

# OutFlip: Generating Out-of-Domain Samples for Unknown Intent Detection with Natural Language Attack

DongHyun Choi<sup>1,2</sup>, Myeong Cheol Shin<sup>1</sup>, EungGyun Kim<sup>1</sup>, and Dong Ryeol Shin<sup>2</sup>

<sup>1</sup>Kakao Enterprise

{*heuristic.c,index.sh,jason.ng*}@kakaenterprise.com

<sup>2</sup>Sungkyunkwan University

*drshin@skku.edu*

## Abstract

Out-of-domain (OOD) input detection is vital in a task-oriented dialogue system since the acceptance of unsupported inputs could lead to an incorrect response of the system. This paper proposes OutFlip, a method to generate out-of-domain samples using only in-domain training dataset automatically. A white-box natural language attack method HotFlip is revised to generate out-of-domain samples instead of adversarial examples. Our evaluation results showed that integrating OutFlip-generated out-of-domain samples into the training dataset could significantly improve an intent classification model’s out-of-domain detection performance<sup>1</sup>.

## 1 Introduction

Intent classification is crucial for task-oriented dialogue systems such as Google DialogFlow or Amazon Lex. It is vital for an intent classifier not only to map an input utterance into the correct label but also to detect *out-of-domain* (OOD) inputs. An accepted OOD input will lead the dialogue system to give erroneous responses.

Approaches for OOD detection in text classification could be classified into two major categories. Outlier detection approaches (Fei and Liu, 2016; Hendrycks and Gimpel, 2017; Shu et al., 2017; Lin and Xu, 2019; Yan et al., 2020; Xu et al., 2020) try to find out the boundaries of known classes in feature space. They need no labeled OOD dataset, but it is hard for them to deal with boundary cases. ( $n + 1$ )-way classification approaches (Kim and Kim, 2018; Larson et al., 2019; Ryu et al., 2018; Zheng et al., 2020) train classifiers for OOD detection using (pseudo)-labeled OOD samples. In practice, it is difficult and expensive to collect a large

number of labeled OOD samples with an open-world environment.

This paper proposes OutFlip, a method to generate OOD samples from in-domain training dataset automatically. For a given training dataset  $T$  and a reference intent classification model  $M$  which is trained with  $T$ , the OutFlip generates a set of OOD samples  $O$ . The generated OOD samples  $O$  could be used to train  $M$  iteratively to improve its OOD detection performance. Since the OutFlip does not require any modifications to the model architecture, it could be used with other OOD detection approaches to further improve the OOD detection performance.

The generated OOD samples should satisfy two conditions. First, they should be “hard-enough”; if the generated examples are too easy to distinguish from in-domain intents, they will be useless in training the OOD detector. Second, they should not belong to any in-domain intents. With a given reference model  $M$  and a set of in-domain labels  $I$ , this could be considered as finding a sentence  $\mathbf{x}_o$  with truth label  $y \notin I$  and model classification  $y' \in I$ . In this point of view, the OOD sample generation task could be considered as a variant of natural language attack on model  $M$ ; the goal of natural language attack on  $M$  is to find a  $\mathbf{x}_a$  with truth label  $y \in I$  and model classification  $y' \neq y$ . We revised HotFlip (Ebrahimi et al., 2018), a natural language attack method, to generate such OOD samples.

Our evaluation results showed that the generated OOD samples could significantly improve the OOD detection performances of the reference models. We also showed that applying OutFlip with other OOD detection approaches could further improve the model’s OOD detection performance. The evaluation results also suggest that the generated OOD samples could train the models other than the reference model to improve their OOD detection perfor-

<sup>1</sup>The source code is available at <https://github.com/kakaoenterprise/OutFlip>

mances.

Our contributions are summarized as follows:

- We proposed OutFlip, a simple and efficient OOD sample generation method using only in-domain training samples.
- We experimentally showed the effectiveness of our proposed approach using the intent classification benchmarks.
- We showed that the generated OOD samples could also improve the OOD detection performances of models other than the reference model.

