

# Generating Hard-Negative Out-of-Scope Data with ChatGPT for Intent Classification

Zhijian Li, Stefan Larson, Kevin Leach

Vanderbilt University, Nashville, TN, USA

{zhijian.li, stefan.larson, kevin.leach} @ vanderbilt.edu

## Abstract

Intent classifiers must be able to distinguish when a user's utterance does not belong to any supported intent to avoid producing incorrect and unrelated system responses. Although out-of-scope (OOS) detection for intent classifiers has been studied, previous work has not yet studied changes in classifier performance against hard-negative out-of-scope utterances (i.e., inputs that share common features with in-scope data, but are actually out-of-scope). We present an automated technique to generate hard-negative OOS data using ChatGPT. We use our technique to build five new hard-negative OOS datasets, and evaluate each against three benchmark intent classifiers. We show that classifiers struggle to correctly identify hard-negative OOS utterances more than general OOS utterances. Finally, we show that incorporating hard-negative OOS data for training improves model robustness when detecting hard-negative OOS data and general OOS data. Our technique, datasets, and evaluation address an important void in the field, offering a straightforward and inexpensive way to collect hard-negative OOS data and improve intent classifiers' robustness.

**Keywords:** intent classification, out-of-scope, hard-negative data, data generation

## 1. Introduction

Task-oriented dialog systems rely on robust intent classification models to produce appropriate responses based on the user utterances. During deployment, the intent classifiers need to not only accurately classify the user utterances, but also must identify if user utterances do not belong to any supported intents. Brittle intent classifiers that fail to reliably distinguish OOS (out-of-scope) utterances from the INS (in-scope) utterances ultimately lead to poor user experiences, wasted time and resources, and potential safety concerns. Thus, it is imperative to develop techniques to ensure robustness against OOS utterances.

When developing intent classifiers, developers typically begin by collecting and labelling a large amount of INS training data often acquired through crowd-sourcing (Kang et al., 2018; Larson and Leach, 2022b). Collecting OOS datasets to improve the classifier's OOS detection capability is not a common practice (Larson et al., 2019b). Currently, there is a dearth of large public OOS datasets, and the difficulty posed to intent classifiers has not been rigorously examined. General OOS datasets typically contain mostly samples that exhibit minimal similarity with the INS samples. While many models can distinguish such OOS samples and the INS samples during testing, using such OOS samples to evaluate models' OOS detection capabilities can produce misleading results. During deployment, the model may encounter OOS utterances that closely resemble the INS utterances but convey entirely different meanings. Due to the high similarity that those *hard-negative out-of-scope* utterances share with the

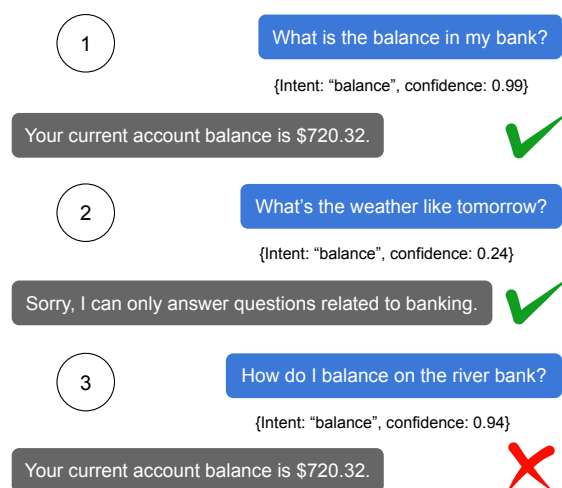


Figure 1: Example exchanges between a user (blue, right side) and a task-driven dialog system for personal finance (grey, left side). The system correctly identifies the user's utterance as in-scope in ①, and correctly identifies the user's utterance as out-of-scope and gives a valid response in ②. In ③, the system incorrectly identifies the hard-negative OOS user utterance as in-scope and provides an incorrect response.

INS data, the models are more susceptible to misclassifying them into a supported intent with high confidence. Therefore, collecting hard-negative OOS data is pivotal to ensure that the intent classifiers can reliably distinguish all OOS utterances, regardless of their resemblance to the INS samples.

Figure 1 shows example utterances and responses of a task-oriented dialog system driven

by an intent classifier trained only on utterances for personal finance. In the first user-system exchange, the system correctly categorizes the user utterance as an in-scope `balance` utterance with appropriately high confidence. In the second exchange, the user provides a general OOS utterance, and the system successfully identifies that the user utterance cannot be understood because the confidence score for all the intents are low. In the third exchange, the user presents a hard-negative OOS utterance. Since the utterance includes the words “balance” and “bank”, and appears similar to the in-scope training data for the `balance` intent, the system incorrectly classifies the utterance as in-scope for `balance` and produces an incorrect response.

Obtaining OOS data that sufficiently challenges the intent classifiers is difficult. Often, this is done through crowd-sourcing with platforms like Amazon Mechanical Turk (e.g., [Larson et al. \(2019b\)](#)). However, this approach is costly and time-consuming, and it requires careful verification to make sure that the samples are indeed out-of-scope and challenging. Furthermore, data collection via crowd-sourcing introduces quality control problems as the the collected data are often error prone ([Larson et al., 2020a](#); [Shah and Zhou, 2016](#)). In datasets containing a large number of intents, verifying each utterance to be irrelevant from all the intents poses significant difficulty for human crowd-sourcing workers. For example, the Clinc-150 dataset ([Larson et al., 2019b](#)) encompasses 150 intents; relying on the crowd-workers to verify all sentences to be OOS may be challenging and could lead to erroneous results.

A cost-effective alternative is to generate hard-negative OOS data using Large Language Models such as ChatGPT. [Larson et al. \(2019b\)](#) paid crowd workers \$0.20 US to produce three paraphrases of an utterance, and current best practices involve payment that extrapolates to a fair wage ([Kummerfeld, 2021](#)). In comparison, the GPT-3.5 turbo API costs \$0.0015 per 1K tokens in the prompt and \$0.003 per 1K tokens for the output, a substantial potential savings.

