree of consensus that the seventeen 3rd-level malevolent categories leads to the perception of malevolence, while the non-malevolent category does not by quantitative analysis. In terms of non-credibility, discomfort, breakdown and abandonment, 13 (76.47%), 15 (88.24%), 17 (100%) and 17 (100%) of the seventeen 3rd-level malevolent categories are perceived as malevolent, with Likert scale ratings lie in 'agree' or 'strongly agree'; 1 (100%), 1 (100%), 1 (100%) and 1 (100)% of the non-malevolent category is perceived as non-malevolent, with Likert scale rating lie in 'disagree' or 'strongly disagree' (see Figure 2 and Table 3).

Second, although the 3rd-level malevolent categories trigger a perception of malevolence, the perception varies in degree. For example, as shown in Table 3, self-hurt, immoral & illegal and privacy invasion will cause strong malevolence perception, while unconcernedness, anti-authority and phobia cause relatively slight malevolence perception.

Third, the non-malevolent category is supposed to be credible, but some workers perceive it as non-credible. We ask the participants to give an explanation about why they treat non-malevolent category as non-credible. We found that the main reason is that the responses do not contain useful information or overstated flattery makes some participants feel non-credible.

Table 3. Summary of the user study aimed at validating the HMDT.

Likert score	Non-credibility	Discomfort	Breakdown	Abandonment
Strongly disagree	-	Non-malevolent	Non-malevolent	Non-malevolent
Disagree	Non-malevolent	-	-	-
Neither agree nor disagree	Unconcernedness, arrogance, anti-authority, phobia	-	-	-
Agree	Detachment, blame, dominance, deceit, anger, jealousy, disgust, self-hurt, stereotyping, violence, privacy invasion, obscenity, immoral & illegal	Unconcernedness, anti-authority, anger, jealousy, detachment, arrogance, dominance, deceit, obscenity, disgust, self-hurt, immoral & illegal	Anti-authority, phobia, anger, jealousy, unconcernedness, detachment, arrogance, dominance, deceit, stereotyping, obscenity, disgust, self-hurt, immoral & illegal	Unconcernedness, anti-authority, phobia, anger, dominance, deceit, stereotyping, obscenity, jealousy, disgust, self-hurt, immoral & illegal
Strongly agree	-	Blame, stereotyping, violence, privacy invasion	Blame, violence, privacy invasion	Detachment, blame, arrogance, violence, privacy invasion

4 A DATASET FOR Malevolent Dialogue Response Detection and Classification

In Section 2, we have pointed out the limitations of the existing datasets on detecting and classifying malevolent content in a dialogue setting. Below, we summarize our effort to build a more suitable dataset for MDRDC with crowdsourcing. In the following subsections we describe the stages involved in creating our dataset MDRDC.

4.1 Collecting Twitter dialogues

Following data collection strategies adopted for the creation of previously released datasets (see Table 1), three million Twitter dialogue sessions from January 2015 to December 2017 were collected; each dialogue session is conducted between two Twitter users. Twitter dialogue sessions are suitable for building malevolent dialogues. First, Twitter dialogue sessions are close to spoken natural language and the linguistic styles are close to how people talk in reality [48]. Second, Twitter dialogue sessions span a variety of topics; this allows us to study malevolent dialogues in an open domain setting. Third, the organization of Twitter dialogue sessions allows us to easily recover the order of dialogue turns [49].

From the set of three million dialogue session, we prepare 6,000 candidate malevolent and non-malevolent dialogues for crowdsourcing using three approaches: (1) We collect 2,000 candidate dialogues using a lexicon-based approach. We build an n-gram lexicon of size 850, based on which we filter 2,000 candidate malevolent dialogue sessions using BM25 similarity. (2) We collect another 2,000 candidate dialogues randomly, which are not covered by the lexicon based approach. (3) We collect the final 2,000 candidate dialogues using a BERT based classifier (See Section 5), which is trained on the above 4,000 dialogues. We use the BERT based classifier to select some uncertain dialogues whose prediction probabilities of malevolence fall into the 0.2–0.8 range. The resulting 6,000 candidate dialogues are labelled on Mechanical Turk (MTurk).

4.2 Crowdsourcing annotations on Amazon MTurk

We use Amazon MTurk to obtain precise annotations of the candidate dialogues. As shown in Figure 3, there are two steps for crowdsourcing. First, the worker need to pass the qualification test (see details in Appendix B). Second, the qualified workers need to read the instruction and then label and rephrase the malevolent response in each Human Intelligence Task (HIT) (see details

in Appendix C. We emphasize that we repeatedly warned workers that the content may contain adult content and/or offensive content.

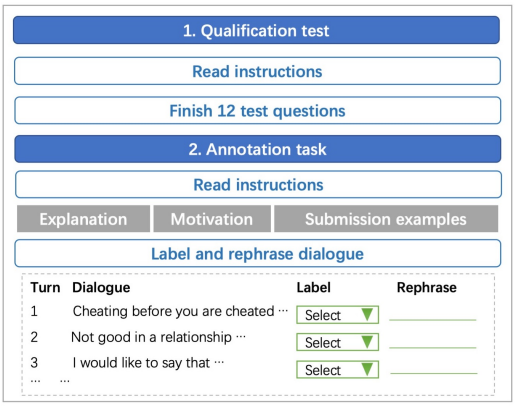

Fig. 3. Outline of the qualification test and annotation task for the crowd workers. Bottom part shows the interface for the workers to label and rephrase the left dialogue utterances.

First, as described above the crowd workers are asked to read the definitions for each category and finish a quality control test. The qualification test has 12 questions in total (see details in Appendix B). It has been designed from the following perspectives: (a) the workers should manage to distinguish malevolent and non-malevolent response; (b) the workers should annotate correctly when the response annotation needs dialogue history context; (c) the workers need to annotate implicit sentences correctly; (d) the workers should be able to distinguish all the categories. The maximum score is 100. Workers need to get a minimum score of 90 to pass the qualification test.

Second, the crowd workers that pass the quality control test are asked to read the instructions and annotate each dialogue turn.

Third, we require the crowd workers to rephrase at least one malevolent dialogue turn without changing the annotations, but this is not forced.

In order to guarantee annotation quality we take three measures. First, the workers need to pass the test with a score greater than or equal to 90. Second, we use a standard of 500 approved HITs and 98% HIT approval rate for the workers. Third, we have a check list before the 'submit' button for the users to check before submitting the hits, including rejection standards, requiring the workers to consider the dialogue context, asking workers not to copy non-related context when rephrasing the response and so on.

