

Modeling Discriminative Representations for Out-of-Domain Detection with Supervised Contrastive Learning

Zhiyuan Zeng^{1*}, Keqing He^{2*}, Yuanmeng Yan¹, Zijun Liu¹, Yanan Wu¹
Hong Xu¹, Huixing Jiang², Weiran Xu^{1*}

¹Pattern Recognition & Intelligent System Laboratory

¹Beijing University of Posts and Telecommunications, Beijing, China

²Meituan Group, Beijing, China

{zengzhiyuan, yanyuanmeng, liuzijun, yanan.wu, xuhong}@bupt.edu.cn
{hekeqing, jianghuixing}@meituan.com, {xuweiran}@bupt.edu.cn

Abstract

Detecting Out-of-Domain (OOD) or unknown intents from user queries is essential in a task-oriented dialog system. A key challenge of OOD detection is to learn discriminative semantic features. Traditional cross-entropy loss only focuses on whether a sample is correctly classified, and does not explicitly distinguish the margins between categories. In this paper, we propose a supervised contrastive learning objective to minimize intra-class variance by pulling together in-domain intents belonging to the same class and maximize inter-class variance by pushing apart samples from different classes. Besides, we employ an adversarial augmentation mechanism to obtain pseudo diverse views of a sample in the latent space. Experiments on two public datasets prove the effectiveness of our method capturing discriminative representations for OOD detection. ¹

1 Introduction

Detecting Out-of-Domain (OOD) or unknown intents from user queries is an essential component in a task-oriented dialog system (Gnewuch et al., 2017; Akasaki and Kaji, 2017; Tulshan and Dhage, 2018; Shum et al., 2018). It aims to know when a user query falls outside their range of predefined supported intents to avoid performing wrong operations. Different from normal intent detection tasks, we do not know the exact number of unknown intents in practical scenarios and can barely annotate extensive OOD samples. Lack of real OOD examples leads to poor prior knowledge about these unknown intents, making it challenging to identify OOD samples in the task-oriented dialog system.

*The first two authors contribute equally. Weiran Xu is the corresponding author.

¹Our code is available at <https://github.com/arZival27/supervised-contrastive-learning-for-out-of-domain-detection>.

Previous methods of OOD detection can be generally classified into two types: supervised and unsupervised OOD detection. Supervised OOD detection (Scheirer et al., 2013; Fei and Liu, 2016; Kim and Kim, 2018; Larson et al., 2019; Zheng et al., 2020; Zeng et al., 2021b) represents that there are extensive labeled OOD samples in the training data. In contrast, unsupervised OOD detection (Bendale and Boult, 2016; Hendrycks and Gimpel, 2017; Shu et al., 2017; Lee et al., 2018; Ren et al., 2019; Lin and Xu, 2019; Xu et al., 2020; Zeng et al., 2021a) means no labeled OOD samples except for labeled in-domain data. Specifically, for supervised OOD detection, Fei and Liu (2016); Larson et al. (2019), form a $(N + 1)$ -class classification problem where the $(N + 1)$ -th class represents the unseen intents. Further, Zheng et al. (2020) uses labeled OOD data to generate an entropy regularization term to enforce the predicted distribution of OOD inputs closer to the uniform distribution. However, these methods heavily rely on large-scale time-consuming labeled OOD data. Compared to these supervised methods, unsupervised OOD detection first learns discriminative intent representations via in-domain (IND) data, then employs detecting algorithms, such as Maximum Softmax Probability (MSP) (Hendrycks and Gimpel, 2017), Local Outlier Factor (LOF) (Lin and Xu, 2019), Gaussian Discriminant Analysis (GDA) (Xu et al., 2020) to compute the similarity of features between OOD samples and IND samples. In this paper, we focus on the unsupervised OOD detection.