## 2 Related Work

Previous OOD detection works could be classified into two major categories. Outlier detection approaches find boundaries of known classes in feature space. Fei and Liu (2016) computes a center for each class and transforms each document into a vector of similarities to the center. A binary classifier is built using the transformed data for each class. For deep learning-based systems, Hendrycks and Gimpel (2017) proposed the baseline of using softmax score as a threshold. Shu et al. (2017) trained the intent classifier using the sigmoid function and used the standard distribution to set each class’s score threshold. Lin and Xu (2019) first trained the classifier using Large Margin Cosine Loss (LMCL) (Nalisnick et al., 2018), and applied Local Outlier Factor (Breunig et al., 2000) to detect the OOD inputs. Yan et al. (2020) proposed a semantic-enhanced Gaussian mixture model to gather vectors of the same classes closely. Xu et al. (2020) calculated the mean and covariance of training samples for each class and used Mahalanobis distance as a distance function.

$(n + 1)$ -way classification approaches train the intent classifier with one additional class, where  $(n + 1)$ -th class represents the unseen intent. Kim and Kim (2018) proposed joint learning for in-domain and out-of-domain speeches. Larson et al. (2019) manually collected OOD samples to train intent classifiers. Ryu et al. (2018) generated OOD feature vectors using generative adversarial network (Goodfellow et al., 2014) to train an OOD detector. Since the approach proposed in Ryu et al. (2018) only works on continuous feature space, it highly depends on the feature encoder, which transforms inputs into feature vectors. Zheng

et al. (2020) also generated OOD feature vectors, but they also used unlabeled examples to enhance classification performance further. Although the  $(n + 1)$ -way classification approaches are easy to adopt without modification in the classification model, it is incredibly costly and time-consuming to collect the appropriate OOD samples.

The proposed OutFlip automatically generates OOD samples using the only in-domain training set, significantly reducing the cost of manually collecting OOD samples. Also, the OutFlip does not depend on the feature encoder.

The goal of adversarial attack in text classification is to fool a given text classification model  $M$ , by generating an adversarial example  $\mathbf{x}_a$  with truth label  $y$  and model classification  $y' \neq y$ . Many successful attacks first take a correctly classified example  $\mathbf{x}$  and replace its important words or characters to get an adversarial sample  $\mathbf{x}_a$ . In a white-box scenario, the attacker has access to the target model’s structure; thus, the important word or characters could be easily selected by inspecting the gradient of model  $M$ . HotFlip (Ebrahimi et al., 2018) estimates the best change of characters by maximizing the first-order approximation of the change in the loss.

In a black-box scenario, the attacker is not aware of the model or training data; the attacker is only capable of querying the target model with supplied inputs and obtaining the output predictions and their confidence scores. Alzantot et al. (2018) randomly selects a word from sentence  $\mathbf{x}$  and selects a suitable replacement word that has a similar semantic meaning. Jin et al. (2020) proposed a word importance score, which is used to find the word to be replaced. Li et al. (2020) applied BERT (Devlin et al., 2018) pre-trained language model to find a replacement word.

The proposed OutFlip first extracts important words using the algorithm proposed in Jin et al. (2020), and applies a variant of HotFlip (Ebrahimi et al., 2018) to generate OOD samples which are hard to distinguish from the in-domain intents by the given reference model.

## 3 Proposed Approach

In this section, the proposed OOD sample generation approach OutFlip is described in more detail.

### 3.1 HotFlip

We first introduce the white-box adversarial example generation method HotFlip (Ebrahimi et al., 2018). Let  $M$  be a text classification model,  $V$  be the word vocabulary set,  $\mathbf{x} = \{x_1; \dots; x_n\}$  be a sentence with  $n$  words where  $x_i \in \{0, 1\}^{|V|}$  denotes one-hot vector representing the  $i$ -th word, and  $L_M(\mathbf{x}, \mathbf{y})$  be the loss of  $M$  on input  $\mathbf{x}$  with true output  $\mathbf{y}$ . For a given sentence  $\mathbf{x}$ , a flip of the  $i$ -th word from  $w_a$  to  $w_b$  is represented by the following vector:

$$\vec{v}_{ib} = (\vec{0}; \dots; (0, \dots, -1, \dots, 1, 0, \dots, 0)_i; \dots; \vec{0}) \quad (1)$$

where -1 and 1 are in the corresponding positions for words  $w_a$  and  $w_b$  in the word vocabulary, respectively. A first-order approximation of the change in loss  $L_M(\mathbf{x}, \mathbf{y})$  can be obtained from a directional derivative along this vector:

$$\nabla_{\vec{v}_{ib}} L_M(\mathbf{x}, \mathbf{y}) = \nabla_x L_M(\mathbf{x}, \mathbf{y})^T \cdot \vec{v}_{ib} \quad (2)$$

