

# CONDA: a CONtextual Dual-Annotated dataset for in-game toxicity understanding and detection

Henry Weld    Guanghao Huang    Jean Lee    Tongshu Zhang    Kunze Wang  
 Xinghong Guo    Siqiu Long    Josiah Poon    Soyeon Caren Han\*

School of Computer Science, The University of Sydney, NSW, Australia

{hgua3108, tzha6458, kwan4418, xguo2796, slon6753}@uni.sydney.edu.au

{henry.weld, jean.lee, josiah.poon, caren.han}@sydney.edu.au

## Abstract

Traditional toxicity detection models have focused on the single utterance level without deeper understanding of context. We introduce CONDA, a new dataset for in-game toxic language detection enabling joint intent classification and slot filling analysis, which is the core task of Natural Language Understanding (NLU). The dataset consists of 45K utterances from 12K conversations from the chat logs of 1.9K completed *Dota 2* matches. We propose a robust dual semantic-level toxicity framework, which handles utterance and token-level patterns, and rich contextual chatting history. Accompanying the dataset is a thorough in-game toxicity analysis, which provides comprehensive understanding of context at utterance, token, and dual levels. Inspired by NLU, we also apply its metrics to the toxicity detection tasks for assessing toxicity and game-specific aspects. We evaluate strong NLU models on CONDA, providing fine-grained results for different intent classes and slot classes. Furthermore, we examine the coverage of toxicity nature in our dataset by comparing it with other toxicity datasets.<sup>1</sup>

## 1 Introduction

As the popularity of multi-player online games has grown, the phenomenon of in-game toxic behavior has taken root within them. Toxic behavior is strongly present in recent online games and is problematic to the gaming industry (Adinolf and Turkay, 2018). For instance, 74% of US players of such games report harassment with 65% experiencing severe harassment. (ADL, 2019).

In the past few years, Natural Language Processing (NLP) researchers have proposed several online game/community toxicity analysis frameworks

\*Corresponding author (caren.han@sydney.edu.au)

<sup>1</sup>The dataset and lexicons are available at <https://github.com/usydnlp>.

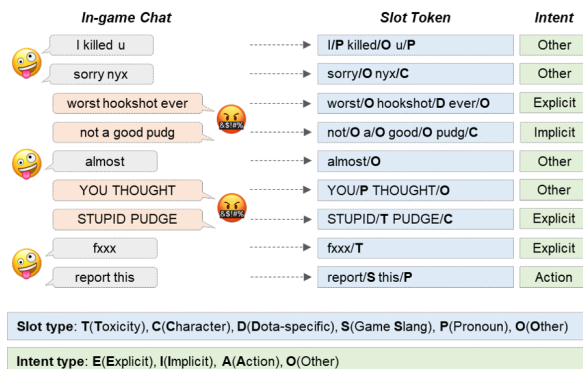


Figure 1: An example intent/slot annotation from the CONDA (CONtextual Dual-Annotated) dataset.

(Kwak et al., 2015; Murnion et al., 2018; Wang et al., 2020) and datasets (Märtens et al., 2015; Stoop et al., 2019). However, existing datasets (1) focus only on the single utterance level without deeper understanding of context in the whole conversation/chat, and (2) do not explicitly use semantic clues from the words within the utterance.

The chat in online games and communities is similar in nature to spoken language, an area studied by Natural Language Understanding (NLU). NLU research aims to best represent human communication by extracting semantic structure in the form of intent and slot analysis. Intent detection is the classification of the desired outcome of an utterance (or sentence), and slot filling is the labeling of each token (or word) in the utterance with the type of semantic information it carries. In recent literature, these two tasks are trained jointly to capture synergies between them, and these jointly trained models give better results (Zhang et al., 2019b). Furthermore, researchers have made available joint task datasets that contain the context of a multi-turn conversation (Budzianowski et al., 2018; Schuster et al., 2019)

Inspired by this NLU research progress, we

propose CONDA, an in-game toxicity detection dataset, with a robust dual-level annotation which enables intent detection and slot filling. Our dataset consists of 45k utterances from the chat logs of 1.9k *Dota 2* matches, labeled with 4 intent classes and 6 slot classes to address toxicity and the game-specific vocabulary. Figure 1 illustrates an example of CONDA including raw data (in-game chat) and processed data with slot and intent labels. In order to enable the dual semantic-level framework, we conduct lexicon-based automation for token-level data and human annotation for utterance-level data.

We investigate the CONDA dataset through an in-depth analysis. The large portion of game-specific classes in the dual levels enables the dataset to be more sophisticated in detecting toxicity in games. The combination of each intent with each slot class shows that dual annotation can help determine toxicity from gamer slang when used in both toxic and non-toxic situations. We also find more toxic utterances appear pre-game and post-game rather than during the games, especially peaking post-game due to the chat for post-victory celebration and recrimination.