In this paper, we aim to investigate the following research questions:

1. Can ChatGPT generate OOS utterances that do not fall into any of the system-supported INS intents?
2. Do the generated hard-negative OOS utterances lead to high-confidence predictions from intent classifiers?
3. Can the generated hard-negative OOS utterances be used in training to improve the intent classifiers’ OOS detection ability and decrease the models’ confidence on the sup-

ported intents when encountering OOS utterances?

To answer these questions, we select five large public datasets and introduce a method to generate 3,732 hard-negative OOS utterances — that is, utterances intended to closely resemble the in-scope data for a given intent.

Our method works by analyzing important words that are likely to have the biggest influence on the intent classifiers’ predictions for each intent from the INS training data. Then, the method prompts ChatGPT to generate OOS utterances that include specific important words for each intent. Since the generated OOS utterances contain the important words from supported intents, they are more likely to confuse the model and produce high confidence. In doing so, we can increase the rigor and challenge of a given dataset by partially automating the creation of hard negative inputs.

We evaluate the performance of three benchmark transformer models for intent classification on our generated datasets. The hard-negative OOS utterances produced using our method consistently resulted in high confidence (but incorrect) predictions across all five datasets. Notably, intent classifiers struggled to distinguish the hard-negative OOS utterances from the INS utterances. In particular for Clinc-150 and Banking77, model confidence scores are substantially higher for our generated hard-negative OOS datasets compared to the general OOS dataset. We see various improvements in model performance when incorporating hard-negative OOS data in training. For instance, for Banking77 evaluated with BERT, the AUROC for hard-negative OOS and INS improved to 0.996 and the AUROC for general OOS and INS improved to 0.989 when the model is trained on both INS and hard-negative OOS data. Our study shows that more attention must be given to curating hard-negative OOS datasets to enhance model robustness in deployment.

## 2. Related Work

In this section, we discuss prior work related to data collection for intent classification tasks and ChatGPT’s data generation capabilities.

### 2.1. Hard-Negative Data

Several studies have highlighted the significance of using hard-negative samples during training to improve model robustness. For instance, [Zhan et al. \(2021\)](#) and [Nguyen et al. \(2023\)](#) used hard-negative examples during training to aid retrieval models to better discriminate between relevant and irrelevant documents. Hard-negatives can

also facilitate contrastive learning for image classification (e.g., Kalantidis et al. (2020)) and image retrieval (e.g., Hughes et al. (2018); Melekhov et al. (2016)). However, the use of hard-negatives for intent classification remains understudied. In this paper, we use our generated hard-negative OOS datasets to improve model robustness for intent classification.

## 2.2. Data Collection for Intent Classification

**General Data Collection.** Prior strategies for data collection for constructing training and evaluation data for intent classifiers include the use of crowd-sourcing to generate queries by either (1) paraphrasing input query prompts or (2) responding to scenarios with queries. This prior work includes Coucke et al. (2018); Gupta et al. (2018); Kang et al. (2018); Larson et al. (2019b); Liu et al. (2019a). Recently, prior work has investigated using large language models (LLMs) to generate this type of training data. This work includes Rosenbaum et al. (2022), who used AlexaTM, and Sahu et al. (2022), who used GPT-3, and Cegin et al. (2023), who used ChatGPT to generate paraphrases.

**Out-of-Scope and Challenging Data Collection.** Most prior work on data collection for intent classification does not consider the production of OOS queries for evaluating a model’s ability to distinguish between in- and out-of-scope inputs. Exceptions to this include Larson et al. (2019b), who’s Clinc-150 dataset includes out-of-scope utterances that were generated via crowd-sourcing. Other work in this space includes Larson and Leach (2022a), whose out-of-scope data is constructed by sampling from other datasets, and Zhang et al. (2022), who constructed “in-domain, out-of-scope” splits of Clinc-150 and Banking77 where training data includes a set of intents  $I_{in-scope}$  from the original dataset  $I$ , but evaluation data includes a set of intents  $I_{out-of-scope}$  that are in the same domain<sup>1</sup> as  $I_{in-scope}$  but  $I_{in-scope} \cap I_{out-of-scope} = \emptyset$ .

Similarly, Khosla and Gangadharaiah (2022) created challenging datasets for testing model robustness to “covariate shifted” data. In that work, covariate shifted data means data that was generated from a different distribution with respect to some other original data distribution, but where both data distributions generate data for the same intent category. Considering the `weather` intent from Clinc-150, covariate shifted data includes data that was originally generated for the Snips or HWU64 datasets (i.e., the `get_weather` in-

tent from Snips, and HWU64’s `weather_query` intent).

Other work that focuses on directly generating new challenge data includes techniques from Larson et al. (2019a), which prompted crowd workers to paraphrase only unique queries from a dataset, and Larson et al. (2020b), who prompted crowd workers to paraphrase seed phrases but constrained the crowd workers from using certain keywords that were found to be indicative of certain intents. The techniques from both of these works can be seen as ways to generate covariate shifted in-scope data.

## 2.3. Adversarial Examples

Our work on generating hard-negative data has similarities to prior work on generating adversarial examples. Motivated by adversarial example generation work in computer vision that applies “imperceptible” perturbations to images to provoke incorrect model classification (e.g., Goodfellow et al. (2015)), prior work in natural language processing on adversarial example generation has revolved around perturbing texts using character-level alterations (e.g., Ebrahimi et al. (2018); Gao et al. (2018)) and word synonym and phrase replacement (e.g., Alzantot et al. (2018); Garg and Ramakrishnan (2020); Ren et al. (2019)). Similar work in dialog data has been done by Liu et al. (2021); Peng et al. (2021); Sengupta et al. (2021), where utterances are modified (e.g., by introducing typos, word synonyms, ASR errors, etc.) in order to test the robustness of models. The method we present in this paper can be seen as way to generate adversarial examples to test the robustness of intent classification models, but differs in that we do not apply perturbations to existing samples.