A key challenge of unsupervised OOD detection is to learn discriminative semantic features via IND data. We hope to cluster the same type of IND intents more tightly and separate different types of IND intents further. Traditional softmax loss (Hendrycks and Gimpel, 2017) only focuses on whether the sample is correctly classified, and does not explicitly distinguish the relationship between

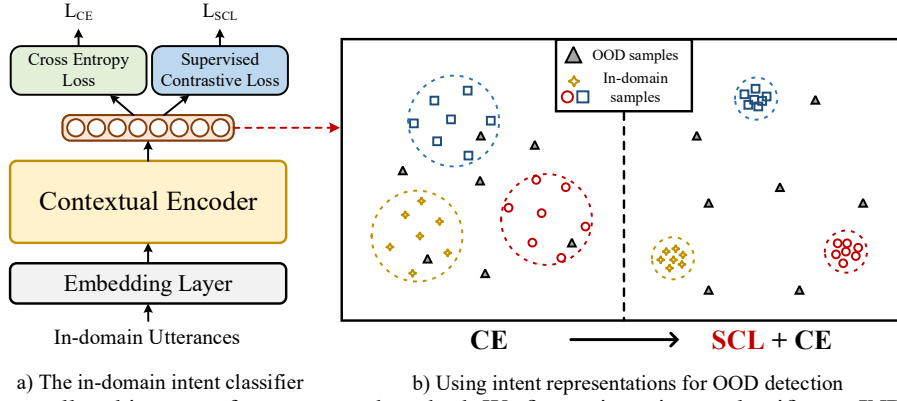


Figure 1: The overall architecture of our proposed method. We first train an intent classifier on IND data using CE or SCL+CE objectives. Then, we extract the intent representation of a test sample to detect OOD.

categories. Further, Lin and Xu (2019) proposes a large margin cosine loss (LMCL) (Wang et al., 2018) which maximizes the decision margin in the latent space. LMCL forces the model to not only classify correctly but also maximize inter-class variance and minimize intra-class variance. Following the similar motivation, we aim to pull intents belonging to the same class together while simultaneously pushing apart samples from different classes to further model discriminative semantic features.

In this paper, we propose a supervised contrastive learning (SCL) model to learn discriminative semantic intent representation for OOD detection. SCL aims to minimize intra-class variance by pulling together IND intents belonging to the same class and maximize inter-class variance by pushing apart samples from different classes. Empirical results demonstrate the effectiveness of discriminative representation for OOD detection. Besides, to enhance the diversity of data augmentation in SCL, we employ an adversarial attack mechanism to obtain pseudo hard positive samples in the latent space by computing model-agnostic adversarial worst-case perturbations to the inputs. Our contributions are three-fold: (1) To the best of our knowledge, we are the first to apply supervised contrastive learning to OOD detection. (2) Compared to cross-entropy (CE) loss, SCL+CE can maximize inter-class variance and minimize intra-class variance to learn discriminative semantic representation. (3) Extensive experiments and analysis on two public datasets demonstrate the effectiveness of our method.

2 Methodology

Overall Architecture Fig 1 shows the overall architecture of our proposed method. As Fig 1(a) displays, we first train an IND intent classifier us-

ing CE or SCL+CE objectives in the training stage. Then in the test stage, we extract the intent feature of a test sample and employ the detection algorithms MSP (Hendrycks and Gimpel, 2017), LOF (Lin and Xu, 2019) or GDA (Xu et al., 2020) to detect OOD.² Fig 1(b) demonstrates the effectiveness of our method capturing discriminative intent representations, where SCL+CE can maximize inter-class variance and minimize intra-class variance.

Supervised Contrastive Learning We first review the classic cross-entropy (CE) loss and its improved version, large margin cosine loss (LMCL). Then we explain our supervised contrastive loss (SCL) in detail. Given an IND sample x_i and its intent label y_i , we adopt a BiLSTM (Hochreiter and Schmidhuber, 1997) or BERT (Devlin et al., 2019) encoder to get the intent representation s_i . The CE loss and LMCL are defined as follows³:

$$\mathcal{L}_{CE} = \frac{1}{N} \sum_i -\log \frac{e^{W_{y_i}^T s_i / \tau}}{\sum_j e^{W_j^T s_i / \tau}} \quad (1)$$

$$\mathcal{L}_{LMCL} = \frac{1}{N} \sum_i -\log \frac{e^{W_{y_i}^T s_i / \tau}}{e^{W_{y_i}^T s_i / \tau} + \sum_{j \neq y_i} e^{(W_j^T s_i + m) / \tau}} \quad (2)$$

where N denotes the number of training samples, y_i is the ground-truth class of the i -th sample, τ is the temperature factor, W_j is the weight vector of the j -th class, and m is the cosine margin. Compared to CE, LMCL adds a normalized decision margin on the negative classes and forces the model to explicitly distinguish positive class and negative classes. Our experiment 3.2 shows LMCL can slightly improve the performance of OOD detec-

²In this paper, we focus on the first training stage. Thus we dive into the details about the detection algorithms MSP, LOF and GDA in the appendix.