We provide five strong baseline NLU models and compare the toxicity detection performance over our dataset. For evaluation, we apply four NLU metrics to assess performance in toxicity and game specific aspects. Results vary across models, indicating a challenge for improvement. Furthermore, we perform a transfer learning experiment with existing toxicity datasets. We find that the nature of toxicity in our dataset can generalize to other proposed taxonomies, including hatefulness, sexism and racism. Beyond this commonality, our experiment illustrates that CONDA is distinguished from other toxicity datasets due to game-specific characteristics. This paper then makes the following contributions:

- To the best of our knowledge, this is the first attempt to build a toxicity detection dataset with joint Natural Language Understanding aspects of intent classification and slot filling;
- We propose a robust dual semantic-level toxicity framework, which handles utterance and token-level patterns with rich in-game chatting history;
- We formalise NLU metrics for toxicity detection, evaluate strong NLU models on our dataset, and further conduct transfer learning experiments with other toxicity datasets.

## 2 Related Work

**Toxicity Datasets in Online Games** In multi-player online games, prior research focused on analysis of anti-social or disruptive behavior, so-called toxic behavior (Blackburn and Kwak, 2014; de Mesquita Neto and Becker, 2018) including cyberbullying (Kwak et al., 2015) and grieving (Murnion et al., 2018). Although these terms contain similar elements, a single definition of toxic behavior is yet to emerge. Some studies have conducted data annotation using pre-defined lexicon categories (Märtens et al., 2015) or toxic player information (Stoop et al., 2019). These annotation methods are not robust enough to handle unlabelled toxicity words or unreported toxic players.

**Toxicity Datasets in Online Community** An extensive body of work has focused on datasets to detect toxicity including hate speech (Waseem and Hovy, 2016; Davidson et al., 2017; ElSherief et al., 2018) and abusive language (Nobata et al., 2016; Founta et al., 2018). However, the majority of toxicity datasets do not consider the context of a conversation, instead simply analysing a single utterance. Even if a model uses contextual information (Gao and Huang, 2017), it is limited to meta-information (e.g. news title or user name) which is not sufficient to understand a conversation. In our research, context is defined as linguistic contextual information, particularly previous single or multiple utterances. Along similar lines, recent studies have focused on conversation aiming to discover warning signals (Zhang et al., 2018), to generate intervention responses (Qian et al., 2019), or to measure the importance of context (Pavlopoulos et al., 2020). Existing toxicity datasets mainly focus on annotating at utterance-level, whereas ours conducts a dual-level annotation at utterance and token-level, while also providing a conversation history (see Table 1). These extra features are what distinguish CONDA.

**NLU Datasets and Models** In-game chat has similar characteristics to multi-turn dialogue in NLU. The approaches used in multi-turn dialogue analysis have not yet been observed in toxicity datasets. In NLU, generally, intent classification (IC) is treated as a semantic utterance classification task and slot filling (SF) is treated as a sequential token labelling task (Zhang and Wang, 2016). By conducting a joint model for the two tasks, a synergistic effect can be achieved (Zhang et al.,

Dataset	Approach	Domain	Labels	Conv.
(Märtens et al., 2015)	utterance-level	Game (Dota 2)	toxic, non-toxic	N
(Waseem and Hovy, 2016)	utterance-level	Twitter	racist, sexist, normal	N
(Nobata et al., 2016)	utterance-level	Yahoo News	clean, hate, derogatory, profanity	N
(Davidson et al., 2017)	utterance-level	Twitter	hateful, offensive, neither	N
(Gao and Huang, 2017)	utterance-level	Fox News	hate, non-hate	N
(ElSherief et al., 2018)	utterance-level	Twitter	hate, non-hate / hate instigator, hate target	N
(Founta et al., 2018)	utterance-level	Twitter	offensive, abusive, hateful speech, aggressive, cyberbullying, spam, normal	N
(Zhang et al., 2018)	utterance-level	Wikipedia	toxic, non-toxic	Y
(Stoop et al., 2019)	utterance-level	Game (LoL)	toxic, non-toxic	Y
(Qian et al., 2019)	utterance-level	Gab & Reddit	hate, non-hate	Y
(Pavlopoulos et al., 2020)	utterance-level	Wikipedia	toxic, non-toxic	Y
<b>CONDA (our dataset)</b>	<b>dual-level (utterance and token)</b>	Game (Dota 2)	- utterance level (intent): explicit toxicity, implicit toxicity, action, others - token level (slot): toxicity, character, dota-specific, slang, pronoun, other	<b>Y</b>

Table 1: Comparison of CONDA with other toxicity datasets (Conv.: Conversation).

2019b). To build multi-turn dialogue datasets, most studies have recruited workers via crowd-sourcing to collect task-oriented dialogues across different domains (e.g. in-car assistant (Eric et al., 2017), navigation and events (Gupta et al., 2018), multi-domains (Budzianowski et al., 2018), personal notifications (Schuster et al., 2019)). Recently, deep learning models have also been extensively studied in order to capture the contextual signals from multiple sequential inputs. (e.g. BiLSTM with attention (Wang et al., 2019), GRU with self-attention and context-fusion (Gupta et al., 2019)). The models listed all show an increase in semantic detection performance when the context is included in the analysis.