Then, the HotFlip chooses the vector with the biggest increase in loss:

$$\begin{aligned} \max \nabla_x L_M(\mathbf{x}, \mathbf{y})^T \cdot \vec{v}_{ib} &= \max_{ib} \frac{\partial L_M^{(b)}}{\partial x_i} - \frac{\partial L_M^{(a)}}{\partial x_i} \\ \text{subject to } \text{sim}(w_a, w_b) &\geq T_{sim} \\ \text{and } POS(w_a) &= POS(w_b) \end{aligned} \quad (3)$$

where  $T_{sim}$  is a similarity threshold between two words, and  $POS(w_a)$  is the Part-of-Speech tagging of  $w_a$ . The two constraints are added to ensure that  $\mathbf{x}_a$  is semantically similar to the original input  $\mathbf{x}$ . With equation 3, the HotFlip could determine the flip position  $i$  and the replacement word  $w_b$ .

### 3.2 OutFlip

For a given reference model  $M$  and an in-domain sample  $\mathbf{x}$  with true output  $\mathbf{y}$ , the main idea of OutFlip is to flip the most important word of  $\mathbf{x}$ ,  $w_M(\mathbf{x})$ , to a semantically different word, while *minimizing* the change of loss  $L_M(\mathbf{x}, \mathbf{y})$ . By doing so, the OutFlip expects to get a sample  $\mathbf{x}_o$  whose truth label is different from the truth label of  $\mathbf{x}$ , while the model classifications of  $\mathbf{x}_o$  and  $\mathbf{x}$  are the same.

The word importance score proposed in Jin et al. (2020) is defined as follows:

$$I_{x_i}(M, \mathbf{x}) = o_y(M, \mathbf{x}) - o_y(M, \mathbf{x}_{\setminus x_i}) \quad (4)$$

---

#### Algorithm 1 OutFlip

---

**Input** In-domain training corpus  $T = \{\mathbf{x}_1, \dots, \mathbf{x}_t\}$ , in-domain labels  $Y$ , vocabulary  $V$ , reference model  $M$ , similarity threshold  $T_{sim}$

**Output** A set of OOD samples  $O$

```

1: for  $y \in Y$  do
2:   for  $\mathbf{x}$  with truth label  $y$  do
3:     Calculate  $w_M(\mathbf{x})$ 
4:   end for
5:    $C_T(y) \leftarrow$  Top 5 most frequent  $w_M(\mathbf{x})$ 
6: end for
7:
8: for  $\mathbf{x} \in T$  do
9:    $y \leftarrow$  truth label of  $\mathbf{x}$ 
10:  if  $w_M(\mathbf{x}) \in C_T(y)$  then
11:     $i \leftarrow$  the position of  $w_M(\mathbf{x})$ 
12:    Sort  $V$  in ascending order of  $\frac{\partial L_M^{(b)}}{\partial x_i}$ 
13:    Candidate  $\leftarrow$  Top 1% of  $V$ 
14:    Remain only words whose similarity
      with  $w_M(\mathbf{x})$  is less than  $T_{sim}$ 
15:    Randomly select  $w_b$  among candidates
16:    Replace  $w_M(\mathbf{x})$  with  $w_b$  to get  $\mathbf{x}_o$ 
17:    if  $M$  classifies  $\mathbf{x}_o$  to  $y$  then
18:      Add  $\mathbf{x}_o$  to  $O$ 
19:    end if
20:  end if
21: end for

```

---

where  $y$  is the truth label of  $\mathbf{x}$ ,  $o_y(M, \mathbf{x})$  is the logit output of the target model  $M$  for label  $y$ , and  $\mathbf{x}_{\setminus x_i}$  is the sentence after masking  $x_i$ . The most important word  $w_M(\mathbf{x})$  is defined as the word with the largest importance score in  $\mathbf{x}$ .

For each in-domain label  $y$  of the training dataset  $T$ , we define Core Class Token (CCT)  $C_T(y)$  as the top 5 most frequent  $w_M(\mathbf{x})$  among the training samples with truth label  $y$ . Since the importance score is calculated based on the reference model  $M$ , the OutFlip could select a wrong token as the most important token due to the model error. If the OutFlip flips such a word, the generated sentence's truth label will remain unchanged, leading to an erroneous OOD sample. To prevent such case, the OutFlip simply disregards  $\mathbf{x}$  during OOD generation process if  $w_M(\mathbf{x}) \notin C_T(y)$ .