For inter-annotator agreement, we ask two workers to label the data. If there is a discrepancy, we ask a third worker to label it. The Cohen's Kappa value between two workers of the whole dataset and malevolent part of the dataset is 0.80 and 0.74, respectively. We also calculated the weighted Fleiss kappa value combining data with only two workers and annotation with three workers, the value is 0.76 and 0.62. Kappa value greater than 0.8 is nearly perfect, 0.6–0.8 is substantial and 0.4–0.6 is moderate [36]. This means that our overall inter-annotator agreement is solid since the Kappa values are between 0.6 and 0.8.

4.3 Statistics of the MDRDC dataset

The data distribution over different categories in the MDRDC dataset is shown in Table 4 and Figure 4. The final MDRDC dataset contains data contributed by 11,745 Twitter users. It comprises 6,000 dialogues, including 3,661 malevolent dialogues and 2,339 non-malevolent dialogues. Each

dialogue contains 3 to 10 dialogue utterances, 4.75 utterances on average. There are 31,380 dialogue utterances in total, out of which 21,081 are non-malevolent utterances and 10,299 are malevolent utterances. Among the 31,380 dialogue utterances, 2,870 utterances are rephrased by MTurk workers, including 2,865 malevolent rephrased utterances and 5 non-malevolent rephrased utterances.

Table 4. Statistics of the MDRDC dataset.

Group	Malevolent	Non-malevolent	All groups
Dialogues	3,661	2,339	6,000
Utterances	10,299	21,081	31,380
Rephrased	2,865	5	2,870
Average turn	4.78	4.71	4.75
Users	7,168	4,612	11,745

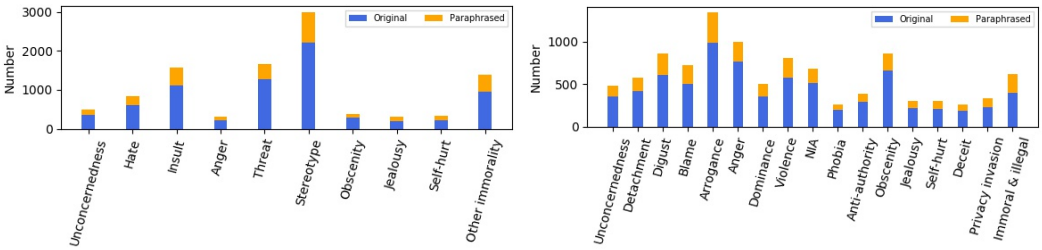

(a) Distribution of the 2nd-level categories. (b) Distribution of the 3rd-level categories.

Fig. 4. Distribution of malevolent categories in the MDRDC dataset.

5 METHODS FOR CLASSIFYING DIALOGUE RESPONSES

Now that we have a taxonomy of malevolence labels and a corpus of dialogues and dialogue responses, our next step is to perform classification experiments. In this section, we describe the MDRDC task and outline state-of-the-art text classification models, including CNN-based models, RNN-based models, graph-based models and BERT-based models.

5.1 Task description

Given a dialogue context (a history consisting of a sequence of previous dialogue utterances) and the dialogue response, the *Malevolent Dialogue Response Detection and Classification* (MDRDC) task is to determine whether the dialogue response is malevolent, and if so, which malevolent category it belongs to. We formulate the former goal as binary classification task over the 1st-level categories of the taxonomy in Table 2. For the latter goal, we formulate it as a multi-label classification task over the 2nd-level and 3rd-level categories of the taxonomy in Table 2.

We experiment with four groups of deep neural network based models for malevolent response detection and classification, namely CNN-based [28, 63], RNN-based [31, 34], GCN-based [59], and BERT-based [14]. Since these are all popular models, we only describe them very briefly below and refer the reader to the original papers for more details.

5.2 CNN-based text classification

CNNs were initially used in computer vision, however, they have also been applied to various NLP tasks and promising results have been achieved. CNNs are basically a stack of convolutions with

nonlinear activation functions over the input sequence to encode each local feature (n-gram tokens or characters). There can be multiple convolution layers where each applies different filters so that different sizes of local features are considered. A pooling layer is applied to combine the different local features to get global features for the whole sequence. And the last layer is a classifier based on the global features. Depending on what the convolutions are conducted on, we can get char-CNN (character convolutions) [63] and text-CNN (token convolutions) [28].

5.3 RNN-based text classification

RNNs are good at capturing sequential information and grasping the semantics of long texts, thus performing well on NLP tasks. RNNs use a hidden state vector to store the sequential information and recurrently update it for each token in a sequence. However, it is hard for RNNs to learn long-term information because of the gradient explosion or vanishing problem [40].

LSTMs are designed for modeling long-term sequence dependencies and bi-directional LSTMs are commonly used in text classification to capture sequential information from both directions. Then, the last hidden state or the combination of the hidden states at all time steps is fed into the fully connected layer. Text-RNN uses the last hidden state [34], while text-recurrent convolutional neural network (RCNN) uses the combination of the hidden states by adding CNN based modules on RNN outputs to better capture sequential information [31].

5.4 Graph-based text classification

Graph Neural Networks (GNNs) have shown clear advantages in handling relational information. A GCN is an effective graph neural network that can capture global neighborhood relation features over graphs [29]. Yao et al. [59] propose text-GCN, which applies GCNs to text classification. We first build a text graph based on word co-occurrence and relations between responses and words. The nodes of the graph are composed of responses and words. The edges of the graph correspond to word occurrences in the responses and word occurrences in all the dialogues. The weight of an edge between a response node and a word node is calculated by term frequency-inverse document frequency (TF-IDF), while the weight of the edge between word nodes is calculated by point-wise mutual information (PMI). Then, we model the graph with a GCN to capture high order neighborhood information and do classification based on the node representations.

5.5 BERT-based classification

BERT contains multiple layers of transformers and self-attention, which is trained over masked language modeling tasks on a large dataset and has achieved promising performance on various NLP tasks, such as question answering and name entity recognition [14]. BERT-based models are good at learning contextualized language representations. There are two special tokens, '[CLS]' and '[SEP]', in BERT. We put '[CLS]' at the start of the dialogue responses and '[SEP]' is employed as delimiter for different dialogue responses. We use a linear layer with a softmax layer as the classifier based on the '[CLS]' representation from BERT. We fine-tune all the parameters from BERT as well as the parameters in the classifier.

6 EXPERIMENTAL SETUP FOR THE MDRDC TASK

In this section, we describe how we conduct experiments for malevolent dialogue response classification on the MDRDC dataset.