## 2.4. Data Collection and Annotation with ChatGPT

In a recent study, Cegin et al. (2023) showed that large language models such as ChatGPT can generate more lexically and syntactically diverse INS data by paraphrasing existing corpora. Cegin et al. (2023) also showed that ChatGPT can follow prompted restrictions and used ChatGPT in lieu of crowd-sourcing to generate “taboo” paraphrases in the manner of Larson et al. (2020b), where certain words are avoided in paraphrases in order to promote diversity. However, Cegin et al. (2023) also observed several issues with open-sourced models such as Falcon-40b: duplicated outputs, erroneous outputs, and lack of instruction following. Another study showed ChatGPT’s performance varies for sentiment analysis on tweets depending on the topics (Zhu et al., 2023). For tasks such as classifying the political affiliation of

<sup>1</sup> Here, a domain is a semantically meaningful group of intents like “banking”, “travel”, etc.

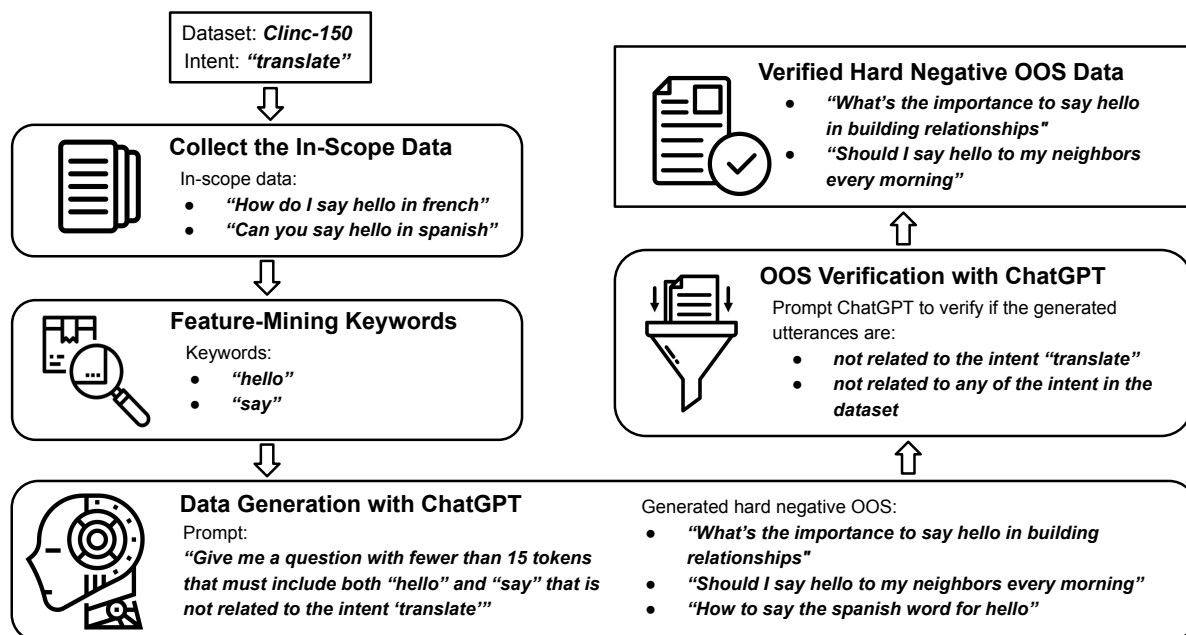


Figure 2: An overview of the hard-negative OOS generation process, including examples. The third generated utterance is filtered out during the two-step OOS verification.

Twitter users, ChatGPT outperforms human crowd-source workers (Törnberg, 2023). ChatGPT’s performance for text-annotation tasks has also been found to exceed that of crowd workers (e.g., Gilardi et al. (2023); He et al. (2024)). Sahu et al. (2022) prompted GPT-3 to generate labeled training data, and the generated data significantly improves the intent classifiers when the intents are distinct. Using ChatGPT to rephrase sentences, models trained with AugGPT outperform state-of-the-art text data augmentation methods to generate data for scarce intents in a few-shot learning setting (Dai et al., 2023).

### 3. Methods

We introduce an automated method for generating hard-negative OOS utterances using ChatGPT’s API. In this study, we generate 3,732 hard-negative OOS queries for five benchmark datasets. Our objective is to generate hard-negative OOS utterances by producing utterances that are likely to contain words that heavily influence the intent classifiers’ predictions.

The hard-negative OOS datasets are generated following these steps, which we discuss throughout this section. Example utterances produced with our approach are shown in Figure 2.

1. Use feature-mining to select  $n$  keywords by analyzing the most frequently appearing words for each intent for every selected dataset.

2. Select a combination of  $m$  keywords  $k_1 \dots k_m$  from the  $n$  keywords for an intent  $i$ .
3. Show ChatGPT the name of the intent  $i$  and five in-scope utterances from  $i$ .
4. Prompt ChatGPT to generate  $x$  questions (i.e., utterances) that must contain  $k_1 \dots k_m$ , and is not related to the intent  $i$ .
5. Prompt ChatGPT to verify that each of the  $x$  generated questions is not related to the intent  $i$ .
6. Prompt ChatGPT to verify that each of the  $x$  generated questions is not related to any of the intents in the entire dataset.

In this study, we select  $n$  to be 5,  $m$  to be 2, and  $x$  to be 4. When the prompts include only one keyword, we noticed that the generated data shares less similarity to the INS data compared to when  $m = 2$ . When prompts include three or more keywords, ChatGPT frequently struggles to include all the keywords or produce an OOS utterance. Larger  $n$  and  $x$  can be selected to generate more hard-negative OOS data.