In summary, the OutFlip chooses the replace-

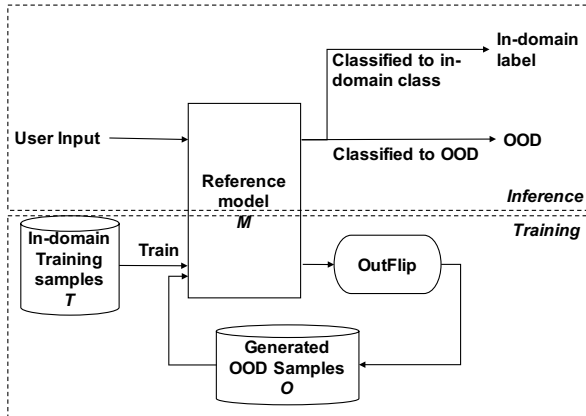


Figure 1: Applying the OutFlip to iteratively train the reference model  $M$  with newly generated OOD samples.

ment word  $w_b$  using the following equation:

$$\begin{aligned} \min \nabla_x L_M(\mathbf{x}, \mathbf{y})^T \cdot \vec{v}_{ib} &= \min_{ib} \frac{\partial L_M^{(b)}}{\partial x_i} - \frac{\partial L_M^{(a)}}{\partial x_i} \\ \text{subject to } \text{sim}(w_a, w_b) &\leq T_{\text{sim}} \\ \text{and } w_a &= w_M(\mathbf{x}) \\ \text{and } w_a &\in C_T(y) \end{aligned} \quad (5)$$

Since we do not need the generated OOD samples to be fluent, the part-of-speech condition is removed. The OutFlip randomly chooses  $w_b$  among the top 1% of the vocabulary in the ascending order of the loss change to generate more diverse samples. We used cosine similarity as the similarity measure.

The truth label of the generated sample  $\mathbf{x}_o$  could be an in-domain label different from  $y$  by chance. The OutFlip checks the model classification result of  $\mathbf{x}_o$  to see if it remains the same as  $\mathbf{x}$ . If the classification result changes, the OutFlip disregards  $\mathbf{x}_o$ . Algorithm 1 shows the pseudocode of the proposed OutFlip.

### 3.3 Iteratively Populating OOD samples

The reference model  $M$  could be iteratively trained with the generated OOD samples to improve its OOD detection performance. Figure 1 shows the overall framework. For each iteration, the set of generated OOD samples  $O$  is randomly split into training and dev set and used for the next train iteration.

Since the OutFlip does not require any change in model architecture, the OutFlip could be applied independently with other OOD detection algorithms

Dataset	ATIS	SNIPS	Kakao
Language	English	English	Korean
Vocab Size	938	12,054	22,831
Avg. Length	11.21	9.36	8.88
# Train	4,478	13,784	90,692
# Dev	500	700	11,310
# Test	893	-	12,711
# of Classes	18	7	48
Is Balanced	X	O	O

Table 1: Dataset statistics.

that require modifications on model architecture or loss function, such as Shu et al. (2017) or Lin and Xu (2019). In such cases, those OOD detection algorithms are applied to the examples classified as in-domain to filter out the OOD samples further.

## 4 Experiments

In this section, experimental settings and evaluation results are shown.

### 4.1 Datasets

Experiments are conducted on 3 real task-oriented dialogue datasets, SNIPS (Coucke et al., 2018), ATIS (Hemphill et al., 1990) and Kakao dialogue corpus<sup>2</sup> (Choi et al., 2020). SNIPS and ATIS are well-known English benchmarks. Kakao dialogue corpus is a Korean intent classification benchmark. We evaluated the proposed OutFlip with the Kakao dataset to see if it could be applied to different languages. Table 1 summarizes the statistics of the datasets. The ATIS dataset is highly imbalanced; more than 70% samples belong to one class, while three classes have less than 10 samples. The SNIPS and Kakao datasets are relatively balanced.

Since the SNIPS dataset does not have a test set, we randomly selected 30% of the training set and used them as the test set.

### 4.2 Baselines

We implemented two sentence encoders to show the generality of the proposed approach. LSTM (Hochreiter and Schmidhuber, 1997)-based encoder applies one-layer BiLSTM with output dimension 128 on the word embeddings of the given input; a self-attention layer with attention dimension 10 is followed to get the feature vector. CNN-based model applies the algorithm proposed in Kim

<sup>2</sup>Since the Kakao dataset is not publicly available, we contacted the authors to get the Kakao dataset and Korean GloVe vectors.