³For brevity, we omit the L2 normalization on both features and weight vectors for LMCL.

Models		CLINC-Full				CLINC-Small			
		IND		OOD		IND		OOD	
		ACC	F1	Recall	F1	ACC	F1	Recall	F1
LSTM	CE	86.34	87.73	63.72	65.23	84.24	84.30	60.40	61.07
	LMCL	86.83	87.90	64.14	65.79	84.46	84.87	60.72	61.89
	SCL+CE(ours)	87.01	88.28	66.80	67.68	85.73	86.61	63.96	64.44
	SCL+LMCL(ours)	87.37	88.60	66.92	68.04	85.93	87.02	64.16	64.70
BERT	CE	88.13	88.98	64.24	66.17	86.68	86.20	61.64	62.58
	LMCL	88.57	89.12	64.76	66.80	86.76	86.64	62.20	63.11
	SCL+CE(ours)	88.97	89.57	66.84	68.03	87.65	88.07	64.44	64.52
	SCL+LMCL(ours)	89.20	90.03	67.28	68.21	87.87	88.30	64.64	65.01

Table 1: Performance comparison on CLINC-Full and CLINC-Small datasets ($p < 0.05$ under t-test).

tion. To further model discriminative intent representations, motivated by recent contrastive learning work (Chen et al., 2020; He et al., 2020; Khosla et al., 2020; Gunel et al., 2020), we propose a supervised contrastive learning objective to minimize intra-class variance and maximize inter-class variance:

$$\mathcal{L}_{SCL} = \sum_{i=1}^N -\frac{1}{N_{y_i} - 1} \sum_{j=1}^N \mathbf{1}_{i \neq j} \mathbf{1}_{y_i = y_j} \log \frac{\exp(s_i \cdot s_j / \tau)}{\sum_{k=1}^N \mathbf{1}_{i \neq k} \exp(s_i \cdot s_k / \tau)} \quad (3)$$

where N_{y_i} is the total number of examples in the batch that have the same label as y_i and $\mathbf{1}$ is an indicator function. Note that we only perform SCL on the IND data since we focus on the unsupervised OOD detection where no labeled OOD data exists. As Fig 1(b) shows, SCL aims to pull together IND intents belonging to the same class and pushing apart samples from different classes, which helps recognize OOD intents near the decision boundary. In the implementation, we first pre-train the intent classifier using SCL, then finetune the model using CE or LMCL, both on the IND data. We compare iterative training and joint training in the appendix.

Adversarial Augmentation Chen et al. (2020) has proved the necessity of data augmentation for contrastive learning. However, there is no simple and effective augmentation strategy in the NLP area, which requires much handcrafted engineering. Thus, we apply adversarial attack (Goodfellow et al., 2015; Kurakin et al., 2016; Jia and Liang, 2017; Zhang et al., 2019; Yan et al., 2020) to generate pseudo positive samples to increase the diversity of views for contrastive learning. Specifically, we need to compute the worst-case perturbation δ that maximizes the original cross-entropy loss \mathcal{L}_{CE} : $\delta = \arg \max_{\|\delta'\| \leq \epsilon} \mathcal{L}_{CE}(\theta, x + \delta')$, where θ represents the parameters of the intent classifier and x denotes a given sample. ϵ is the norm bound of

the perturbation δ . We apply Fast Gradient Value (FGV) (Rozsa et al., 2016) to approximate the perturbation δ :

$$\delta = \epsilon \frac{g}{\|g\|}; \text{ where } g = \nabla_x \mathcal{L}_{CE}(f(x; \theta), y) \quad (4)$$

We perform normalization to g and then use a small ϵ to ensure the approximate is reasonable. Finally, we can obtain the pseudo augmented sample $x_{adv} = x + \delta$ in the latent space. The pseudo samples are applied to augment positive views per anchor in SCL. Ablation study 3.3 shows adversarial augmentation significantly improves the performance of SCL for OOD detection.