### 3 CONDA

#### 3.1 Data Collection

Our annotated dataset, CONDA, is based on the *Defense of the Ancients 2 (Dota 2)* data dump available at Kaggle<sup>2</sup>. *Dota 2* is a multiplayer online game where teams of five players attempt to destroy their opponents’ ancient structure. The raw data is compiled from game matches including players, duration, match outcomes, and complete chat logs. In order to curate data, we select 50,000 utterances in complete chat logs from 1,921 matches.

#### 3.2 Data Processing

Our data processing is designed to enable dual annotation, making utterance-level data suitable for human annotators and generating token-level data for lexicon-based automation. The main processes are creation of conversations, restructuring utter-

ances while keeping original context, and generation of tokens.

We generate conversations to give human annotators a context of previous utterances when labelling the current utterance. We identify the beginning of a conversation as the first utterance in the match, or an utterance that occurs greater than 60 seconds after the previous utterance in the match. While the raw data is largely in English, other languages appear occasionally including Russian, Chinese, Spanish, etc. We exclude conversations with chat in non-English.

For the utterance-level data, we maintain the original form such as punctuation and case in order to keep context. In addition, we merge consecutive utterances by a single user within a conversation. These are combined into one utterance with a special token, [SEPA], added to denote the separation point (e.g. ‘*easiest [SEPA] game [SEPA] of my life*’). For the token-level data, we use contraction restoration (e.g. ‘*I’m*’ -> ‘*I am*’), whitespace tokenise each utterance, retain emoticons, but remove punctuation. This token-level processing is used for lexicon-based automated slot annotation.

Our final CONDA dataset (Table 2) consists of 44,869 utterances and 1,921 matches. We further create a subset, equivalent to about 10% of the full dataset, for a preliminary round of utterance-level annotation.

Dataset Feature	CONDA
Matches	1,921
Conversations	12,152
Utterances	44,869
Avg. utterances per match	23.3

Table 2: CONDA statistics.

<sup>2</sup><https://www.kaggle.com/devinanzelmo/dota-2-matches>

### 3.3 Annotation

**Dual Aspects** Inspired by NLU, we provide a dual-level annotation approach to detect toxicity, which often relies on context. This allows one to find toxic intent even though an utterance does not contain any toxic words, or to determine non-toxic intent even if an utterance has toxic words. For example, Figure 1 shows an utterance “*not a good pudg*”, which does not contain any toxic words. However, considering the previous utterance of “*worst hookshot ever*”, we can identify hidden or implicit toxicity. As an example of the other way around, an utterance of “*happy fuck you day*” contains a toxic word but it is used for cheering after saying “*gg*” (*good game*).

**Token-level Slot Annotation** With the processed token-level data, an automated slot labelling is performed. Initially, we create six distinct slot labels: T (**T**oxicity), C (**C**haracter), D (**D**ota-specific), S (game **S**lang), P (**P**ronoun) and O (**O**ther). To construct the T lexicon, we combine several toxicity lexicons (see Section 8 Ethics) and remove overlaps. We also use the supplemental data sourced by Märtens et al. (2015) for the game-related lexicons (C, D and S) and carefully modify it. The P lexicon (e.g. ‘*u*’, ‘*ur*’) is constructed by this research because in-game chat is extremely abbreviated. Then, we perform lexicon-based automation by exact matching each lower-cased token against the lexicons. Anything not matching a lexicon is labelled O. We contrast this with typical NLU slot labelling where a semantic concept can stretch over a span of words. In comparison to other toxicity datasets, our lexicon-based slot labelling enables deeper understanding of game context.

**Utterance-level Intent Annotation** Given tokens with slot labelling and utterance-level data, we perform a test run on the subset of utterances using six annotators. Four annotators are game players and two are non-game players. This preliminary round is for fine-tuning annotation policy and analysing annotator agreement to inform final annotation for the full dataset. The annotators manually classified the utterances into four labels: E (**E**xplicit toxicity), I (**I**mplicit toxicity), A (**A**ction) and O (**O**ther). The label details are explained in the annotator instructions.

**Annotator Instructions** Each annotator was required to consider the earlier conversation, particularly, to detect implicit toxic behavior or to identify

non-toxic behavior in the utterance having toxic-labelled tokens. The annotators worked independently of one another. The guidelines for human annotators were as follows:

**Explicit toxicity:** Typically contains toxic word(s). The intent is to insult or humiliate others, or to make others want to leave the conversation or quit the game. There is no need to consider the context (e.g. ‘*fuck off*’). May include one or more of the following aspects:

- Strong toxicity - blatant insulting or disrespecting others is obviously seen in the text, normally with severely toxic wording;
- Normal toxicity - impolite, rudely worded and unreasonable comment that insults or humiliates others;
- Cursing others with the intent to insult or humiliate them (e.g. ‘*noob*’<sup>3</sup>);
- Sexual wording or talk about sex-related behavior;
- Use of negative or hateful words to describe others (e.g. ‘*useless*’);
- Racist language that is targeted at insulting others (e.g. ‘*Peruvians*’, ‘*fucking russians*’);
- Inflammatory language, insulting others and trying to start a conversational fight.

**Implicit toxicity:** Hidden toxicity that normally cannot be seen from the text itself. The text might be factual or even positive (e.g. sarcasm). However, based on the utterance or conversation context, the intent of insulting or humiliating others can be inferred. Typically contains no toxic word (e.g. ‘*u are poor dude*’).