#### 3.1. Feature-Mining Keywords

To ensure that the generated OOS queries resemble INS queries for each intent, we need to identify important keywords that are likely to influence the intent classifiers’ prediction. Among our five selected datasets, we collect the most frequently occurring words in the INS training samples from each intent. We lemmatize all words using NLTK’s



WordNetLemmatizer (Bird and Loper, 2004) to prevent getting multiple versions of the same word. In addition, we discard the stop words and tokens that contain less than three characters. We have also explored alternative methods to identify keywords. For instance, using the Python ELI5 package,<sup>2</sup> we determined the words with the highest weights after training an SVM classification model on each of the datasets, as was done in Larson et al. (2020b). For larger datasets such as Clinc-150 and HWU64, training multi-class SVM on a substantial number of intents is time consuming. We utilized the LinearSVC from Scikit-Learn (Pedregosa et al., 2011) which is implemented as One-vs-All, resulting in a model requiring  $n$  classifiers for  $n$  intents. After removing stopwords and lemmatizing the tokens, we are not able to obtain at least five keywords for many intents. For example 65 of 150 intents from Clinc-150 produced less than 5 keywords using the ELI5 method. Therefore, we use the original token frequency-based approach for keywords collection.

### 3.2. Data Generation with ChatGPT

Recall that we prompt ChatGPT to produce hard-negative OOS samples. We use the GPT 3.5-turbo model through the chat completion API. (All experiments using ChatGPT were done June-August 2023.) For interaction with the API, we make use of three distinct “roles”: (1) The “system” role, which allows the developers to guide ChatGPT’s behavior throughout the conversation. (2) The “assistance” role, which grants ChatGPT’s API the access to previous conversations, enabling ChatGPT to recall previously generated hard-negative OOS utterances to avoid generating duplicates. ChatGPT can also retain the intent information, removing the need for us to repeatedly inform ChatGPT the intent in every prompt. (3) The “user” role, which lets the developers to prompt questions for ChatGPT’s API to answer.

When generating hard-negative OOS utterances, we first set the role to “system” and guide ChatGPT to answer with only the hard-negative sentence. This approach prevents the API from padding the responses with unnecessary tokens like “Sure, I’d be happy to help.” Next, we set the role to “user” and show ChatGPT the name of each intent and five INS samples for that intent. This enables ChatGPT to understand the semantics of the intent, especially when the intent name alone does not provide sufficient context. Then, we use the “assistance” role to record the dialog between the developer and ChatGPT.

Subsequently, we prompt ChatGPT to generate an utterance that must be unrelated to the intent

(i.e., OOS) and must contain a combination of two keywords collected during feature-mining.<sup>3</sup> This process allows us to guide ChatGPT to produce an utterance that contains commonly used words associated with a given intent but that is not related to that intent — that is, a hard-negative OOS sample. Next, we discuss how we validate each such generated utterance is in fact OOS.

### 3.3. OOS Verification with ChatGPT

To ensure that the hard-negative OOS data does not contain any INS samples, we use a two-step verification method using ChatGPT. After ChatGPT generates every hard-negative OOS utterance, we immediately prompt ChatGPT to assess whether each utterance belongs to the intent that the utterance should *not* relate to. For example, after prompting ChatGPT to generate a question containing “hello” and “french” that is not related to the `translate` intent, we subsequently prompt ChatGPT to determine whether the generated utterance is related to “translate”. If ChatGPT determines that the utterance is related to “translate”, then it is discarded. In the second step of verification, the remaining utterances are then checked to be OOS with respect to the entire dataset. To implement this step, we provide ChatGPT with the name of all INS intent categories using the “system” role, and prompt ChatGPT to verify if each utterance is OOS with the ‘user’ role.

We note that ChatGPT occasionally mislabels a small portion of the utterances during this two-step verification process. However, since our goal is to generate hard negative OOS data, discarding such inadvertently mislabeled data is not critical. Finally, we manually check the hard-negative OOS datasets collected after the two-step verification stage to ensure the label accuracy. For utterances that do not clearly fall into the INS or OOS categories, we compare the utterance with the INS samples and discuss amongst the research team to conclude the correct label. We discard the utterances if no consensus is reached regarding label opinions.

## 4. Evaluation

We design experiments to assess the level of difficulty that intent classifiers have when discerning between in-scope (INS) and hard-negative out-of-scope (OOS) data. As a baseline, we also measure models’ abilities to discern INS data from “general” (i.e., not hard-negative) OOS data. Additionally, we evaluate the improvements in mod-

<sup>2</sup> <https://github.com/eli5-org/eli5>

<sup>3</sup> Preliminary experiments revealed that ChatGPT will often output utterances with mostly the same tokens if prompted to generate multiple utterances at once.

els' OOS detection abilities after including hard-negative OOS and general OOS data in training. We hypothesize that the hard negative OOS data we generate with our approach will lead to performance decreases because the model will confuse these samples as being INS.

#### 4.1. Data

**In-Scope Data.** We use INS data from **ClinC-150** (Larson et al., 2019b), **Banking77** (Casanueva et al., 2020), **ATIS** (Dahl et al., 1994; Hemphill et al., 1990; Hirschman et al., 1993, 1992), **Snips** (Coucke et al., 2018), and **HWU64** (Liu et al., 2019a). Different from the other four datasets, HWU64 consists of two intents `General_Quirky` and `QA_Factoid` that span a wide spectrum of semantics. We excluded these two “catch-all” intents from HWU64 for training and testing to allow for sufficient samples that are considered OOS.

**Hard-Negative Out-of-Scope Data.** We use our method to generate 3,732 hard-negative out-of-scope samples targeting intents from the five datasets listed above. Section 5.1 discusses the generated data in detail.

**General Out-of-Scope Data.** Instead of hard-negative OOS data, the baseline approach uses “general” OOS data. We use the ClinC-150 companion OOS data as this general OOS data, which consists of 1,000 test utterances. We then filtered out any utterances that belonged to any of the in-scope intents from the other four datasets.