Dataset % of known intents	ATIS			SNIPS			Kakao		
	25%	50%	75%	25%	50%	75%	25%	50%	75%
<b>MSP</b> <sub>cnn</sub>	64.61	71.07	62.00	34.24	63.26	73.14	66.23	82.32	89.60
<b>MSP</b> <sub>lstm</sub>	65.82	71.52	61.18	34.12	63.30	73.28	64.23	81.37	88.28
<b>DOC</b> <sub>cnn</sub>	61.99	57.82	38.46	49.66	70.76	77.16	71.61	84.72	<b>90.55</b>
<b>DOC</b> <sub>lstm</sub>	62.76	58.15	38.35	49.77	71.11	77.37	63.36	78.86	85.22
<b>LMCL</b> <sub>cnn</sub>	71.67	74.89	68.73	61.51	<b>84.37</b>	88.31	80.27	<b>87.15</b>	<b>90.60</b>
<b>LMCL</b> <sub>lstm</sub>	72.25	<b>77.90</b>	<b>73.18</b>	69.52	83.32	87.53	76.31	<b>85.79</b>	89.39
<b>OutFlip</b> <sub>cnn</sub>	<b>74.18</b>	<b>79.23</b>	<b>69.37</b>	<b>79.20</b>	<b>84.25</b>	<b>88.99</b>	<b>81.78</b>	85.45	86.85
<b>OutFlip</b> <sub>lstm</sub>	<b>73.85</b>	74.50	68.30	<b>79.26</b>	84.00	<b>88.67</b>	<b>81.96</b>	84.64	86.42

Table 2: Comparisons of the OutFlip and previous OOD detection works. The top 2 results for each metric are marked in bold.  $T_{sim}$  is set to 0.3, and the OutFlip is applied for three iterations.

(2014). More precisely, one-dimensional convolutions with kernel sizes 2, 3, 4, 5 and filter size 32 are applied on top of the word embeddings. The results are max-pooled to get the feature vector. For both encoders, a dense layer is applied to the feature vector to get the logit of each class.

We also implemented three baseline OOD detection systems, as follows:

1. **Maximum Softmax Probability (MSP)** (Hendrycks and Gimpel, 2017) considers the maximum softmax probability of a sample as the rejection score. If the probability is below a certain threshold, the sample is classified as OOD. We used the threshold of 0.5, as the authors suggested.
2. **Deep Open Classification (DOC)** (Shu et al., 2017) replaces softmax with sigmoid activation as the final layer to calculate the score for each class separately. It also calculates the threshold for each class through a statistical approach.
3. **Large Margin Cosine Loss (LMCL)** (Lin and Xu, 2019) replaces the softmax loss with large margin cosine loss (Nalisnick et al., 2018), to force the model to maximize inter-class variance and minimize intra-class variance. After training, it applies Local Outlier Factor (LOF) (Breunig et al., 2000) on training features vectors to detect outliers as OOD. We set the scaling factor  $s = 30$  and cosign margin  $m = 0.35$ , following the authors.

By combining two feature encoders and three baseline OOD detection systems, we implemented eight baseline reference models, six with an OOD detection system and two without.

Pre-trained Embeddings	Accuracy
<b>GloVe</b> (Pennington et al., 2014)	83.17 %
<b>Korean GloVe</b> (Choi et al., 2020)	51.33 %

Table 3: Evaluation results of GloVe embeddings on the language-independent set of word analogy corpus.

### 4.3 Experimental Setup

Word embeddings are initialized with GloVe (Pennington et al., 2014) pre-trained word vectors. We downloaded the pre-trained embeddings containing 1.9M words trained on 42B tokens from the author’s homepage. For Korean, Korean pre-trained GloVe embedding vectors proposed in Choi et al. (2020) are used. The dimensions of both pre-trained embeddings are 300.

We removed some classes from the train/dev set during training and integrated them back during testing, following the evaluation settings of Fei and Liu (2016); Shu et al. (2017); Lin and Xu (2019). We varied the number of known intents in the training dataset as 25%, 50%, and 75% of the intents, and used all intents for testing. We randomly select known intents by weighted random sampling without replacement in the training set. Note that the samples belonging to the unknown intents are removed during training and validation.

Following Fei and Liu (2016); Shu et al. (2017); Lin and Xu (2019), macro F1 score is used to evaluate the models. For each known intent selection, the F1 score for each class is calculated separately. Then the results are macro-averaged across all classes. We reported the average of 10 random known intent selections for each evaluation.