**Action:** Doesn’t belong to I or E, but contains an action such as report, commend, pause, stop, or exit game.

**Other:** Doesn’t belong to I or E or A. May or may not contain toxic words. Includes curses, self-deprecation or any other emotional expression that is NOT targeted at others (e.g. ‘*kill the fucking helicopter*’).

**Findings in Annotation** The preliminary round was useful for enabling discussion around annotation. For example, we decided that the token “*ez*”<sup>4</sup>

<sup>3</sup>“Noob” is a slang term for a newcomer, commonly used to insult someone inexperienced in games.

<sup>4</sup>“Ez” is an abbreviation for easy. It is often used to irritate other players in games, indicating “You are just way too easy”.

or its variations are an implied slur against the opposition’s quality and would generally be part of an **I** label utterance. Similarly, “g” is often a contraction for “go” and would be part of an **A** label utterance. Overall, we observed that the agreement measure for utterance classification was higher for gamer annotators only (Fleiss’ kappa = 0.785) versus the whole group (Fleiss’ kappa = 0.755). The lower inter-rater agreement in the whole group is because non-gamer annotators have low understanding of the game context and domain-specific language. Therefore, annotation of the whole dataset was performed by gamers only. Based on our annotation guidelines, they collectively manually annotated the utterances for the full dataset.

## 4 Dataset Analysis

The CONDA dataset consists of 9 columns - match ID, conversation ID, player ID, player slot, chat time, utterance, slot tokens (cleaned tokens with slot labelling), intent class, and slot classes. For example, the utterance “gg wp” for “good game well played” is shown in the slot tokens column as “gg (S), wp (S)”. Each column is further explained in Appendix A.

Intent	%	Mean L.
E (Explicit)	13.3	6.14
I (Implicit)	6.4	4.16
A (Action)	6.4	4.40
O (Other)	73.9	3.18
<b>Total</b>	<b>100.0</b>	<b>3.71</b>

Slot	%
T (Toxicity)	4.9
C (Character)	5.4
D (Dota-specific)	1.4
S (Game Slang)	11.2
P (Pronoun)	13.5
O (Other)	63.6
<b>Total</b>	<b>100.0</b>

Table 3: Intent labelling statistics (left) and slot labelling statistics (right). % is proportion of the dataset. Mean L. is mean number of tokens after cleaning.

**Dual Annotation Proportion** Table 3 gives the proportional breakdown of the CONDA dataset by intent and slot labels. Together the toxic utterance classes make up 19.7% of the data, emphasizing their prevalence in game chat. The proportion of the I (6.4%) and A (6.4%) intent classes, together 12.7% of all utterances, shows the more granular non-binary class structure captures an aspect of online games. Additionally, the average length of utterances of the E class (6.14) is greater than for each other class, indicating players strongly emphasize emotional frustration. In the proportion of slot labels, we can see the S (11.2%) class is more than double the C (5.4%) class and 8 times

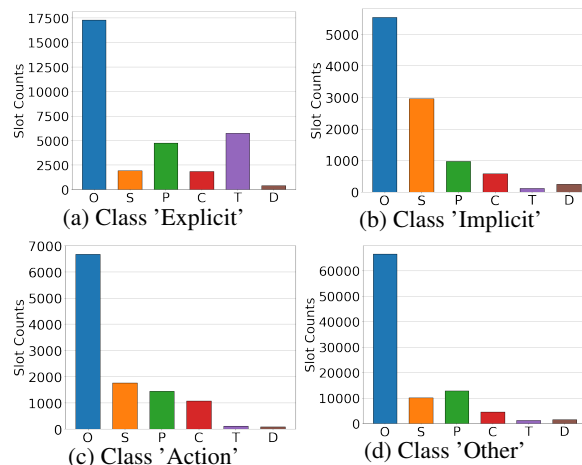


Figure 2: Slot class distributions for each intent class.

Rank	S	T
1	gg (239)	noob (878)
2	report (237)	fuck (807)
3	ez (191)	fucking (593)
4	mid (169)	shit (546)
5	go (114)	idiot (222)

(a) Class "Explicit"

Rank	S	T
1	ez (1,932)	wtf (32)
2	mid (287)	fucking (11)
3	gg (169)	dead (9)
4	report (47)	hook (6)
5	go (38)	fuck (5)

(b) Class "Implicit"

Rank	S	T
1	report (992)	wtf (16)
2	afk (184)	fucking (11)
3	gg (134)	abuse (10)
4	go (57)	noob (8)
5	wp (37)	shit (4)

(c) Class "Action"

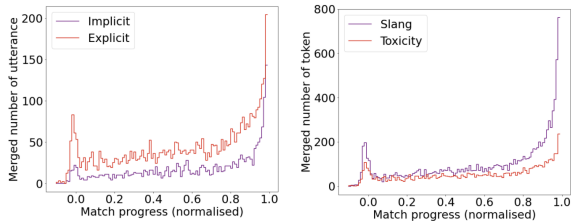
Rank	S	T
1	gg (3,735)	wtf (331)
2	wp (1,115)	dead (89)
3	ggwp (776)	fucking (84)
4	mid (413)	hook (69)
5	go (383)	shit (39)

(d) Class "Other"

Table 4: Top 5 keywords in the S (game Slang) and T (Toxicity) slot classes, for each intent class. The number in brackets is the token count in that combination of classes.

the D (1.4%) class, indicating general gamer slang is used for communication more than terms specific to the game being played. Overall, the large portion of these game-specific classes enables the dataset to be more sophisticated in detecting toxicity in games.