**Data Splits.** We split the INS data into training and testing. Each selected model is trained on the training data and evaluated with the INS testing data, the generated hard-negative OOS data, the general OOS data from ClinC-150. To examine whether hard-negative OOS data can be used during training to improve the models' OOS detection capability, we separated both hard-negative OOS data and general OOS data into 80% training and 20% testing splits and compared the models' confidence for the OOS datasets when we included hard-negative OOS, general OOS, and both OOS corpora in training.

#### 4.2. Intent Classification Models

We use the following intent classifiers in our experiments: **BERT**, a neural network that uses a transformer model to capture the context of each word and fine-tuned on the training data for NLP tasks (Devlin et al., 2019). In this study, we use bert-base-uncased. **RoBERTa**, a variant of BERT

that provides better contextual representation of text (Liu et al., 2019b). **DistilBERT**, a lightweight variant of BERT that uses fewer parameters than BERT and processes texts faster (Sanh et al., 2019). We used the Hugging Face implementations of these models (Wolf et al., 2020).

Common approaches to dealing with OOS inputs involve the use of confidence scores to differentiate between INS and OOS inputs. Normally, a desirable model is one that assigns higher confidence to INS inputs, and lower confidence to OOS inputs. We use two functions to produce the confidence scores when evaluating our hard-negative OOS datasets: (1) **Softmax**, where we use the highest softmax confidence score for each prediction (Hendrycks and Gimpel, 2016). (2) **Energy**, where we compute the energy score (Liu et al., 2020) for each prediction using  $T = 1$ .<sup>4</sup> These two methods are commonly used in prior work on out-of-distribution detection (e.g., Larson et al. (2022)).

#### 4.3. Metrics

We consider several performance metrics to evaluate the quality of our hard-negative OOS datasets, assigning the INS predictions as the positive class and OOS predictions as the negative class: (1) **AUROC**, the Area Under the Receiver Operating Characteristic curve can be interpreted as measuring the overlap between two distributions and is commonly used for benchmarking OOD performances (Fort et al., 2021). An AUROC score of 0.5 means the model is unable to produce confidence scores that distinguish between INS and OOS inputs, and an AUROC closer to 1.0 indicates better OOS detection capability. (2) **AUPR**, the Area Under the Precision and Recall Curve. A higher AUPR indicates a more robust model. (3) **FPR95**, the false positive rate at 95% recall. A lower FPR95 indicates a more robust model. (4) **F1 score with confidence thresholds**, the F1 score for a range of confidence thresholds from 0.5 to 0.95 for softmax confidence scores. Predictions with higher confidence score than the confidence threshold are considered as positive predictions. Higher F1 captures the classifier's ability to distinguish INS and OOS data at different confidence thresholds since most intent classifiers in deployment utilize a confidence threshold to determine whether an utterance is not understood.

## 5. Results

This section discusses the results of our data generation and experiments: Section 5.1 discusses the results of generating hard-negative OOS with

---

<sup>4</sup> We considered alternative values for  $T$  and the results are similar for different  $T$ .

Dataset	Intent	Keywords	Hard-Negative OOS
Clinc-150	find_phone	find, locate	<i>how do i <u>find</u> and <u>locate</u> a lost pet</i>
Clinc-150	change_language	french, english	<i>are <u>french</u> and <u>english</u> official languages in canada</i>
Clinc-150	distance	take, long	<i>does it <u>take long</u> to find a seat on the bus</i>
Banking77	exchange_rate	foreign, know	<i>do you <u>know</u> any foreign exchange officers nearby</i>
Banking77	age_limit	age, children	<i>what are <u>age-appropriate</u> money lessons for <u>children</u></i>
Banking77	card_arrival	track, sent	<i>can i <u>track</u> my <u>sent</u> documents</i>
HWU64	iot_hue_lightchange	color, change	<i>can temperature <u>change</u> the <u>color</u> of a blue flame</i>
HWU64	alarm_set	wake, tomorrow	<i>what are some tips to <u>wake up</u> refreshed <u>tomorrow</u></i>
HWU64	iot_coffee	make, machine	<i>how do i <u>make</u> a homemade washing <u>machine</u></i>
ATIS	flight_time	time, open	<i>what <u>time</u> does the golden gate bridge <u>open</u></i>
ATIS	aircraft	type, used	<i>what <u>type</u> of currency is <u>used</u> in boston</i>
Snips	BookRestaurant	book, reservation	<i>can i <u>book</u> a <u>reservation</u> for the flight</i>
Snips	GetWeather	like, going	<i>what's the vibe <u>like</u> <u>going</u> there</i>

Table 1: Example hard-negative OOS data generated by ChatGPT across different intents and datasets.

Dataset	Total	Step 1	Step 2	Final
Clinc-150	6,000	4,442	2,278	2,266
Banking77	1,440	1,169	742	734
ATIS	640	458	226	220
Snips	280	200	95	90
HWU64	2,720	1,903	438	422
Overall	11,080	8,172	3,779	3,732

Table 2: Results of prompting ChatGPT to generate hard-negative OOS utterances for each dataset. “Step 1” and “Step 2” show how many utterances remained valid after the ChatGPT OOS Verification. The total count of valid hard-negative OOS utterances is displayed in the “Final” column.

our method. Section 5.2 discusses results of experiments determining the effectiveness of hard-negative OOS data vis-à-vis general OOS data. Section 5.3 discusses results of experiments on determining the effectiveness of training with hard-negative OOS data.

## 5.1. Generation Results

In this study, we prompted ChatGPT to generate a total of 11,080 hard-negative OOS utterances from five different datasets. Roughly 74% (8,172 samples) of the generated hard-negative passed the first round of OOS verification. Then, 3,779 (34%) remained to be valid after the second round of OOS verification where each utterance is prompted to ChatGPT to determine whether it is related to any known in-scope intent of the dataset. Finally, when we manually examined the hard-negative OOS datasets after the two-step verification method, we found that 47 (1.2% of the hard-negative OOS after the two-step verification) are INS and mislabelled by ChatGPT, resulting in 3,732 valid hard-negative OOS samples. Examples of generated hard-negative data are shown in

Table 1. Precise counts of verified hard-negative OOS data for each dataset is shown in Table 2.