For each OutFlip iteration, 90% of the generated OOD samples are added to the training set, and

Dataset % of known intents	ATIS			SNIPS			Kakao		
	25%	50%	75%	25%	50%	75%	25%	50%	75%
<b>MSP</b> <sub><i>cnn</i></sub>	77.04 (+12.43)	74.66 (+3.59)	62.44 (+0.44)	78.02 (+43.78)	84.48 (+21.22)	89.66 (+16.52)	82.65 (+16.42)	85.74 (+3.42)	87.26 (-2.34)
<b>MSP</b> <sub><i>lstm</i></sub>	71.37 (+5.55)	76.17 (+4.65)	61.97 (+0.79)	79.88 (+45.76)	83.98 (+20.68)	88.92 (+15.64)	82.22 (+17.99)	85.21 (+3.84)	86.25 (-2.03)
<b>DOC</b> <sub><i>cnn</i></sub>	73.45 (+11.46)	59.03 (+1.21)	41.24 (+2.78)	80.72 (+31.06)	85.05 (+14.29)	89.42 (+12.26)	82.20 (+10.59)	85.45 (+0.73)	86.44 (-4.11)
<b>DOC</b> <sub><i>lstm</i></sub>	70.20 (+7.44)	59.69 (+1.54)	41.09 (+2.74)	81.99 (+32.22)	85.08 (+13.97)	89.69 (+12.32)	80.14 (+16.78)	81.28 (+2.42)	80.78 (-4.44)
<b>LMCL</b> <sub><i>cnn</i></sub>	74.13 (+2.46)	76.35 (+1.46)	68.82 (+0.09)	81.04 (+19.53)	85.90 (+1.53)	90.10 (+1.79)	81.81 (+1.54)	84.75 (-2.40)	85.74 (-4.86)
<b>LMCL</b> <sub><i>lstm</i></sub>	75.44 (+3.19)	76.37 (-1.53)	71.34 (-1.84)	80.86 (+11.34)	84.56 (+1.24)	88.76 (+1.23)	82.24 (+5.93)	84.19 (-1.60)	85.35 (-4.04)

Table 4: Evaluation results of the baselines with the OutFlip applied.  $T_{sim}$  is set to 0.3, and the OutFlip is applied for three iterations. Small numbers below the macro F1 score represents the performance improvement compared to the same baseline without applying the OutFlip.

the remaining 10% are added to the dev set. The populated train/dev sets are used for the next train iteration.

Adam optimizer (Kingma and Ba, 2015) with an initial learning rate of 0.001 is used to train the model. The training batch size is set to 128. Exponential learning rate decay with a decay rate of 0.8 is applied for every two epochs. On each epoch, the trained classifier is evaluated against the dev set, and the training stops when the dev accuracy is not improved for five consequent epochs.

#### 4.4 Evaluation Results

Table 2 shows the evaluation results of the proposed OutFlip and other baseline systems. As can be observed from the table, the proposed OutFlip outperforms other baselines when the number of known intents is small. The small number of known intents is the most similar case to real-world applications, since in the open-world environment, the number of unknown intents is much larger than the number of known intents. The OutFlip also gives comparable results for ATIS and SNIPS corpus with the larger number of known intents.

For the Kakao corpus, the OutFlip performance is lower compared to the other baselines. To figure out the reason, the qualities of English GloVe embeddings and Korean GloVe embeddings are compared. We used 4 out of 14 categories in the word analogy corpus (Mikolov et al., 2013) for fair comparison; capital-common-countries, capital-world, currency and a subset of

family. We removed all the syntactic questions since they cannot be translated into Korean words one-to-one. Part of family category is removed for the same reason. We also removed categories that give an advantage on English pre-trained embeddings; for example, the city-in-state category is removed because it contains relationships between US cities and US states. The remaining 6,168 questions are manually translated into Korean.

Table 3 shows the evaluation results of the English and Korean GloVe vectors on our subset of the word analogy corpus. As can be observed, the accuracy of Korean GloVe is much lower compared to the English GloVe vectors. Since the OutFlip relies on the cosine similarities between pre-trained embedding vectors to generate the OOD samples, the quality of embedding vectors is critical to the OOD generation performance.

Next, we applied the proposed OutFlip to other OOD detection baselines to see if the OutFlip could further improve their performance. Table 4 shows the evaluation results. In most cases applying the OutFlip to other OOD detection approaches resulted in performance improvement. The performance improvement was significant when the dataset is balanced, and the number of known intents is small.