**Dual Annotation Distribution** To understand the effect of dual annotation on the toxicity context, we look at the distribution of the slot labels within each intent class. As seen in Figure 2b, the I intent class shows the highest proportion of the S slot class among non-O classes. Similarly, Figure 2c shows a relatively high proportion of the S slot class in the A intent class. This suggests that the combination of S slot and intent classes provides useful information because slang performs the function of carrying game-specific context. As a result, we focus more on T and S slot classes joined with other intent classes in order to investigate toxicity natures in games carried out from dual annotation.



(a) Class “Explicit, Implicit” (b) Class “Toxicity, Slang”

Figure 3: In-game chat histogram for intent (E,I) and slot (T,S) classes. Match progress is bucketed position within a match whose duration is normalised in  $[0,1]$ , with  $<0$  indicating pregame chat. Merged number of utterance/token is the count of all utterances/tokens from all matches in that match progress bin.

**Keywords in Dual Annotation** Table 4 shows the top 5 keywords by frequency from the T and S slot classes, for each intent class. In the combinations of S class, we observe the prominent position of “ez” in the E and I intent class. This indicates dual annotation captures toxicity from the slang largely used in games. In addition, “gg” appears in all combinations because it may have some toxicity attached via sarcasm. As dual annotation uses the conversational history, it is able to classify the same utterance in different intents.

**Toxicity Analysis Over Time** We further analyse in-game chat over time associated with the E, I intent classes and the T, S slot classes. As shown in Figure 3a, more toxic utterances appear pre-game and post-game rather than during the games. In pre-game, sometimes players are upset due to the selected hero characters if their desired hero is taken by others, or are stressed by planning game strategies in a limited time. Toxic utterance frequency gradually rises towards the end, and peaks in post-game due to chat for post-victory celebration and recrimination. Interestingly, Figure 3b displays a similar pattern. Particularly, tokens for the S slot class increase sharply to the end, indicating significant amounts of slang are used to celebrate wins or humiliate defeated opponents.

**Comparison with Other Datasets** In Figure 4, we compare our dataset with other toxicity detection datasets using the metric of relative frequency of toxic utterances of each length. The datasets we compare with are 1) **Waseem** (Waseem and Hovy, 2016) which consists of 16.2k tweets binary classified as racism/sexism or other, 2) **FoxNews** (Gao and Huang, 2017) which is 1.5k sentences from Fox News discussion threads classified as hateful/non-

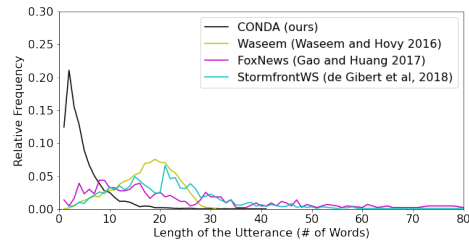


Figure 4: Distribution of toxic utterance length across similar datasets.

hateful, and 3) **StormfrontWS** (de Gibert et al., 2018) which is 10.7k conversation sentences from white supremacist website Stormfront classified as hate speech/non-hate speech. For this analysis, we merge classes into toxic/non-toxic as required. As an example of the CONDA dataset, we combine E and I intent classes into a toxic class, and A and O into a non-toxic class.

The distribution in CONDA is different to the other datasets in that the toxic utterances are shorter. This is due to the terseness of in-game chat during playing, with longer utterances occurring in pre-game and post-game discussion. Waseem has a particular distribution due to the character limit in Twitter (140 characters at the time). FoxNews and StormfrontWS are forums which foster the use of longer sentences.

## 5 Baseline Experiment

To explore the toxicity detection from an NLU perspective, we selected five baseline NLU models and compared their detection performance over our proposed dataset.

### 5.1 Data Preparation

We split the data into train/validation/test sets in the proportions of 0.6/0.2/0.2, or in samples 26,921/8,974/8,974. The data passed to the models is the tokenised utterances with punctuation removed, and for training the slot and intent labels.