3 Experiments

3.1 Setup

Datasets We use two benchmark OOD datasets, CLINC-Full and CLINC-Small (Larson et al., 2019). We report IND metrics: Accuracy(Acc) and F1, and OOD metrics: Recall and F1. OOD Recall and F1 are the main evaluation metrics in this paper. **Baselines** We adopt LSTM and BERT as our intent classifier and compare SCL with CE and LMCL. Since only using SCL can't classify in-domain intents directly, we first pre-train the classifier using SCL, then finetune the model using CE or LMCL, both on the IND data. We use three OOD detection algorithms MSP, LOF and GDA to verify the generalization capability of SCL. We present dataset statistics, implementation details, and results on MSP and LOF in the appendix.

3.2 Main Results

Tab 1 displays the main results on GDA. Combining SCL and CE/LMCL significantly outperforms all the baselines, both on OOD and IND metrics. For OOD metrics, using SCL+CE in LSTM outperforms CE by 3.08%(Recall) and 2.45%(F1) on CLINC-Full, 3.56%(Recall) and 3.37%(F1) on CLINC-Small. Similar improvements based on

models		min	max	mean	median
LSTM	CE	1.13E-07	2.63E-04	4.23E-05	1.61E-05
	SCL+CE	4.35E-08	1.85E-04	3.23E-05	1.39E-05
BERT	CE	8.26E-08	2.23E-04	3.84E-05	1.56E-05
	SCL+CE	2.86E-08	1.67E-04	3.05E-05	1.36E-05

Table 2: Intra-class variance statistics.

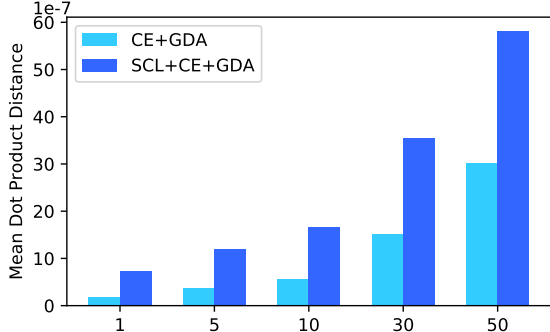


Figure 2: Comparison between inter-class distances.

LMCL are observed. The results prove the effectiveness of SCL for OOD detection. For IND metrics, using SCL+CE in LSTM outperforms CE by 0.67%(ACC) and 0.55%(F1) on CLINC-Full, 1.49%(ACC) and 2.31%(F1) on CLINC-Small. The results confirm SCL also helps IND intent detection. The difference between OOD and IND improvements is probably attributed to metric scale and data imbalance in the original test set. Besides, SCL gains higher improvements on CLINC-Small than CLINC-Full, which displays the advantage of our approach in the few-shot scenario (see details in Section 3.3). SCL also gets consistent improvements on BERT by 2.60%(Recall) and 1.86%(F1) on CLINC-Full OOD metrics, 0.84%(ACC) and 0.59%(F1) on CLINC-Full IND metrics, substantiating our method is model-agnostic for different OOD detection architectures.

3.3 Analysis

Analysis of IND feature distribution. We analyze the representation distribution of IND data on CLINC-Full dataset from two perspectives, intra-class and inter-class. We choose SCL+CE based on GDA to perform analysis. Tab 2 shows the statistics of intra-class variance, which can indicate the degree of clustering of intra-class data representations. Specifically, we average the variances of each sample normalized representation with the same intent label to its cluster center in the test set as cluster intra-class variance, then report min/max/mean/median values on all cluster intra-class variances. Results show SCL effectively decreases intra-class variances, especially in terms of max and mean values, which confirms SCL can

Proportion		10%	20%	30%	40%	50%
IND F1	CE	63.31	70.77	77.84	81.55	84.30
	SCL+CE	69.50	75.14	81.45	84.18	86.61
	Relative↑	9.78%	6.17%	4.64%	3.23%	2.74%
OOD F1	CE	42.16	48.34	53.00	57.92	61.07
	SCL+CE	50.10	54.43	58.61	62.12	64.44
	Relative↑	18.83%	12.60%	10.58%	7.25%	5.52%

Table 3: Effect of training data size.

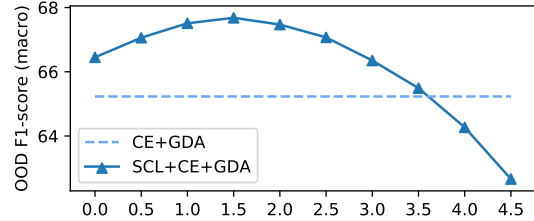


Figure 3: Effect of Adversarial Perturbation Norm.

converge intra-class intent representations.