We conducted a set of experiments to find out the best OutFlip iteration number and  $T_{sim}$  value. Figure 2 and Figure 3 shows the OutFlip performances on the benchmark datasets with changing




Figure 2: Evaluation results of **OutFlip**<sub>cnn</sub> with changing  $T_{sim}$ . We fixed the number of OutFlip iteration to 2.




Figure 3: Evaluation results of **OutFlip**<sub>cnn</sub> with changing OutFlip iteration number. We fixed  $T_{sim}$  to 0.3.

$T_{sim}$  values and OutFlip iteration numbers, respectively. Increasing  $T_{sim}$  value would cause the OutFlip to generate more challenging examples, but the chance of developing wrong OOD samples would also increase.

Figure 2 suggests that the balanced dataset like SNIPS could easily recover the errors introduced from large  $T_{sim}$  values. In contrast, the ATIS dataset’s macro F1 score decreases with the increased  $T_{sim}$  values when many intents are known. Since 3 out of 18 ATIS intents have less than ten sentences, one or two erroneous OOD samples could lead to a performance drop. The macro F1 score of the balanced Kakao dataset does not increase with the  $T_{sim}$  values larger than 0.3. Since the quality of Korean GloVe is relatively low, large  $T_{sim}$  values introduce more errors compared to the English datasets.

As can be observed from Figure 3, in most cases, the macro F1 score converges with two to three OutFlip iterations. Additional OutFlip iterations give small or no performance improvements for balanced datasets and decrease macro F1 score for unbalanced dataset ATIS by introducing more errors.

One advantage of the OutFlip is that the generated OOD samples could be used to train and improve the OOD detection performance of the models other than the reference model without applying additional OutFlip iterations. We trained the BERT-base and BERT-large models (Devlin et al.,

2018) with the ATIS and SNIPS benchmarks. As the same as previous experiments, the unknown intents are removed during training and integrated back during testing. Besides, we added OOD samples generated using reference models **OutFlip**<sub>cnn</sub> and **OutFlip**<sub>lstm</sub> with three OutFlip iterations and  $T_{sim}$  value 0.3, while training the BERT models.

Table 5 shows the evaluation results of the BERT models trained with the OutFlip-generated OOD samples. As can be observed from the table, the OutFlip-generated OOD samples significantly improved the OOD detection performances of BERT models, regardless of the reference models used to generate the OOD samples.

#### 4.5 Error Analysis

We randomly selected 200 samples from the OOD examples generated by **OutFlip**<sub>cnn</sub> with three iterations and  $T_{sim} = 0.3$  for the ATIS dataset, when 75% of the intents are known. The number of newly generated OOD samples for each iteration is shown in Table 6. We manually analyzed the selected examples for errors.

Among the 200 examples, 186 of them were correctly generated OOD sentences. Out of 14 error cases, 12 were due to the wrongfully extracted Core Class Tokens. Some ATIS intents have too few examples to extract Core Class Tokens; for example, intent `atis_restriction` has only six samples. Also, an entity which shows up too frequently could also lead to wrong Core Class Token

Model	Reference model for OutFlip	ATIS			SNIPS		
		25%	50%	75%	25%	50%	75%
BERT-base	<i>None</i>	62.14	74.73	72.61	34.66	60.76	70.49
	<b>OutFlip<sub>cnn</sub></b>	70.60	76.08	72.45	<b>74.51</b>	<b>81.09</b>	<b>86.92</b>
	<b>OutFlip<sub>lstm</sub></b>	<b>71.58</b>	<b>78.22</b>	<b>72.68</b>	73.02	81.04	86.50
BERT-large	<i>None</i>	62.17	74.01	73.32	34.56	61.07	70.51
	<b>OutFlip<sub>cnn</sub></b>	<b>76.02</b>	80.68	<b>79.54</b>	70.95	<b>80.53</b>	<b>87.06</b>
	<b>OutFlip<sub>lstm</sub></b>	74.55	<b>81.46</b>	75.38	<b>73.15</b>	80.31	86.60

Table 5: Results of applying OutFlip-generated samples to train models other than the reference model. Reference model *None* means no OutFlip-generated OOD samples are added during training.