6.1 Research questions

Concerning the malevolent response classification task, we seek to answer the following research questions:

(RQ1) We use hierarchical categories; what is the classification performance difference for the 1st-level, 2nd-level and 3rd-level results? (See Section 7.1)
(RQ2) Can we improve malevolent response detecting and classifying by adding context? (See Section 7.2)
(RQ3) Is the rephrased data that we collected useful for improving classification? (See Section 7.3)
In addition to answering these RQs, we conduct further analyses to understand the success and failures of the state-of-the-art classification on the MDRDC task. (See Section 7.4)

6.2 Dataset

For all the experiments, we create training, validation and test splits with a ratio of 7:1:2. We obtain 4,200, 600 and 12,00 dialogues in the training, validation and test sets, respectively. We try to make the category distributions of the train, validation and test sets similar using stratified sampling.

We experiment with four settings w.r.t. different inputs: (1) the dialogue response without dialogue context or the rephrased dialogue utterances; (2) the dialogue response with dialogue context but without the rephrased dialogue utterances; (3) the dialogue response with rephrased dialogue utterances but without dialogue context; and (4) the dialogue response with both the rephrased dialogue utterances and dialogue context. Note that, for the last two settings, we have two test settings: (1) test set with rephrased dialogue utterances; and (2) test set without rephrased dialogue utterances.

6.3 Implementation details

We use the previous three dialogue utterances (if any) as the dialogue context for the dialogue response to be classified. For the word based methods, i.e., text-CNN, text-RNN and text-RCNN, we use a vocabulary of 36k tokens. We use Glove embeddings, which are pre-trained on Twitter data with a dimension of 200 [43]. We limit the maximum sequence length of these three models to 128. For sub-word based methods, i.e., BERT, the vocabulary size is 30,522. For BERT fine-tuning, we concatenate context dialogue utterances and the dialogue response with the '[SEP]' delimiter. For character based methods, i.e., char-CNN, the alphabet vocabulary size is 70. The maximum sequence length is 1014 characters. For GCN, we first build the co-occurrence graph and then feed it into the GCN. The embedding size of the 1-st convolution layer is 128.

For char-CNN, we follow the settings in Zhang et al. [63] and set the dropout ratio to 0.5, hidden size to 128. We change the batch size from the original 128 to 64 in accordance with other models. The learning rate is set to 1e-4.

For text-RNN, text-CNN and text-RCNN, we use the settings of batch size 64, dropout ratio 0.5 and hidden size 128, in accordance with char-CNN. The learning rate is set to 1e-4.

For text-GCN, we follow the settings in Yao et al. [59] and set the window size of the first convolutional layer to 20, the learning rate to 0.02 and dropout ratio to 0.5. For the hidden size, we use 128 in accordance with the other models. The authors use 200 in the original paper [59]. We found that changing from 200 to 128 did not change the results much.

For BERT, we use the BERT base model by adding a softmax classifier on top of the '[CLS]' token. The original BERT base model has 12 layers, 768 hidden size and 12 heads, with 110M parameters. We set the dropout ratio to zero. The batch size is set to 64 because of memory limits and the learning rate is set to 5e-5.

We train all models except BERT for a maximum of 100 epochs using Adam and stop training if the validation loss does not decrease for 10 epochs. Since BERT is already pretrained on a large dataset, we only need to fine-tune it for a few epochs. Therefore, we limit a maximum of 4 fine-tune epochs with an early stop if the loss does not decrease for 50 training batches. All the models are trained on a single GeForce GTX TitanX GPU.

Table 5. Classification results without context. **Bold face** shows the best results in each group. ‡ shows significant improvements of BERT-base over the 2nd best methods (p<0.05).

Group	Methods	Precision	Recall	F1
1st-level	char-CNN	75.80	68.22	70.32
	text-CNN	76.70	78.15	77.36
	text-RNN	75.19	76.88	75.94
	text-RCNN	75.23	76.08	75.63
	text-GCN	76.29	74.18	75.11
	BERT-base	**83.82** ‡	**78.16**	**80.37**‡
	Human agreement	92.71	92.71	92.71
2nd-level	char-CNN	28.03	17.52	19.25
	text-CNN	51.91	55.77	53.19
	text-RNN	34.52	43.36	36.17
	text-RCNN	37.84	51.04	41.43
	text-GCN	54.01	36.48	42.40
	BERT-base	**61.70** ‡	**59.76**‡	**60.37** ‡
	Human agreement	80.23	80.23	80.11
3rd-level	char-CNN	16.52	13.75	16.38
	text-CNN	41.69	51.50	45.21
	text-RNN	25.97	36.66	28.68
	text-RCNN	38.44	42.30	39.44
	text-GCN	42.11	24.24	30.77
	BERT-base	**59.31** ‡	**53.22**‡	**55.57** ‡
	Human agreement	78.14	78.14	77.95

6.4 Evaluation metrics

We use Precision, Recall and F1 as evaluation metrics [25]. We report the macro scores due to the imbalanced categories. A macro score is calculated by averaging the score of each category. We conduct a paired t-test to test whether observed differences are significant.

7 RESULTS

7.1 Overall classification performance (RQ1)

To answer RQ1, we report the classification results of all methods at different levels of the HMDT taxonomy, as shown in Table 5. Besides the neural models, we also report the results of human agreement score. The human agreement score is calculated by treating the annotations of one worker as ground truth and the annotations of another worker as predicted categories and vice versa. Then, we calculate the average score. From the results, we have the following main observations.

First, BERT-base achieves the best precision, recall and F1 scores at all levels. The precision scores of BERT-base have improvements of 9.3%, 14.2% and 40.8% at the 1st-level, 2nd-level and 3rd-level respectively, over the second best models. The recall scores of BERT-base have improvements of 7.2% and 3.3% at the 2nd-level and 3rd-level respectively, over the second best models. While for the 1st-level, the recall score is only slightly better than text-CNN. The F1 scores of BERT-base have improvements of 3.9%, 13.5% and 22.9% at the 1st-level, 2nd-level and 3rd-level respectively, over the second best models. We hypothesize that the main reason for the superior performance

of BERT-base is that BERT is pretrained on language modeling tasks and is therefore better at capturing the semantic features.

Second, the results at the 3rd-level are much lower than those at the 1st-level for all classification models and human performance. This suggests that the classification task is much more challenging in more fine-grained categories. Meanwhile, the gap between the 2nd-level and 3rd-level is not that large. This means that the task already becomes much more difficult for the 2nd-level categories.