## 5.2. Classifier Performance

We evaluate our five generated hard-negative OOS datasets with BERT, RoBERTa, and DistilBERT trained on only INS data. Table 3 shows the AUROC, AUPR, and FPR95 of all datasets evaluated with BERT.

For Banking77 and ClinC-150, the softmax and energy confidence scores for the INS predictions are substantially closer to the hard-negative OOS predictions than the general OOS predictions for all three models. For instance, when evaluated with BERT, Figure 3 (a) shows that the distribution of the softmax confidence scores for the hard-negative OOS is much closer to the INS data compared to the general OOS data. Figure 3 (b) demonstrates that the distribution of energy confidence scores for the INS data is also more similar to that of the hard-negative OOS data compared to the general OOS data. Figure 3 (c) displays the F1 score for the INS vs hard-negative OOS data is lower than that for the INS vs. general OOS data at all confidence thresholds, indicating the model is worse at distinguishing hard-negative OOS than general OOS from the INS data. RoBERTa and DistilBERT resulted in distributions of confidence scores comparable to BERT. Figure 4 displays the distribution of softmax confidence score for RoBERTa after training on the in-scope data from ClinC-150.

During the evaluation for ATIS, Snips, and HWU64, all three models predict both hard-negative OOS data and general OOS data with a substantial number of high softmax confidence scores. The AUROC of the softmax and energy confidence scores for our hard-negative OOS with INS are lower than those for general OOS with INS, indicating that our generated hard-negative OOS

Dataset	AUROC		AUPR		FPR95	
	General OOS	HN OOS	General OOS	HN OOS	General OOS	HN OOS
Clinc-150	0.968	0.914	0.982	0.914	0.121	0.326
Banking77	0.964	0.810	0.959	0.826	0.205	0.639
ATIS	0.964	0.953	0.972	0.989	0.179	0.245
Snips	0.956	0.864	0.982	0.993	0.223	0.400
HWU64	0.921	0.917	0.971	0.988	0.392	0.462

Table 3: Performance comparison between OOS data with the INS data of a BERT model evaluated via softmax confidence scores.




Figure 3: Results for Banking77 evaluated with BERT. (a) shows the distribution of softmax confidence scores. (b) shows the distribution of energy confidence scores. (c) shows the F1 score of softmax confidence score for hard-negative OOS and general OOS with in-scope at different confidence thresholds.




Figure 4: Distribution of softmax confidence scores for Clinc-150 evaluated with RoBERTa.

utterances has a larger overlap with the INS utterances compared to the general OOS and INS's overlap in confidence scores. The FPR95 for both softmax and energy confidence for hard-negative OOS with INS are higher than those for general OOS with INS across all models and datasets, highlighting that the confidence scores can more effectively differentiate the INS data from the general OOS data than our hard-negative OOS data.

Our results suggest that the hard-negative OOS utterances generated with our approach are, at

minimum, as challenging as the general OOS dataset and frequently result in high-confidence, incorrect predictions from intent classifiers. Specifically for Clinc-150 and Banking77, the hard-negative OOS utterances are substantially more difficult to differentiate as be OOS in comparison to the general OOS utterances. Therefore, our proposed hard-negative OOS datasets and other hard-negative OOS datasets generated following our approach will challenge intent classifiers and scrutinize classifiers' robustness against such data.

### 5.3. Using Hard-Negative OOS in training

When we train transformer-based intent classifiers only on INS data (Section 5.2), the models predicted high confidence for hard-negative OOS utterances and varying general OOS confidence across all five benchmark datasets. To determine if using hard-negative OOS in *training* can improve model robustness, we compare the model confidence after training BERT with (1) INS and general OOS data, (2) INS and hard-negative OOS data, (3) INS, general OOS, and hard-negative OOS data, using a 80/20 train test split for the datasets that are used in training.

Table 4 presents results for the AUROC of the



Dataset	INS		General OOS		Hard-Negative OOS		Both OOS	
	General OOS	HN OOS	General OOS	HN OOS	General OOS	HN OOS	General OOS	HN OOS
Clinic-150	0.968	0.914	0.989	0.916	0.981	0.988	0.984	0.990
Banking77	0.964	0.810	0.997	0.874	0.989	0.996	0.996	0.992
ATIS	0.964	0.953	0.998	0.974	0.996	0.996	1.000	1.000
Snips	0.956	0.864	0.999	0.919	0.968	0.999	0.998	0.993
HWU64	0.921	0.917	0.965	0.927	0.924	0.985	0.961	0.986

Table 4: The AUROC of the OOS data with the INS data of a BERT model. If the AUROC is closer to 1, then the model is less likely to classify an OOS utterance as INS. We compare how using general OOS, hard-negative OOS, and both OOS datasets in training improves the model’s OOS detection capability.

softmax confidence scores to be INS for hard-negative OOS and OOS data with the INS data for each model across five datasets. Since OOS data are included in training, we consider “oos” to be a new label. We calculate the confidence score for each prediction by taking the highest softmax score for any in-scope intent. This confidence score indicates how confident the classifier models are in predicting that utterance belongs to a known intent and is in-scope.

For Clinic-150, Banking77, and ATIS, when the intent classifiers are trained on OOS data, the confidence scores for general OOS utterances drop substantially compared to models trained only on in-scope data, but the confidence score for hard-negative OOS utterances remain high.

In comparison, when we add hard-negative OOS in the training corpora, the models produce low confidence predictions for both hard-negative OOS and OOS utterances, shown by the higher AUROC in Table 4.