### 5.2 Baseline NLU Models

The five NLU models are as follows:

- **RNN-NLU** (Liu and Lane, 2016) is an attention-based bi-directional recurrent neural network model that jointly predicts the current slot and the intent at each time step using shared hidden states and attention.
- **Slot-gated** (Goo et al., 2018) is an attention-based BiLSTM model which builds on sepa-

Model	Metrics												JSA
	UCA	U-F1(E)	U-F1(I)	U-F1(A)	U-F1(O)	T-F1	T-F1(T)	T-F1(S)	T-F1(C)	T-F1(D)	T-F1(P)	T-F1(O)	
RNN-NLU (Liu and Lane, 2016)	0.905	0.813	0.720	0.783	0.944	0.970	0.931	0.981	0.930	0.718	0.991	0.987	0.854
Slot-gated (Goo et al., 2018)	0.894	0.806	0.694	0.773	0.938	<b>0.991</b>	<b>0.978</b>	<b>0.992</b>	<b>0.982</b>	<b>0.952</b>	0.997	<b>0.994</b>	0.875
Inter-BiLSTM (Wang et al., 2018)	0.869	0.719	0.590	0.728	0.923	0.865	0.871	0.889	0.869	0.788	0.942	0.924	0.711
Capsule NN (Zhang et al., 2019a)	0.876	0.735	0.706	0.643	0.926	<b>0.991</b>	0.975	0.991	<b>0.982</b>	0.949	0.997	<b>0.994</b>	0.855
Joint BERT (Castellucci et al., 2019)	<b>0.921</b>	<b>0.872</b>	<b>0.768</b>	<b>0.800</b>	<b>0.954</b>	0.989	0.972	<b>0.992</b>	0.979	0.914	<b>0.998</b>	0.993	<b>0.895</b>

Table 5: Joint intent classification and slot labeling performance on CONDA for the five NLU baseline models. It is measured in the four multi-level metrics including: UCA (Utterance Classification Accuracy); the break-down U-F1 for each intent class - E (Explicit), I (Implicit), A (Action), O (Other); the overall T-F1 and breakdown for each slot class - T (Toxicity), S (game Slang), C (Character), D (Dota-specific), P (Pronoun), O (Other); and JSA (Joint Semantic Accuracy).

rate attended context for slot filling and intent classification while explicitly feeding the intent context into the process of slot filling via a gating mechanism.

- **Inter-BiLSTM** (Wang et al., 2018) combines two inter-connected BiLSTMs performing slot filling and intent classification respectively. The information flow between the two tasks occurs by passing the hidden states at each time step from each side to the other to support the decoding process.
- **Capsule NN** (Zhang et al., 2019a) is a capsule-based neural network that explicitly captures the semantic hierarchical relationship among words, slots and intents via a dynamic routing-by-agreement schema.
- **Joint BERT** directly utilizes the merit of pre-trained BERT (Devlin et al., 2019) and non-recursively conducts the joint prediction over the [CLS] token embedding for intent and the sequence of token embeddings for slots.

### 5.3 Evaluation Metrics

We propose to use the following four metrics for conducting a multi-aspect evaluation. The first two follow the existing traditional abusive language detection research for utterance level detection evaluation while the others are the metrics used for slot-level prediction evaluation and the joint task in NLU models.

- (1) **UCA: Utterance Classification Accuracy** measures the sentence-level classification performance based on the ratio of the number of correctly predicted utterance to the total number of utterances.
- (2) **U-F1: Utterance F1** score calculates the F1 score for each utterance class.

- (3) **T-F1: Token F1** score focuses on the prediction performance for slot tokens and calculates an F1 for each class and the token-based micro-averaged F1 score over all classes excluding label O.
- (4) **JSA: Joint Semantic Accuracy** measures the overall prediction performance over the semantic hierarchy. An utterance is deemed correctly analysed only if both utterance-level and all the token-level labels including Os are correctly predicted.

### 5.4 Implementation details

Links to the source code are given in Appendix C. For Joint BERT, Slot Gated SLU and Capsule NN, we set the number of epochs as 2, 8 and 60, respectively. For the RNN-NLU model, the global step is 1,200 and bidirectional RNN is used with the attention mechanism. For other hyper-parameters, the configuration for the best model in the official GitHub implementation of the baseline models is used. All the experiments are conducted on 16GB Tesla V100-SXM2 GPU with CUDA 10.1.

### 5.5 Baseline results

The experiment result is provided in Table 5. In columns 2 to 6, the metrics associated with utterance labels are shown. The UCA ranges between 0.87 (Inter-BiLSTM) and 0.92 (Joint BERT), and the U-F1 also illustrates a variance of results for each intent class. We observe that class O always achieves the highest F1 score due to its dominance in numbers throughout the dataset. Comparatively, class I presents relatively low F1 scores due to its subtle nature and reliance on understanding context. The variance in U-F1(I) implies potential improvement in implicit toxicity detection.

Columns 7 to 13 present the metrics related with token labels and show much higher overall performance than utterance labels. This indicates the lexicon-based slot automation gives underlying patterns the model can learn easily. Even so, slot class D always has a lower T-F1 score than other slot classes, indicating game-specific tokens in class D have flexible and variant forms, which increases the difficulty of detection. In the last column, the JSA which jointly handles utterance and token labels is shown. Due to the limitation on utterance level intent classification, it presents comparably low JSA scores, indicating a challenge for improvement.

Amongst the models, the non-recursive Joint BERT model performs the best due to the rich linguistic information learned in pre-training. Joint BERT has an implicit influence between the intent sub-task and the slot sub-task based on a joint loss, whereas the recursive models Slot-gated and Capsule NN have explicit influence flowing from intent to slot, leading to similar slot performance. These explicit lines of influence from one task to the other have shown to be successful in NLU and could be explored further in the toxicity detection task.