Fig 2 shows the inter-class distances. We average dot product distances between each class center to its k nearest class centers, then average results of all classes as inter-class distance. The X-axis denotes the number of k. We observe a significant increase in SCL+CE compared to CE. When k is smaller, the increase is more obvious. It verifies SCL can maximize inter-class variance and distinguish intent classes. We also provide visualization analysis in the appendix. In summary, SCL can pull together IND intents belonging to the same class and push apart samples from different classes, which makes representations more discriminative. **Effect of IND Training Data Size.** Tab 3 shows the effect of IND training data size. We randomly choose training data with a certain proportion from CLINC-Full IND data and use the original test set for evaluation. We use the LSTM+GDA setting. Results show SCL+CE consistently outperforms CE. Besides, with the decrease of training data size, the relative improvements gradually increase. It proves SCL has strong robustness for improving OOD detection, especially in the few-shot scenario. **Analysis of Adversarial Perturbation Norm.** Fig 3 shows the effect of adversarial perturbation norm ϵ on OOD detection performance. We conduct the experiments on CLINC-Full dataset, using LSTM and GDA. The X-axis denotes the value of ϵ . The CE+GDA dashed line means no SCL pre-training and $\epsilon = 0.0$ in the SCL+CE+GDA solid line means no adversarial augmentation. In general, both SCL and adversarial augmentation contribute to the improvements and $\epsilon \in (1.0, 2.0)$ achieves better performances. Compared with the baseline without SCL, the SCL+CE method with a smaller adversarial perturbation can still obtain

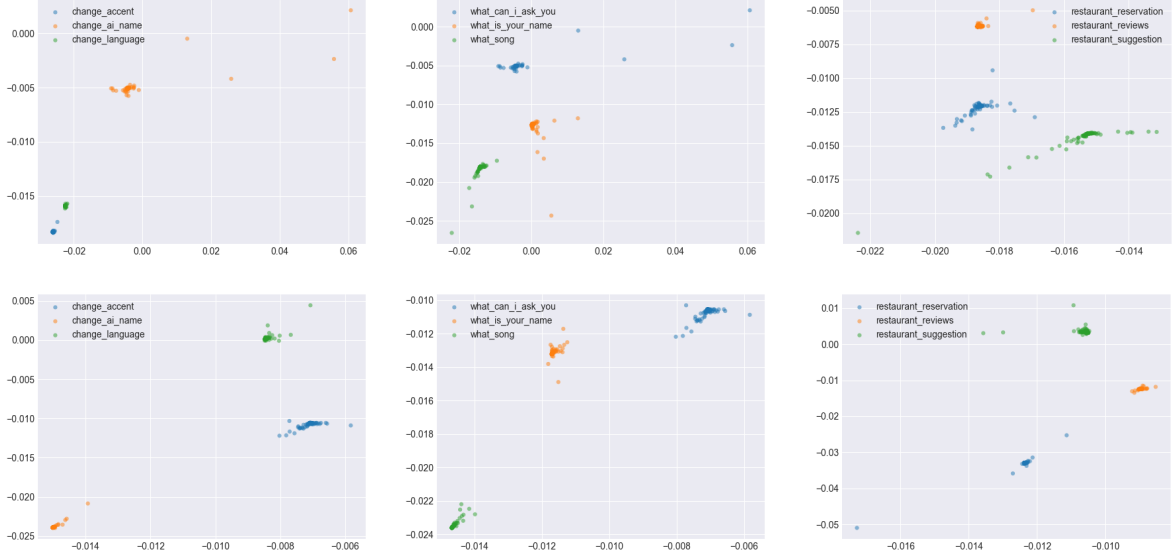


Figure 4: Visualization of in-domain representation distribution.

settings		IND		OOD	
dimension	batch size	ACC	F1	Recall	F1
128	50	87.01	88.28	66.80	67.68
128	100	87.52	88.60	67.08	68.12
128	200	88.10	89.05	67.56	68.63
256	50	87.17	88.40	66.96	67.92
256	100	87.85	88.96	67.32	68.37
256	200	88.37	89.24	67.76	68.94
512	50	87.35	88.78	67.32	68.47
512	100	88.14	89.22	67.64	68.69
512	200	88.54	89.50	68.00	69.27

Table 4: Parameter analysis of batch size and representation dimension.

better results but lower than the results with an optimal range of perturbation, while large norms tend to damage the effect of SCL. Our method still performs well with a broad range of adversarial perturbation and is insensitive to hyperparameters.