For Snips and HWU64, incorporating general OOS in training results in high confidence predictions for hard-negative OOS, and using hard-negative OOS in training results in high confidence predictions for the general OOS. In this case, incorporating only hard-negative OOS in training is not enough to ensure model robustness. When using both hard-negative OOS and general OOS in training, the model predicts OOS utterances with low confidence to be INS.

This result indicates that models trained solely with INS and general OOS data are still prone to predicting hard-negative OOS data with high confidence. When the generated hard-negative OOS datasets are incorporated in the training data, the models are much less likely to produce high confidence scores for hard-negative OOS utterances and displays varying improvements for detecting general OOS utterances. Although our experiments demonstrate that training with both hard-negative OOS and general OOS greatly reduces confidence on OOS utterances, we note that there is no way to guarantee that every OOS intent can be covered. That is, in deployment, we cannot

guarantee that the distribution of inputs will follow the OOS data generated with our approach. Nonetheless, our approach can be used to help improve model robustness and to improve benchmarking and data quality.

## 6. Conclusion

We present a new approach to generating hard-negative OOS data using ChatGPT. After manually reviewing and verifying generated data, we show our approach can generate data that are OOS data with infrequent mislabels. Our evaluation shows that models trained only on INS data are brittle when tested against hard-negative OOS utterances generated with our approach and often result in overconfident but incorrect predictions. Our results indicate that the hard-negative OOS utterances are more challenging to differentiate from the INS utterances when compared to the general OOS utterances. Furthermore, we show that using hard-negative OOS data in training improves model robustness against hard-negative OOS utterances substantially and general OOS utterances to varying degrees. Models trained on general OOS data still struggle with hard-negative OOS utterances to a noticeable extent across all five datasets. Since collecting hard-negative OOS data with ChatGPT is substantially less costly than traditional crowd-sourcing methods, we hope that our technique and analysis will lead to more robust intent classifiers.

All code used for generating and verifying hard-negative OOS data with ChatGPT, datasets generated and used, results, and additional figures are available at [github.com/frank7li/Generating-Hard-Negative-Out-of-Scope-Data-with-ChatGPT-for-Intent-Classification](https://github.com/frank7li/Generating-Hard-Negative-Out-of-Scope-Data-with-ChatGPT-for-Intent-Classification).

## Acknowledgements

We acknowledge support from Vanderbilt University’s Summer Research and Seeding Success programs.

## References

- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. [Generating natural language adversarial examples](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*.
- Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. [Efficient intent detection with dual sentence encoders](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*.
- Jan Cegin, Jakub Simko, and Peter Brusilovsky. 2023. [ChatGPT to replace crowdsourcing of paraphrases for intent classification: Higher diversity and comparable model robustness](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibault Lavril, et al. 2018. [Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces](#). *arXiv preprint arXiv:1805.10190*.
- Deborah A. Dahl, Madeleine Bates, Michael Brown, William Fisher, Kate Hunicke-Smith, David Pallett, Christine Pao, Alexander Rudnicky, and Elizabeth Shriberg. 1994. [Expanding the scope of the ATIS task: The ATIS-3 corpus](#). In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Haixing Dai, Zhengliang Liu, Wenxiong Liao, Xiaoke Hugna, Yihan Cao, Zihao Wu, Lin Zhao, shaochen Xu, Wei Liu, Ninghao Liu, Sheng Li, Dajiang Zhu, Hongmin Cai, Lichao Sun, Quanzheng Li, Dinggang Shen, Tianming Liu, and Xiang Li. 2023. [AugGPT: Leveraging chatGPT for text data augmentation](#). *arXiv preprint arXiv:2302.13007*.
- 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 (NAACL-HLT)*.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. [HotFlip: White-box adversarial examples for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Stanislav Fort, Jie Ren, and Balaji Lakshminarayanan. 2021. [Exploring the limits of out-of-distribution detection](#). *Advances in Neural Information Processing Systems*, 34:7068–7081.
- Ji Gao, Jack Lanchantin, Mary Lou Soffa, and Yanjun Qi. 2018. [Black-box generation of adversarial text sequences to evade deep learning classifiers](#). In *2018 IEEE Security and Privacy Workshops (SPW)*.
- Siddhant Garg and Goutham Ramakrishnan. 2020. [BAE: BERT-based adversarial examples for text classification](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. [ChatGPT outperforms crowd-workers for text-annotation tasks](#). *Proceedings of the National Academy of Sciences*, 120(30).
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and harnessing adversarial examples](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- 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 (EMNLP)*.
- Zeyu He, Huang Chieh-Yang, Chien-Kuang Cornelia Ding, Shaurya Rohatgi, and Ting-Hao 'Kenneth' Huang. 2024. [If in a crowd-sourced data annotation pipeline, a GPT-4](#). *arXiv preprint arXiv:2402.16795*.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. [The ATIS spoken language systems pilot corpus](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.