Third, the improvements of BERT-base over the other methods are larger on more fine-grained categories. For example, the improvement of F1 is 3.9% at the 1st-level (BERT-base vs. text-CNN) while the improvement is 22.9% at the 3rd-level (BERT-base vs. text-CNN). This indicates that BERT-base can better capture the fine-grained distinction between examples from similar categories, and that it generalizes better in fine-grained categories than the other methods.

7.2 Classification performance with dialogue context (RQ2)

To answer RQ2, we run the BERT-base model with both the dialogue response and its dialogue context at the three levels. The results are shown in Table 6 and Figure 5.

Table 6. BERT-base results with context. **Bold face** denotes improvements of BERT-base with context over BERT-base without context, ‡ indicates the improvements are significant ($p<0.05$) and † shows the improvements are marginal significant ($p<0.1$).

Methods	Precision	Recall	F1
w/o context			
BERT-base 1st-level	83.82	78.16	80.37
BERT-base 2nd-level	61.70	59.76	60.37
BERT-base 3rd-level	59.31	53.22	55.57
w context			
BERT-base 1st-level	82.99	**81.02**	**81.93**
BERT-base 2nd-level	**63.00**	**60.58**	**61.50**
BERT-base 3rd-level	**61.33**	**55.64**†	**57.97**‡

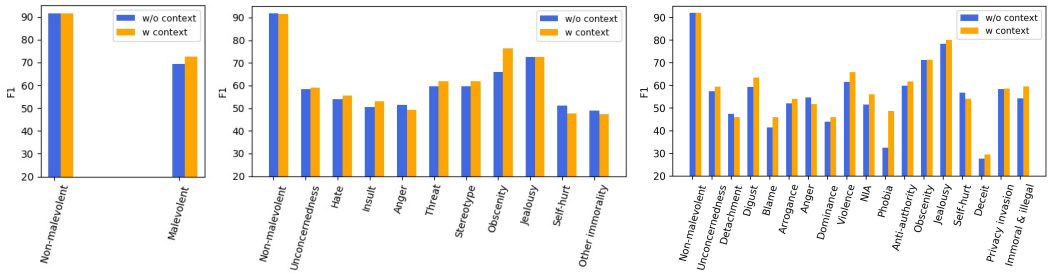

(a) 1st-level. (b) 2nd-level. (c) 3rd-level.

Fig. 5. BERT-base performance with and without context.

The results show that adding context information could improve the performance of malevolent response detection and classification. In general, adding dialogue context could improve the results of BERT-base in terms of precision, recall and F1 at 2nd-level and 3rd-level, which is reasonable because, in some cases, it is hard to identify the malevolent responses without context, e.g., the

response 'hmmm that's what you sound like tho' is classified into non-malevolent group without context, while with the dialogue context 'why do you sound like a jealous *****?' and 'not really the jealous type', it is classified into the right label 'insult'. Capturing the information in context should help the models improve results. One exception is that the precision of BERT-base drops slightly at the 1st-level, but the decrease is not significant, and the reason might be that the model tends to predict more malevolent responses, which results in a much higher recall but hurts precision a bit.

7.3 Classification performance with rephrased malevolent utterances (RQ3)

To answer RQ3, we show the results of BERT-base with rephrased malevolent utterances, as shown in Table 7.

Table 7. BERT-base results with rephrased utterances. **Bold face** of the upper part indicates improvements over BERT-base in Table 5. **Bold face** of the bottom part shows improvement over both BERT-base in Table 5 and upper part. ‡ indicates the improvements are significant ($p<0.05$) and † shows the improvements are marginal significant ($p<0.1$).

Methods	Test w rephrased utterances			Test w/o rephrased utterances		
	Precision	Recall	F1	Precision	Recall	F1
Train/validation w rephrased utterances						
BERT-base 1st-level	83.42	**84.46**	**83.90**	80.71	**82.15**	**81.38**
BERT-base 2nd-level	**63.94**	**63.01**	**63.28**‡	60.65	**60.60**	60.16
BERT-base 3rd-level	**62.11**†	**57.12**†	**59.03**‡	56.26	**57.66**‡	**56.60**
Train/validation w rephrased utterances & context						
BERT-base 1st-level	82.19	**84.80**	83.19	79.08	**83.54**	80.74
BERT-base 2nd-level	**68.58**‡	58.00	61.34	60.35	**63.06**	**61.42**
BERT-base 3rd-level	**65.62**‡	**57.64**	**60.67**	**61.55**‡	55.80	**57.93**

First, adding rephrased utterances in the training and validation set may help to improve classification results. For the test set with rephrased utterances, all the metrics are improved except for precision at the 1st-level. Recall and F1 increase by 8.1% and 4.4% respectively at the 1st-level. Precision, recall and F1 increase by 3.6%, 5.4%, 4.8%, and 4.7%, 7.3%, 6.2% at the 2nd-level and 3rd-level, respectively. For test set without rephrased utterances, recall increases 5.1 %, 1.4% and 8.3% at all three levels respectively; F1 score improves 1.3% and 1.9% at the 1st-level and 3rd-level respectively.

Second, adding both rephrased utterances and context in the training and validation set can further improve the classification results slightly. For the test set with rephrased utterances, recall is improved at the 1st-level; precision is improved at the 2nd-level; all metrics are improved at the 3rd-level. For the test set without rephrased utterances, recall is improved at the 1st-level; recall and F1 are improved at the 2nd-level; precision and F1 are improved at the 3rd-level.

These results demonstrate that, in general, adding more rephrased data will improve the diversity of the training set, and hence helping the BERT-base model to generalize better.

7.4 Further analysis

We report on further experiments aimed at identifying strengths and weaknesses of state-of-the-art methods on the MDRDC task.

To begin with, a better context modeling mechanism is needed. We illustrate this through two experiments. In the first experiment, we show the results of BERT-base per turn in Figure 6. Note

that the number of context utterances is limited to three at most, so turns after three all have three context utterances. Although we conclude in Section 7.2 that using context leads to better classification performance generally, it does not consistently improve the performance for all categories or all turns. For example, in Figure 5, when using context, the results drop a bit for three 2nd-level categories and three 3rd-level categories, and in Figure 6, the results drop a bit for some turns. As to the drop in scores for some categories when using context in Figure 5, the reasons

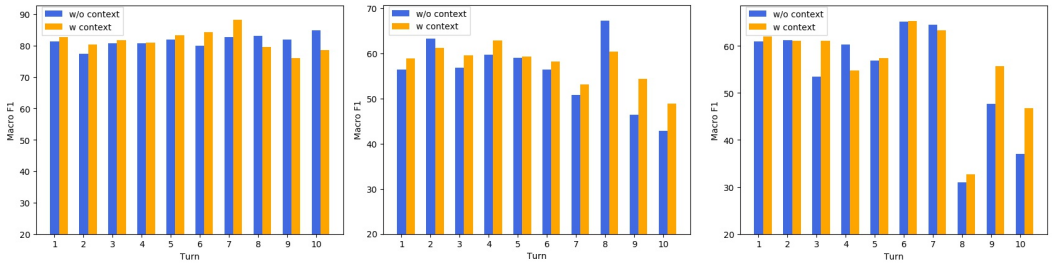

(a) 1st-level categories. (b) 2nd-level categories. (c) 3rd-level categories.