Parameter Analysis. As our proposed SCL is a method involving contrastive learning, we analyze batch sizes and representation dimensions to further verify the effectiveness, whose results are presented in Table 4. We conduct experiments in CLINC-Full dataset, using LSTM and SCL+CE objective for training and GDA for detection. With the increase of batch size and representation dimension, both in-domain and OOD metrics are slightly improved. However, compared with the method proposed in this paper, the improvement is relatively limited. In general, our proposed method is not sensitive to hyperparameters and can show the expected effect under a wide range of reasonable settings.

Feature Visualization. As shown in Fig 4, we extract several groups of similar classes for PCA visualization analysis. The three pictures in the

upper part represent training using only CE, while the three pictures in the lower part use SCL+CE for training. In the same column, we sample the same classes for observation. It is worth noting that the scale of the image has been adjusted adaptively in order to display all the data. The actual distance can be sensed by observing the marking of the coordinate axis. After SCL is added, the distance between similar classes is significantly expanded, and the data in the same classes are more closely clustered.

4 Conclusion

In this paper, we focus on the unsupervised OOD detection where no labeled OOD data exist. To learn discriminative semantic intent representations via in-domain data, we propose a novel supervised contrastive learning loss (SCL). SCL aims to minimize intra-class variance by pulling together in-domain intents belonging to the same class and maximize inter-class variance by pushing apart samples from different classes. Experiments and analysis confirm the effectiveness of SCL for OOD detection. We hope to provide new guidance for future OOD detection work.

Acknowledgements

This work was partially supported by National Key R&D Program of China No. 2019YFF0303300 and Subject II No. 2019YFF0303302, DOCOMO Beijing Communications Laboratories Co., Ltd, MoE-CMCC "Artificial Intelligence" Project No. MCM20190701.

Broader Impact

Task-oriented dialog systems have demonstrated remarkable performance across a wide range of applications, with the promise of a significant positive impact on human production mode and lifeway. However, in scenarios where information is complex and rapidly changing, models usually face input that is meaningfully different from typical examples encountered during training. Current models are prone to make unfounded predictions on these inputs, which may affect human judgment and thus impair the safety of models in practical applications. In domains with the greatest potential for societal impacts, such as navigation or medical diagnosis, models should be able to detect potentially agnostic OOD and be robust to high-entropy inputs to avoid catastrophic errors. This work proposes a novel unsupervised OOD detection method that using supervised contrastive learning to learn discriminative semantic intent representations. The effectiveness and robustness of the model are significantly improved by adding a supervised contrastive learning pre-training stage, which takes a step towards the ultimate goal of enabling the safe real-world deployment of task-oriented dialog systems in safety-critical domains. The experimental results have been reported on standard benchmark datasets for considerations of reproducible research.

References

- Satoshi Akasaki and Nobuhiro Kaji. 2017. Chat detection in an intelligent assistant: Combining task-oriented and non-task-oriented spoken dialogue systems. *ArXiv*, abs/1705.00746.
- Abhijit Bendale and Terrance E. Boult. 2016. Towards open set deep networks. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1563–1572.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. A simple framework for contrastive learning of visual representations. *ArXiv*, abs/2002.05709.
- J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Geli Fei and Bing Liu. 2016. Breaking the closed world assumption in text classification. In *HLT-NAACL*.
- Ulrich Gnewuch, S. Morana, and A. Maedche. 2017. Towards designing cooperative and social conversational agents for customer service. In *ICIS*.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. *CoRR*, abs/1412.6572.
- Beliz Gunel, Jingfei Du, Alexis Conneau, and Ves Stoyanov. 2020. Supervised contrastive learning for pre-trained language model fine-tuning. *ArXiv*, abs/2011.01403.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. Momentum contrast for unsupervised visual representation learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9726–9735.
- Dan Hendrycks and Kevin Gimpel. 2017. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *ArXiv*, abs/1610.02136.
- S. Hochreiter and J. Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9:1735–1780.
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. *arXiv preprint arXiv:1707.07328*.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, A. Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *ArXiv*, abs/2004.11362.
- Joo-Kyung Kim and Young-Bum Kim. 2018. Joint learning of domain classification and out-of-domain detection with dynamic class weighting for satisfying false acceptance rates. *ArXiv*, abs/1807.00072.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*.
- 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, Hong Kong, China. Association for Computational Linguistics.
- Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. 2018. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *ArXiv*, abs/1807.03888.