## 6 Transfer Experiment

We compared our dataset with the toxicity detection datasets introduced in Section 4 in terms of transfer performance over utterance-level binary prediction as toxic or non-toxic. That is, training on one dataset and testing on the others. For simplicity, we solely use the intent classification circuit of the Joint BERT as the prediction model.

### 6.1 Data Preparation

We combine classes into toxic/non-toxic as explained in Section 4. For each dataset, we split into train/test sets in the ratio of 0.9/0.1. The statistics for each dataset are shown in Table 6.

Dataset	Train / Test
Waseem	14,581 / 1,621
StormWS	9,849 / 1,621
FoxNews	1,373 / 1,095
CONDA (ours)	40,382 / 4,487

Table 6: Dataset sample counts for transfer experiment.

### 6.2 Transfer results

The transfer performance measured in UCA is given in Table 7. Firstly, we look at the test performance of each dataset trained on the other datasets, that is to compare results in each column. It can be

seen from column 4 that CONDA’s transferred performance is generally good when trained on each of the other three datasets, ranging from 0.81 to 0.83. This implies that CONDA covers the nature of toxicity that can be generalized from the other toxicity datasets which emphasize hatefulness, sexism and racism.

Changing the perspective to the rows, we compare the test performance for each dataset on a model on another dataset. StormfrontWS and CONDA perform well on Waseem, picking up the racism components there. However Waseem does not perform well when trained on either of those, suggesting their specific hate speech and game speech respectively is too focused. FoxNews training transfers the weakest results indicating its general news nature is too broad.

StormfrontWS performs well when trained on CONDA due to shared toxicity characteristics, but the performance of Waseem and FoxNews when tested on a model trained on CONDA is relatively low. We propose that this is due to two aspects of our dataset previously discussed: the specific game related nature of our language, and the shorter utterances in our set compared to the others.

		Testing			
		Waseem	StormWS	FoxNews	CONDA
Training	Waseem	-	0.8845	0.7287	0.8307
	StormWS	0.7118	-	0.7379	0.8329
	FoxNews	0.6931	0.8241	-	0.8056
	CONDA	0.6955	0.8748	0.6690	-

Table 7: Transfer Learning Result, UCA (Utterance Classification Accuracy).

## 7 Conclusion and Future Work

In this paper, we propose CONDA, a new dataset with dual-level (token and utterance) annotation for understanding in-game chat and to detect toxicity. Compared to previous studies, we draw on the NLU perspective and use the joint token-utterance aspect for detection of toxicity. Accordingly, we formalise a multi-level evaluation system. Through experiments with joint slot and intent NLU models, we show the promising potential of such models for toxicity detection utilizing the dual-level annotation. We also compare our dataset with other benchmark toxicity datasets in the literature through a transfer experiment. In future work, the automated token labelling can be manually adjusted and the size of the dataset can be expanded.



## 8 Ethics/Broader Impact Statement

The study follows the ethical policy set out in the ACL code of Ethics<sup>5</sup> and addresses the ethical impact of presenting a new dataset. In addition, it is approved by our Institutional Review Board (project number : 2019/741). As described in the data collection section, our annotated dataset, CONDA, is based on the *Dota 2* game chat where it can be accessed on Kaggle website (See Section 3.1).

For our automated slot labelling, we generated the game toxicity lexicon by taking the supplemental materials released by Märtens et al. (2015) and ElSherief et al. (2018) and the list of words banned by Google<sup>6</sup>. We then added variants or new toxic words found in the utterances extracted from Kaggle. For intent labelling, all volunteer annotators were recruited from academia and research students. They were informed about toxic behavior in online games before handling the data. Our instructions allowed them to feel free to leave if they were uncomfortable with the content. Due to privacy considerations, we group them by online game experiences and do not take into account annotators' demographic information.

The CONDA dataset is intended for toxicity detection in online games by providing both slot and intent labels. With respect to the potential risks, we note that the subjectivity of human annotation would impact on the quality of the dataset. In order to improve the quality of our dataset, we compared the inter-rater agreements between a gamers' group and a non-gamers' group, and then final annotation of the whole dataset was performed by gamers only.

## References

- Sonam Adinolf and Selen Turkay. 2018. Toxic behaviors in esports games: player perceptions and coping strategies. In *Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts*, pages 365–372.
- ADL. 2019. [Free to play: Hate, harassment and positive social experiences in online games.](#)
- Jeremy Blackburn and Haewoon Kwak. 2014. STFU NOOB! Predicting crowdsourced decisions on toxic behavior in online games. In *Proceedings of the 23rd international conference on World wide web*, pages 877–888.
- <sup>5</sup><https://www.aclweb.org/portal/content/acl-code-ethics>
- <sup>6</sup><https://github.com/RobertJGabriel/Google-profanity-words>
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.
- Giuseppe Castellucci, Valentina Bellomaria, Andrea Favalli, and Raniero Romagnoli. 2019. Multilingual intent detection and slot filling in a joint bert-based model. *arXiv preprint arXiv:1907.02884*.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 11.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*.
- Mai ElSherief, Shirin Nilizadeh, Dana Nguyen, Giovanni Vigna, and Elizabeth Belding. 2018. Peer to peer hate: Hate speech instigators and their targets. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D Manning. 2017. Key-value retrieval networks for task-oriented dialogue. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 37–49.
- Antigoni Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large scale crowdsourcing and characterization of twitter abusive behavior. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12.
- Lei Gao and Ruihong Huang. 2017. Detecting online hate speech using context aware models. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP*, pages 260–266.
- Ona de Gibert, Naiara Perez, Aitor García-Pablos, and Montse Cuadros. 2018. Hate speech dataset from a white supremacy forum. In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, pages 11–20.
- Chih-Wen Goo, Guang Gao, Yun-Kai Hsu, Chih-Li Huo, Tsung-Chieh Chen, Keng-Wei Hsu, and Yun-Nung Chen. 2018. Slot-gated modeling for joint slot filling and intent prediction. In *Proceedings of the 2018 Conference of the NAACL-HLT*, pages 753–757.