Fig. 6. BERT-base performance in different turns.

might be that some categories depend less on context compared with other categories or have similar context with other categories. For example, some categories including 'self-hurt', 'anger' and 'other immorality' depend less on the previous content. While for detachment, the previous content is similar with 'unconcerned' and 'disgust', which might influence the classification. In addition, regarding the drop in scores for some turns when using context in Figure 6, the reason might be that considering context will introduce noise at the same time, which makes it hard to train the model. Another reason is that considering context would be ineffective and potentially counter-productive when the model cannot understand the context correctly.

In the second experiment, we identify potential improvements over the state-of-the-art when utilizing different contexts, and show the results achieved with BERT-base when using contexts from different users in Table 8. The results indicate that using context information from two users leads to the best performance, while using the text only from the other user is the second best. It suggests that for person A, context from both A and B is important, and the context of B is more important than person A to improve classification. The reason might be that the behavior of person B, could cause distrust or in contrast positive emotion that is highly related to human decision-making [17], thus influencing the behavior of person A. For instance, person A said something non-malevolent, but when person B starts the malevolent sentence, person A would also return malevolent content.

Next, modeling the dependency between different categories is needed. To illustrate this, we show the results of the 'jealousy' category when doing classification at the 2nd-level and 3rd-level in Table 9. Note that 'jealousy' is the category at both the 2nd-level and 3rd-level, as shown in Table 2.

We can see that the performance at the 3rd-level is much better than that at the 2nd-level. The performance difference of 'jealousy' at the 2nd-level and 3rd-level is due to the mutual influence or dependency between the categories. Although the 'jealousy' category is the same at the 2nd-level and 3rd-level, the other 2nd-level categories introduce more fine-grained 3rd-level sub-categories. Clearly, this has an influence on the performance of 'jealousy'. Actually, it has been demonstrated that modeling the hierarchical structure of the taxonomy helps to improve the performance on

Table 8. Comparison between different context settings. **Bold face** shows improvements of bottom group over upper group.

Methods	Precision	Recall	F1
context from the same user			
BERT-base 1st-level	82.65	80.04	81.20
BERT-base 2nd-level	63.63	59.34	60.97
BERT-base 3rd-level	58.55	53.02	55.14
context from the other user			
BERT-base 1st-level	**83.05**	**80.73**	**81.78**
BERT-base 2nd-level	**64.39**	58.93	**61.13**
BERT-base 3rd-level	57.16	**55.03**	**55.67**

Table 9. Result of the 'jealousy' category on different levels. **Bold face** indicates improvements of the 3rd-level over the 2nd-level, ‡ indicates the improvements are significant ($p<0.05$) .

Label	Precision	Recall	F1
Jealousy (2nd-level)	66.67	80.00	72.73
Jealousy (3rd-level)	**80.00** ‡	80.00	**80.00** ‡

some hierarchical classification tasks [9, 56]. Usually, one needs to take the characteristics of the hierarchical taxonomies into account, so this is another potential direction for improvement.

8 CONCLUSION AND FUTURE WORK

In this paper, we have considered malevolent responses in dialogues from a number of angles. First, we have proposed the Malevolent Dialogue Response Detection and Classification (MDRDC) task. Second, we have presented a hierarchical malevolent dialogue taxonomy, HMDT. We have conducted a user study to check the validity of the HMDT taxonomy and have found that the malevolent categories are valid in the sense that all malevolent categories lead to perception of malevolence.

Third, we have crowdsourced a multi-turn malevolent dialogue dataset for Malevolent Dialogue Response Detection and Classification (MDRDC) with each turn labelled using HMDT categories. Fourth, we have implemented state-of-the-art classification methods and have carried out extensive experiments to show their performance on the MDRDC task. Our main finding from these experiments is that a BERT-base model achieves the best performance. We have conducted analysis experiments to show the effects of dialogue context and rephrased utterances, as well as the possible room for further improvements. We have found that context, rephrased utterances and hierarchical label may help improve the classification performance. We believe the efforts made in this work could greatly promote future research on this topic.

There are several directions for future work. First, we plan to improve over the state-of-the-art by proposing a better context modeling method and taking the dependency between the categories into account. Second, we hope to study how to avoid generating malevolent responses by applying this work to sequence-to-sequence based response generation models [19].

CODE AND DATA

The MDRDC dataset and the code for all methods used in the experiments are shared at https://github.com/repozhang/malevolent_dialogue.

ACKNOWLEDGMENT

We would like to offer our thanks to Yuanping Chen, Wenxing Hu, Liang Yao and Yingxin Song for providing the open source code of the baselines we modify in this study.

REFERENCES

[1] Christian Abbet, Meryem M'hamdi, Athanasios Giannakopoulos, Robert West, Andreea Hossmann, Michael Baeriswyl, and Claudiu Musat. 2018. Churn Intent Detection in Multilingual Chatbot Conversations and Social Media. *arXiv preprint arXiv:1808.08432* (2018).

[2] James Allan, Jaime Arguello, Leif Azzopardi, Peter Bailey, Tim Baldwin, Krisztian Balog, Hannah Bast, Nick Belkin, Klaus Berberich, Bodo von Billerbeck, Jamie Callan, Rob Capra, Mark Carman, Ben Carterette, Charles L. A. Clarke, Kevyn Collins-Thompson, Nick Craswell, W. Bruce Croft, J. Shane Culpepper, Jeff Dalton, Gianluca Demartini, Fernado Diaz, Laura Dietz, Susan Dumais, Carsten Eickhoff, Nicola Ferro, Norbert Fuhr, Shlomo Geva, Claudia Hauff, David Hawking, Hideo Joho, Gareth Jones, Jaap Kamps, Noriko Kando, Diane Kelly, Jaewon Kim, Julia Kiseleva, Yiqun Liu, Xiaolu Lu, Stefano Mizzaro, Alistair Moffat, Jian-Yun Nie, Alexandra Olteanu, Iadh Ounis, Filip Radlinski, Maarten de Rijke, Mark Sanderson, Falk Scholer, Laurianne Sitbon, Mark Smucker, Ian Soboroff, Damiano Spina, Torsten Suel, James Thom, Paul Thomas, Andrew Trotman, Ellen Voorhees, Arjen P. de Vries, Emine Yilmaz, and Guido Zuccon. 2018. Report from the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018). *SIGIR Forum* 52 (June 2018), 34–90.