- Ting-En Lin and Hua Xu. 2019. Deep unknown intent detection with margin loss. In *ACL*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Jie Ren, Peter J. Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark A. DePristo, Joshua V. Dillon, and Balaji Lakshminarayanan. 2019. Likelihood ratios for out-of-distribution detection. *ArXiv*, abs/1906.02845.
- Andras Rozsa, Ethan M Rudd, and Terrance E Boulton. 2016. Adversarial diversity and hard positive generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 25–32.
- Walter J. Scheirer, Anderson Rocha, Archana Sapkota, and Terrance E. Boulton. 2013. Toward open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35:1757–1772.
- Lei Shu, Hu Xu, and Bing Liu. 2017. Doc: Deep open classification of text documents. In *EMNLP*.
- H. Shum, X. He, and Di Li. 2018. From eliza to xiaoice: challenges and opportunities with social chatbots. *Frontiers of Information Technology & Electronic Engineering*, 19:10–26.
- Amrita S. Tulshan and S. Dhage. 2018. Survey on virtual assistant: Google assistant, siri, cortana, alexa.
- Feng Wang, Jian Cheng, Weiyang Liu, and H. Liu. 2018. Additive margin softmax for face verification. *IEEE Signal Processing Letters*, 25:926–930.
- 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, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Yuanmeng Yan, Keqing He, H. Xu, Sihong Liu, Fanyu Meng, Min Hu, and Weiran Xu. 2020. Adversarial semantic decoupling for recognizing open-vocabulary slots. In *EMNLP*.
- Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Hong Xu, and Weiran Xu. 2021a. Adversarial self-supervised learning for out-of-domain detection. In *NAACL*.
- Zhiyuan Zeng, Hong Xu, Keqing He, Yuanmeng Yan, Sihong Liu, Zijun Liu, and Weiran Xu. 2021b. [Adversarial generative distance-based classifier for robust out-of-domain detection](#). In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7658–7662.
- Huangzhao Zhang, Hao Zhou, Ning Miao, and Lei Li. 2019. Generating fluent adversarial examples for natural languages. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5564–5569.
- 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*, 28:1198–1209.

A Dataset Details

Table 5 shows the details of two benchmark OOD dataset⁴ CLINC-Full and CLINC-Small (Larson et al., 2019). They both contain 150 in-domain intents across 10 domains. It is worth noting that our paper does not use labeled OOD data from the training set in the training stage.

CLINC	Full	Small
Avg utterance length	9	9
Intents	150	150
Training set size	15100	7600
Training samples per class	100	50
Training OOD samples amount	100	100
Development set size	3100	3100
Development samples per class	20	20
Development OOD samples amount	100	100
Testing Set Size	5500	5500
Testing samples per class	30	30
Development OOD samples amount	1000	1000

Table 5: Statistics of the CLINC datasets.

B Baseline Details

We compare many types of unsupervised OOD detection models. Therefore, the model proposed in this paper can be divided into the training stage and detection stage. For each model LSTM or BERT, we use different detection methods to verify its performance. The innovation of this paper focuses mainly on the training stage. Due to the limitation of space, we do not detailed introduce the detection methods in the main body. We will supplement the relevant contents as follows:

MSP (Maximum Softmax Probability)(Hendrycks and Gimpel, 2017) applies a threshold on the maximum softmax probability where the threshold is set to 0.5 according to the dev set.

LOF (Local Outlier Factor)(Lin and Xu, 2019) uses the local outlier factor to detect unknown intents. The motivation is that if an example’s local density is significantly lower than its k-nearest neighbor’s, it is more likely to be considered as the unknown intents.