Iteration	1	2	3	4
<b>ATIS</b>	2003	1648.5	755.5	336.6
<b>SNIPS</b>	4923.5	4644.2	3221.5	1956.2

Table 6: The number of generated OOD samples for each OutFlip iteration with reference model **OutFlip<sub>cnn</sub>**,  $T_{sim} = 0.3$ . 75% of the intents are known. Numbers are the average of 10 known intent selections.

extraction result. For intent `atis_flight`, 797 out of 4,334 samples contain the entity “Denver”.

For one case, the OutFlip-generated sentence accidentally belongs to the other in-domain intent. However, due to the reference model’s error, the OutFlip fails to remove the generated sentence. A training instance “Can you list the cheapest round trip **fare** from Orlando to Kansas City” (truth label `atis_airfare`) is converted to a sentence “Can you list the cheapest round trip **airplane** from Orlando to Kansas City” (truth label `atis_flight`), but the reference model classifies the converted sentence to `atis_airfare`. Since the classification result remains the same, the OutFlip considers the generated sentence as “hard-enough” OOD sample.

The ATIS dataset allows an instance to have multiple labels; two or more labels are assigned to 23 ATIS training instances. The OutFlip failed to properly handle those instances. The remaining one error case is generated from a training instance with two assigned labels.

## 5 Conclusion

In this paper, we proposed OutFlip, a method to generate OOD samples using only in-domain training dataset. Our evaluation results showed that the proposed OutFlip could significantly improve the OOD detection performance of an intent classification model by iteratively generating difficult OOD samples. Since OutFlip does not require any

modifications to model architecture, it could be used with other OOD detection approaches to improve OOD detection performance further. We also showed that the generated OOD samples could be used to train and improve the OOD detection performance of models other than the reference model, without applying additional OutFlip iterations.

Currently, we only focused on generating difficult OOD samples, which can fool the reference model. However, generating meaningful OOD samples could also be beneficial, since then the dialogue engine developer could check the generated OOD samples to find new intents. As our future work, we will focus on generating meaningful, fluent OOD samples.

## 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*, pages 2890–2896.
- Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. 2000. Lof: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data.*, pages 93–104.
- DongHyun Choi, IINam Park, Myeong Cheol Shin, EungGyun Kim, and Dong Ryeol Shin. 2020. Integrated eojeol embedding for erroneous sentence classification in korean chatbots. *Computing Research Repository*, arXiv:2004.05744.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, Maël Primet, and Joseph Dureau. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *Computing Research Repository*, arXiv:1805.10190.



- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- 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 (Volume 2: Short Papers)*, pages 31–36.
- Geli Fei and Bing Liu. 2016. Breaking the closed world assumption in text classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 506–514.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27:2672–2680.
- 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, USA, June 24-27*.
- Dan Hendrycks and Kevin Gimpel. 2017. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *Proceedings of International Conference on Learning Representations*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is bert really robust? a strong baseline for natural language attack on text classification and entailment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8018–8025.
- Joo-Kyung Kim and Young-Bum Kim. 2018. Joint learning of domain classification and out-of-domain detection with dynamic class weighting for satisficing false acceptance rates. In *Proceedings of InterSpeech 2018*, pages 556–560.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1292–1302.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, volume 5.
- 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. 2019. 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)*, pages 1311–1316.
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020. Bert-attack: Adversarial attack against bert using bert. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 6193–6202.
- Ting-En Lin and Hua Xu. 2019. Deep unknown intent detection with margin loss. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5491–5496.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. volume arXiv:1301.3781.
- Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. 2018. Do deep generative models know what they don’t know? In *International Conference on Learning Representations*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Seonghan Ryu, Sangjun Koo, Hwanjo Yu, and Gary Geunbae Lee. 2018. Out-of-domain detection based on generative adversarial network. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 714–718.
- Lei Shu, Hu Xu, and Bing Liu. 2017. Doc: Deep open classification of text documents. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2911–2916.
- Hong Xu, Keqing He, Yuanmeng Yan, Sihong Liu, Zijun Liu, and Weiran Xu. 2020. A deep generative distance-based classifier for out-of-domain detection with mahalanobis space. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1452–1460.
- Guangfeng Yan, Lu Fan, Qimai Li, Han Liu, Xiaotong Zhang, Xiao-Ming Wu, and Albert Y.S. Lam. 2020. Unknown intent detection using gaussian mixture model with an application to zero-shot intent classification. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1050–1060.
- Yinhe Zheng, Guanyi Chen, and Minlie Huang. 2020. Out-of-domain detection for natural language understanding in dialog systems. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, pages 1198–1209.