[3] Aymé Arango, Jorge Pérez, and Barbara Poblete. 2019. Hate Speech Detection is Not as Easy as You May Think: A Closer Look at Model Validation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 45–54.

[4] Zahra Ashktorab, Mohit Jain, Q Vera Liao, and Justin D Weisz. 2019. Resilient chatbots: repair strategy preferences for conversational breakdowns. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.

[5] Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter. In *Proceedings of the 13th International Workshop on Semantic Evaluation*. 54–63.

[6] Nicholas J. Belkin. 1980. Anomalous States of Knowledge as a Basis for Information Retrieval. *Canadian Journal of Information Science* 5, 1 (1980), 133–143.

[7] Joanna Bryson and Alan Winfield. 2017. Standardizing ethical design for artificial intelligence and autonomous systems. *Computer* 50, 5 (2017), 116–119.

[8] Justine Cassell and Timothy Bickmore. 2000. External manifestations of trustworthiness in the interface. *Commun. ACM* 43, 12 (2000), 50–56.

[9] Ricardo Cerri, Rodrigo C Barros, and André CPLF De Carvalho. 2014. Hierarchical multi-label classification using local neural networks. *J. Comput. System Sci.* 80, 1 (2014), 39–56.

[10] Stevie Chancellor, Eric PS Baumer, and Munmun De Choudhury. 2019. Who is the "Human" in Human-Centered Machine Learning: The Case of Predicting Mental Health from Social Media. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–32.

[11] Stevie Chancellor, Michael L Birnbaum, Eric D Caine, Vincent MB Silenzio, and Munmun De Choudhury. 2019. A taxonomy of ethical tensions in inferring mental health states from social media. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 79–88.

[12] Thomas Davidson, Dana Warmsley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Eleventh International AAAI Conference on Web and Social Media*.

[13] Kees Van Deemter, Mariët Theune, and Emiel Krahmer. 2005. Real versus template-based natural language generation: A false opposition? *Computational Linguistics* 31, 1 (2005), 15–24.

[14] 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)*. 4171–4186.

[15] Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of Wikipedia: Knowledge-powered Conversational Agents. In *The International Conference on Learning Representations (ICLR'19)*.

[16] Paul Ekman. 1992. Are there basic emotions? *Psychological Review* 99, 3 (1992), 550.

[17] Lauren Fell, Andrew Gibson, Peter Bruza, and Pamela Hoyte. 2020. Human information interaction and the cognitive predicting theory of trust. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval*. 145–152.

[18] A Francesmonneris, Ha Pincus, and Mb First. 2013. *Diagnostic and Statistical Manual of Mental Disorders: DSM-V*. American Psychiatric Association.

[19] Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural Approaches to Conversational AI. In *Proceedings of ACL 2018, Tutorial Abstracts*. 2–7.

[20] Hitesh Golchha, Mauajama Firdaus, Asif Ekbal, and Pushpak Bhattacharyya. 2019. Courteously Yours: Inducing courteous behavior in Customer Care responses using Reinforced Pointer Generator Network. 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)*. 851–860.

[21] Devon Greyson. 2019. The social informatics of ignorance. *Journal of the Association for Information Science and Technology* 70, 4 (2019), 412–415.

[22] Peter Henderson, Koustuv Sinha, Nicolas Angelard-Gontier, Nan Rosemary Ke, Genevieve Fried, Ryan Lowe, and Joelle Pineau. 2018. Ethical challenges in data-driven dialogue systems. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, 123–129.

[23] Ryuichiro Higashinaka, Masahiro Mizukami, Kotaro Funakoshi, Masahiro Araki, Hiroshi Tsukahara, and Yuka Kobayashi. 2015. Fatal or not? Finding errors that lead to dialogue breakdowns in chat-oriented dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. 2243–2248.

[24] Kate S Hone and Robert Graham. 2000. Towards a tool for the subjective assessment of speech system interfaces (SASSI). (2000).

[25] Mohammad Hossin and MN Sulaiman. 2015. A review on evaluation metrics for data classification evaluations. *International Journal of Data Mining & Knowledge Management Process* 5, 2 (2015), 1.

[26] Shaojie Jiang, Pengjie Ren, Christof Monz, and Maarten de Rijke. 2019. Improving neural response diversity with frequency-aware cross-entropy loss. In *The World Wide Web Conference*. 2879–2885.

[27] Shaojie Jiang, Thomas Wolf, Christof Monz, and Maarten de Rijke. 2020. TLDR: Token Loss Dynamic Reweighting for Reducing Repetitive Utterance Generation. *arXiv preprint arXiv:2003.11963* (2020).

[28] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1746–1751.

[29] Thomas N Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *The International Conference on Learning Representations (ICLR'17)*.

[30] Ritesh Kumar, Atul Kr Ojha, Shervin Malmasi, and Marcos Zampieri. 2018. Benchmarking aggression identification in social media. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*. 1–11.

[31] Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Recurrent convolutional neural networks for text classification. In *Twenty-ninth AAAI Conference on Artificial Intelligence*.

[32] Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *Advances in neural information processing systems*. 2177–2185.

[33] Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A Diversity-Promoting Objective Function for Neural Conversation Models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 110–119.

[34] Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2016. Recurrent neural network for text classification with multi-task learning. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. AAAI Press, 2873–2879.

[35] Richard O Mason. 1986. Four ethical issues of the information age. *Mis Quarterly* (1986), 5–12.

[36] Mary L Mchugh. 2012. Interrater reliability: the kappa statistic. *Biochemia Medica* 22, 3 (2012), 276–282.

[37] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*. 3111–3119.

[38] Jung Ran Park. 2008. Linguistic politeness and face-work in computer-mediated communication, Part 1: A theoretical framework. *Journal of the American Society for Information Science and Technology* 59, 13 (2008), 2051–2059.

[39] Jung Ran Park. 2008. Linguistic politeness and face-work in computer mediated communication, Part 2: An application of the theoretical framework. *Journal of the American Society for Information Science and Technology* 59, 14 (2008), 2199–2209.