GDA (Gaussian Discriminant Analysis)(Xu et al., 2020) is a generative distance-based classifier for out-of-domain detection with Euclidean space. They estimate the class-conditional distribution on feature spaces of DNNs via Gaussian discriminant analysis (GDA) to avoid over-confidence problems

⁴<https://github.com/clinc/oos-eval>

and use Mahalanobis distance to measure the confidence score of whether a test sample belongs to OOD. GDA is the state-of-the-art detection methods till now, so we adopt GDA as our main detection algorithm. We also report MSP and LOF results in Section D.

C Implementation Details

We use the public pre-trained 300 dimensions GloVe embeddings (Pennington et al., 2014)⁵ or bert-base-uncased (Devlin et al., 2019)⁶ model to embed tokens. We use a single-layer BiLSTM as a feature extractor and set the dimension of hidden states to 128. The dropout value is fixed at 0.5. We use Adam optimizer (Kingma and Ba, 2014) to train our model. We set a learning rate to 1E-03 for GloVe+LSTM and 1E-04 for Bert. In the training stage, 100 epochs of supervised contrastive training are first conducted, then 10 epochs of finetune training are conducted with CE or LMCL. Both phases are training only on in-domain labeled data. The training stage has an early stop setting with patience equal to 5. We use the best F1 scores on the validation set to calculate the GDA threshold adaptively. Each result of the experiments is tested 5 times under the same setting and gets the average value. The norms of adversarial perturbation are obtained by the heuristic method, in which MSP and LOF are 1.0 and GDA is 1.5. The training stage of our model lasts about 10 minutes using GloVe embeddings, and 18 minutes using Bert-base-uncased, both on a single Tesla T4 GPU(16 GB of memory). The average value of the trainable model parameters is 3.05M.

D Supplementary Experimental Results

Various Detection Methods In this paper, the experiments and analysis are mainly conducted around the training stage. Different detection models are used to verify the generalization of our proposed method. Due to the limitation of space, we use GDA for most of the presentation in the main body. The main experiments of LOF and MSP using LSTM feature extractor are shown in Table 6. It is worth noting that using different detection methods can obtain the same analysis results as the main experimental in the main body.

Combining two training stages in different ways We display results of different combining

⁵<https://github.com/stanfordnlp/GloVe>

⁶<https://github.com/google-research/bert>

Models		CLINC-Full				CLINC-Small			
		IND		OOD		IND		OOD	
		ACC	F1	Recall	F1	ACC	F1	Recall	F1
LOF	CE	85.46	85.80	57.40	58.78	82.45	82.73	52.88	53.90
	LMCL	85.87	86.08	58.32	59.28	82.83	82.98	53.96	54.63
	SCL+CE(ours)	86.52	86.80	60.72	61.80	83.13	83.39	56.88	57.48
	SCL+LMCL(ours)	86.94	87.15	61.88	63.03	83.40	83.57	57.92	58.60
MSP	CE	85.76	86.27	27.12	34.91	83.81	84.12	20.40	22.76
	LMCL	87.36	87.62	31.28	36.66	85.02	85.30	24.16	25.72
	SCL+CE(ours)	87.44	87.87	33.68	39.34	85.54	85.95	27.24	27.43
	SCL+LMCL(ours)	88.89	89.21	35.40	41.75	86.87	87.20	29.28	31.02

Table 6: Supplementary experimental results of LOF and MSP.

models	IND		OOD	
	ACC	F1	Recall	F1
CE	86.34	87.73	63.72	65.23
CE+SCL	82.29	83.59	61.96	63.40
multitask	86.69	88.02	65.76	67.25
SCL+CE	87.01	88.28	66.80	67.68

Table 7: Results of combining two training stages in different ways

ways of two training stages on CLINC-Full dataset using LSTM and GDA detection method in Table 7. CE is the baseline that only uses the cross-entropy loss function to train the feature extractor. SCL+CE follows the paradigm of pre-training first and then finetuning, which achieves the best performance. Besides, we try two different combinations to explore the relationship between the two training stages. CE+SCL means that we first conduct training to minimize cross-entropy loss, and then conduct supervised contrastive learning. The results show that the subsequent SCL leads to a decline in metrics, especially on in-domain. This is because SCL, while optimizing the representation distribution, compromises the mapping relationship with labels. Multitask means to optimize two losses simultaneously. This setting leads to mutual interference between two tasks, which affects the convergence effect and damages the performance and stability of the model. In general, SCL should be used as a pre-training method and CE as a finetuning method. The best results can be achieved by first using SCL to learn discriminative representation and then finetuning the model by CE.