- Arshit Gupta, Peng Zhang, Garima Lalwani, and Mona Diab. 2019. Casa-nlu: Context-aware self-attentive natural language understanding for task-oriented chatbots. In *Proceedings of the (EMNLP-IJCNLP)*, pages 1285–1290.
- Sonal Gupta, Rushin Shah, Mrinal Mohit, Anuj Kumar, and Mike Lewis. 2018. Semantic parsing for task oriented dialog using hierarchical representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2787–2792.
- Haewoon Kwak, Jeremy Blackburn, and Seungyeop Han. 2015. Exploring cyberbullying and other toxic behavior in team competition online games. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 3739–3748.
- Bing Liu and Ian Lane. 2016. Attention-based recurrent neural network models for joint intent detection and slot filling. *Interspeech 2016*, pages 685–689.
- Marcus Märtens, Siqi Shen, Alexandru Iosup, and Fernando Kuipers. 2015. Toxicity detection in multiplayer online games. In *2015 International Workshop on Network and Systems Support for Games (NetGames)*, pages 1–6. IEEE.
- Joaquim Alvino de Mesquita Neto and Karin Becker. 2018. Relating conversational topics and toxic behavior effects in a moba game. *Entertainment computing*, 26:10–29.
- Shane Murnion, William J Buchanan, Adrian Smales, and Gordon Russell. 2018. Machine learning and semantic analysis of in-game chat for cyberbullying. *Computers & Security*, 76:197–213.
- Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive language detection in online user content. In *Proceedings of the 25th international conference on world wide web*, pages 145–153.
- John Pavlopoulos, Jeffrey Sorensen, Lucas Dixon, Nithum Thain, and Ion Androutsopoulos. 2020. Toxicity detection: Does context really matter? In *Proceedings of the 58th Annual Meeting of the ACL*, pages 4296–4305. Association for Computational Linguistics.
- Jing Qian, Anna Bethke, Yinyin Liu, Elizabeth M. Belding, and William Yang Wang. 2019. A benchmark dataset for learning to intervene in online hate speech. In *Proceedings of the 2019 Conference on EMNLP-IJCNLP*, pages 4754–4763. Association for Computational Linguistics.
- Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. 2019. Cross-lingual transfer learning for multilingual task oriented dialog. In *Proceedings of the 2019 Conference of the NAACL-HLT*, pages 3795–3805. Association for Computational Linguistics.
- Wessel Stoop, Florian Kunneman, Antal van den Bosch, and Ben Miller. 2019. Detecting harassment in real-time as conversations develop. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 19–24.
- Kunze Wang, Dong Lu, Caren Han, Siqun Long, and Josiah Poon. 2020. Detect all abuse! toward universal abusive language detection models. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6366–6376.
- Yu Wang, Yilin Shen, and Hongxia Jin. 2018. A bi-model based rnn semantic frame parsing model for intent detection and slot filling. In *NAACL-HLT (2)*.
- Yufan Wang, Tingting He, Rui Fan, Wenji Zhou, and Xinhui Tu. 2019. Effective utilization of external knowledge and history context in multi-turn spoken language understanding model. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 960–967. IEEE.
- 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*, pages 88–93.
- Chenwei Zhang, Yaliang Li, Nan Du, Wei Fan, and S Yu Philip. 2019a. Joint slot filling and intent detection via capsule neural networks. In *Proceedings of the 57th Annual Meeting of the ACL*, pages 5259–5267.
- Chenwei Zhang, Yaliang Li, Nan Du, Wei Fan, and Philip S. Yu. 2019b. Joint slot filling and intent detection via capsule neural networks. In *Proceedings of the 57th Conference of the ACL*, pages 5259–5267. Association for Computational Linguistics.
- Justine Zhang, Jonathan P. Chang, Cristian Danescu-Niculescu-Mizil, Lucas Dixon, Yiqing Hua, Dario Taraborelli, and Nithum Thain. 2018. Conversations gone awry: Detecting early signs of conversational failure. In *Proceedings of the 56th Annual Meeting of the ACL*, pages 1350–1361. Association for Computational Linguistics.
- Xiaodong Zhang and Houfeng Wang. 2016. A joint model of intent determination and slot filling for spoken language understanding. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI*, pages 2993–2999. IJCAI/AAAI Press.