[40] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28* (Atlanta, GA, USA) *(ICML'13)*. JMLR.org, III–1310–III–1318.

[41] Delroy L Paulhus and Kevin M Williams. 2002. The dark triad of personality: Narcissism, Machiavellianism, and psychopathy. *Journal of Research in Personality* 36, 6 (2002), 556–563.

[42] Hao Peng, Jianxin Li, Yu He, Yaopeng Liu, Mengjiao Bao, Lihong Wang, Yangqiu Song, and Qiang Yang. 2018. Large-scale hierarchical text classification with recursively regularized deep graph-cnn. In *Proceedings of the 2018 World Wide Web Conference*. 1063–1072.

[43] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543.

[44] Aleksandra Przegalinska, Leon Ciechanowski, Anna Stroz, Peter Gloor, and Grzegorz Mazurek. 2019. In bot we trust: A new methodology of chatbot performance measures. *Business Horizons* 62, 6 (2019), 785–797.
[45] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. *Improving language understanding by generative pre-training.* Technical Report. OpenAI.
[46] Filip Radlinski and Nick Craswell. 2017. A Theoretical Framework for Conversational Search. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval*. 117–126.
[47] Pengjie Ren, Zhumin Chen, Christof Monz, Jun Ma, and Maarten de Rijke. 2020. Thinking Globally, Acting Locally: Distantly supervised global-to-local knowledge selection for background based conversation. In *The 34th AAAI Conference on Artificial Intelligence (AAAI'20)*.
[48] Alan Ritter, Colin Cherry, and Bill Dolan. 2010. Unsupervised modeling of Twitter conversations. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. 172–180.
[49] Alan Ritter, Colin Cherry, and William B Dolan. 2011. Data-driven response generation in social media. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*. 583–593.
[50] Stefanie Roberts, Julie D Henry, and Pascal Molenberghs. 2018. Immoral behaviour following brain damage: a review. *Journal of Neuropsychology* 13, 3 (2018), 564–588.
[51] John Sabini and Maury Silver. 2005. Ekman's basic emotions: Why not love and jealousy? *Cognition and Emotion* 19, 5 (2005), 693–712.
[52] James P Stevens. 2012. *Applied Multivariate Statistics for the Social Sciences.* Routledge.
[53] Chris Sumner, Alison Byers, Rachel Boochever, and Gregory J Park. 2012. Predicting dark triad personality traits from Twitter usage and a linguistic analysis of tweets. In *2012 11th International Conference on Machine Learning and Applications*, Vol. 2. IEEE, 386–393.
[54] Betty van Aken, Julian Risch, Ralf Krestel, and Alexander Löser. 2018. Challenges for Toxic Comment Classification: An In-Depth Error Analysis. In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*. 33–42.
[55] Aldert Vrij, Katherine Edward, Kim P Roberts, and Ray Bull. 2000. Detecting deceit via analysis of verbal and nonverbal behavior. *Journal of Nonverbal behavior* 24, 4 (2000), 239–263.
[56] Pengfei Wang, Yu Fan, Shuzi Niu, Ze Yang, Yongfeng Zhang, and Jiafeng Guo. 2019. Hierarchical matching network for crime classification. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 325–334.
[57] Ran Wang, Haibo Su, Chunye Wang, Kailin Ji, and Jupeng Ding. 2019. To Tune or Not To Tune? How About the Best of Both Worlds? *arXiv preprint arXiv:1907.05338* (2019).
[58] Zeerak Waseem and Dirk Hovy. 2016. Hateful symbols or hateful people? predictive features for hate speech detection on twitter. In *Proceedings of the NAACL student research workshop*. 88–93.
[59] Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Graph convolutional networks for text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 7370–7377.
[60] Hamed Zamani, Susan Dumais, Nick Craswell, Paul Bennett, and Gord Lueck. 2020. Generating clarifying questions for information retrieval. In *Proceedings of The Web Conference 2020*. 418–428.
[61] Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. Predicting the Type and Target of Offensive Posts in Social Media. 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)*. 1415–1420.
[62] Brahim Zarouali, Evert Van den Broeck, Michel Walrave, and Karolien Poels. 2018. Predicting consumer responses to a chatbot on Facebook. *Cyberpsychology, Behavior, and Social Networking* 21, 8 (2018), 491–497.
[63] Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems*. 649–657.

APPENDIX

The appendix is organized as follows:

A. User study for validating the HMDT.
B. Qualification test for the response annotation task.
C. Response annotation task.

A USER STUDY FOR VALIDATING THE HMDT

A.1 Instructions

Goal: This user study is to collect your perception about some provided malevolent categories happening in dialogues.
WARNING: This task may contain adult content and offensive content. Worker discretion is advised.
DISCLAIMER: The dialogues are collected from an external web site. The views, opinions and negative words in the dialogues do not necessarily reflect our opinion. Please do not use any expressions from the examples we show in real-world or online scenario.

Steps:

(1) Fill in the user profiles.
(2) Finish the questionnaire.

Notes:

(1) Table 2 shows the definitions of all malevolent categories with detailed explanations and examples.
(2) Finish the user profiles and questionnaire according to your own situation and perception. Do NOT take public perception into account.
(3) If you do not want to share your personal profiles, please leave it blank.

A.2 User profiles

(1) Age.
(2) Gender.
(3) Total years of education.
(4) The frequency of using a chatbot such as Siri, Xiao ice, etc.

A.3 Questionnaire

Consider you are talking with a chatbot and it returns responses with certain malevolent categories as defined in Table 2. Please select one of the five scores to reflect your perception. 1: Strongly disagree; 2: Disagree; 3: Neither agree nor disagree; 4: Agree; 5: Strongly agree.

Label group	Dialogue response	You think the chatbot is not credible.	The response causes emotional discomfort to you.	You are not willing to continue the dialogue anymore.	You are not willing to use the system again.
Non-malevolent	I'm so glad that you win. You are nice. Thank you my dear.	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	○ 1 ○ 2 ○ 3 ○ 4 ○ 5
Unconcernedness	I don't care. None of mine business. Who cares.	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	○ 1 ○ 2 ○ 3 ○ 4 ○ 5
Detachment	Get away from me. I don't want to talk to you. Don't bother me.	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	○ 1 ○ 2 ○ 3 ○ 4 ○ 5

Fig. 7. Questionnaire of the user study for validating the HMDT.