- Dan Hendrycks and Kevin Gimpel. 2016. [A baseline for detecting misclassified and out-of-distribution examples in neural networks](#). In *Proceedings of International Conference on Learning Representations (ICLR)*.
- L. Hirschman, M. Bates, D. Dahl, W. Fisher, J. Garofolo, D. Pallett, K. Hunicke-Smith, P. Price, A. Rudnicky, and E. Tzoukermann. 1993. [Multi-site data collection and evaluation in spoken language understanding](#). In *Human Language Technology: Proceedings of a Workshop Held at Plainsboro, New Jersey, March 21-24, 1993*.
- Lynette Hirschman, Madeleine Bates, Deborah Dahl, William Fisher, John Garofolo, and Kate Hunicke-Smith. 1992. [Multi-site data collection for a spoken language corpus](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*.
- Lloyd Haydn Hughes, Michael Schmitt, and Xiao Xiang Zhu. 2018. [Mining hard negative samples for sar-optical image matching using generative adversarial networks](#). *Remote Sensing*, 10(10).
- Yannis Kalantidis, Mert Bulent Sariyildiz, Noe Pion, Philippe Weinzaepfel, and Diane Larlus. 2020. [Hard negative mixing for contrastive learning](#). *Advances in Neural Information Processing Systems*, 33:21798–21809.
- Yiping Kang, Yunqi Zhang, Jonathan K. Kummerfeld, Lingjia Tang, and Jason Mars. 2018. [Data collection for dialogue system: A startup perspective](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*.
- Sopan Khosla and Rashmi Gangadharaiah. 2022. [Benchmarking the covariate shift robustness of open-world intent classification approaches](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*.
- Jonathan K. Kummerfeld. 2021. [Quantifying and avoiding unfair qualification labour in crowdsourcing](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*.
- Stefan Larson, Adrian Cheung, Anish Mahendran, Kevin Leach, and Jonathan K. Kummerfeld. 2020a. [Inconsistencies in crowdsourced slot-filling annotations: A typology and identification methods](#). In *Proceedings of the 28th International Conference on Computational Linguistics (COLING)*.
- Stefan Larson and Kevin Leach. 2022a. [Redwood: Using collision detection to grow a large-scale intent classification dataset](#). In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*.
- Stefan Larson and Kevin Leach. 2022b. [A survey on intent classification and slot-filling datasets for task-oriented dialog](#). *arXiv preprint arXiv:2207.13211*.
- Stefan Larson, Gordon Lim, Yutong Ai, David Kuang, and Kevin Leach. 2022. [Evaluating out-of-distribution performance on document image classifiers](#). In *Proceedings of the Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Stefan Larson, Anish Mahendran, Andrew Lee, Jonathan K. Kummerfeld, Parker Hill, Michael A. Laurenzano, Johann Hauswald, Lingjia Tang, and Jason Mars. 2019a. [Outlier detection for improved data quality and diversity in dialog systems](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Stefan Larson, Anish Mahendran, Joseph J. Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael A. Laurenzano, Lingjia Tang, and Jason Mars. 2019b. [An evaluation dataset for intent classification and out-of-scope prediction](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Stefan Larson, Anthony Zheng, Anish Mahendran, Rishi Tekriwal, Adrian Cheung, Eric Guldan, Kevin Leach, and Jonathan K. Kummerfeld. 2020b. [Iterative feature mining for constraint-based data collection to increase data diversity and model robustness](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Jiexi Liu, Ryuichi Takanobu, Jiaxin Wen, Dazhen Wan, Hongguang Li, Weiran Nie, Cheng Li, Wei Peng, and Minlie Huang. 2021. [Robustness testing of language understanding in task-oriented](#)



- dialog. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*.
- Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. 2020. [Energy-based out-of-distribution detection](#). In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*.
- Xingkun Liu, Arash Eshghi, Pawel Swietojanski, and Verena Rieser. 2019a. [Benchmarking natural language understanding services for building conversational agents](#). In *Proceedings of the Tenth International Workshop on Spoken Dialog Systems Technology (IWSDS)*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [RoBERTa: A robustly optimized BERT pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Iaroslav Melekhov, Juho Kannala, and Esa Rahtu. 2016. [Siamese network features for image matching](#). In *Proceedings of the International Conference on Pattern Recognition (ICPR)*.
- Thanh-Do Nguyen, Chi Minh Bui, Thi-Hai-Yen Vuong, and Xuan-Hieu Phan. 2023. [Passage-based BM25 hard negatives: A simple and effective negative sampling strategy for dense retrieval](#). In *Proceedings of the 37th Pacific Asia Conference on Language, Information and Computation*.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. [Scikit-learn: Machine learning in Python](#). *Journal of Machine Learning Research*, 12(Oct):2825–2830.
- Baolin Peng, Chunyuan Li, Zhu Zhang, Chenguang Zhu, Jinchao Li, and Jianfeng Gao. 2021. [RADDLE: An evaluation benchmark and analysis platform for robust task-oriented dialog systems](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*.
- Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. [Generating natural language adversarial examples through probability weighted word saliency](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Andy Rosenbaum, Saleh Soltan, Wael Hamza, Yannick Versley, and Markus Boese. 2022. [LINGUIST: Language model instruction tuning to generate annotated utterances for intent classification and slot tagging](#). In *Proceedings of the 29th International Conference on Computational Linguistics (COLING)*.
- Gaurav Sahu, Pau Rodriguez, Issam Laradji, Parmida Atighehchian, David Vazquez, and Dzmitry Bahdanau. 2022. [Data augmentation for intent classification with off-the-shelf large language models](#). In *Proceedings of the 4th Workshop on NLP for Conversational AI*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *arXiv preprint arXiv:1910.01108*.
- Sailik Sengupta, Jason Krone, and Saab Mansour. 2021. [On the robustness of intent classification and slot labeling in goal-oriented dialog systems to real-world noise](#). In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*.
- Nihar Shah and Dengyong Zhou. 2016. [No oops, you won't do it again: mechanisms for self-correction in crowdsourcing](#). In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*.
- Petter Törnberg. 2023. [ChatGPT-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning](#). *arXiv preprint arXiv:2304.06588*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*.
- Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, Min Zhang, and Shaoping Ma. 2021. [Optimizing dense retrieval model training with hard negatives](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Jianguo Zhang, Kazuma Hashimoto, Yao Wan, Zhiwei Liu, Ye Liu, Caiming Xiong, and Philip Yu.



2022. [Are pre-trained transformers robust in intent classification? a missing ingredient in evaluation of out-of-scope intent detection.](#) In *Proceedings of the 4th Workshop on NLP for Conversational AI*.

Yiming Zhu, Peixian Zhang, Ehsan-Ul Haq, Pan Hui, and Gareth Tyson. 2023. [Can ChatGPT reproduce human-generated labels? a study of social computing tasks.](#) *arXiv preprint arXiv:2304.10145*.