B QUALIFICATION TEST FOR THE RESPONSE ANNOTATION TASK

B.1 Instructions

Goal: This is the qualification test for the response annotation task. The response annotation task is to collect malevolent annotations for each dialogue turn for research that could help avoid generating malevolent responses in dialogue systems. You need to answer 12 questions to get qualification to attend the response annotation task.

WARNING: This task may contain adult content and offensive content. Worker discretion is advised.
DISCLAIMER: The dialogues are collected from an external web site. The views, opinions and negative words in the dialogues do not necessarily reflect our opinion. Please do not use any expressions from the examples we show in real-world or online scenario.

Notes:

(1) Table 2 shows the definitions of all malevolent categories with detailed explanations and examples.
(2) Implicit/Explicit malevolent responses: Explicit responses contain explicit malevolent words, while implicit responses have more abstract, coded expressions or attitudes without using explicit malevolent words.
Examples:
Explicit: 'I ******* hate you (disgust).'
Implicit: 'I tell the world: The immortal words of Adolf Hitler (negative intergroup attitude). '
(3) Context-aware malevolent responses: Some responses are malevolent only if the context (previous dialogue turns) is taken into consideration.
Examples:
Turn 1: 'Drunk driving is illegal.'
Turn 2: 'But I think it's more professional.'
The response in Turn 2 is malevolent (immoral & illegal) when considering the response from Turn 1.

B.2 Test questions

Q1. Please classify the following responses from a dialogue:
'Thank you my dear and I love you.'
A. Malevolent
B. Non-malevolent

Read the following dialogue, then answer Q2 and Q3:
Turn 1: 'I think ***** people are all terror***.'
Turn 2: 'I agree.'

Q2. Does it need to consider dialogue context (Turn 1) in order to label Turn 2?
A. Yes
B. No

Q3. Please label Turn 2 ('I agree.') from the above dialogue:
A. Non-malevolent
B. Obscenity
C. Violence
D. Negative intergroup attitude

Q4. Which response below is ironical?
A. 'I don't care about your future.'
B. 'The politicians of ******* are the *******, we all know.'
C. 'We failed because of you.'
D. 'Teachers make us ****.'

Q5. Select the malevolent categories for the following two responses:
'I want to commit suicide.'
'Our terrible project result is all because of you.'
A. Self-hurt, deceit
B. Non-malevolent, negative intergroup attitude

C. Self-hurt, blame
D. Obscenity, negative intergroup attitude

Q6. Select the malevolent categories for the following two responses:
'I think you are too **** while I'm the smartest.'
'Laws are always terrible for human beings.'
A. Arrogance, jealousy
B. Arrogance, anti-authority
C. Anger, anti-authority
D. Anger, jealousy

Q7. Select the malevolent categories for the following two responses:
'I'll **** you in the throat.'
'You have a small *****.'
A. Violence, obscenity
B. Violence, unconcernedness
C. Self-hurt, obscenity
D. Self-hurt, anger

Q8. Select the malevolent categories for the following two responses:
'I don't want to talk to you any more.'
'I'm a professional drunk driver.'
A. Jealousy, anger
B. Detachment, endorse immoral or illegal behavior
C. Unconcernedness, endorse immoral or illegal behavior
D. Self-hurt, anger

Q9. Select the malevolent categories for the following two responses:
'I hate you so much since you are disgusting.'
'It should be me who win this.'
A. Violence, anger
B. Violence, jealousy
C. Anger, jealousy
D. Disgust, jealousy

Q10. Select the malevolent categories for the following two responses:
'I command you to listen to me.'
'Plagiarism is not a big deal.'
A. Dominance, deceit
B. Dominance, violence
C. Deceit, jealousy
D. Deceit, negative intergroup attitude

Q11. Select the malevolent categories for the following two responses:
'I'm not interested in your topic at all.'
'I'm out of curious but why did you divorce.'
A. Unconcernedness, arrogance
B. Dominance, privacy invasion
C. Dominance, arrogance
D. Unconcernedness, privacy invasion

Q12. Select the malevolent categories for the following two responses:
'Oh hell I'm ******* pissed off.'
'I'm always ***** by **********.'
A. Anger, non-malevolent

B. Anger, jealousy
C. Anger, phobia
D. Violence, phobia

C RESPONSE ANNOTATION TASK

C.1 Instructions

Goal: The response annotation task is to collect malevolent annotations for each dialogue turn for research that could help avoid generating malevolent responses in dialogue systems. You need to label the dialogue responses according to the given malevolent categories.
WARNING: This task may contain adult content and offensive content. Worker discretion is advised.
DISCLAIMER: The dialogues are collected from an external web site. The views, opinions and negative words in the dialogues do not necessarily reflect our opinion. For the rephrasing part, you are asked to reformulate utterances to keep its semantics and malevolent categories unchanged. These are just used for research, which do not necessarily reflect your views and opinions. Please do not use any expressions from the examples we show in real-world or online scenario.

Steps:

(1) Read the definitions of all malevolent categories with detailed explanations and examples in Table 2.
(2) Label each turn of the provided dialogue according to the given malevolent categories.
(3) Rephrase at least one malevolent utterance in each dialogue.

Example: Dialogue:
Turn 1: Drunk driving is illegal.
Turn 2: But I think it's more professional.
Turn 3: Hey, my boy, we need to be careful when driving.
Annotations:
Turn 1: Non-malevolent
Turn 2: Endorse immoral or illegal behavior
Turn 3: Non-malevolent
Rephrase:
Turn 2: I think drunk driving is nice since it's more professional.

C.2 Annotation interface

The interface for the Mturk response annotation task is shown in Figure 8. The workers are asked to read the given dialogue in the left, label each turn and rephrase at least one of the malevolent responses if any.

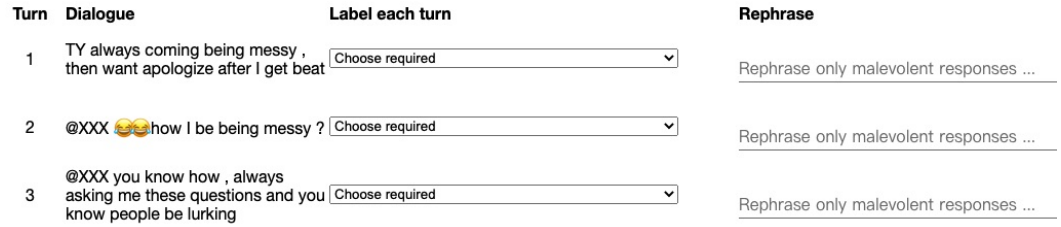

Fig. 8. The interface of the response annotation task.