

Disentangling Confidence Score Distribution for Out-of-Domain Intent Detection with Energy-Based Learning

Yanan Wu^{1*}, Zhiyuan Zeng^{1*}, Keqing He^{2*}, Yutao Mou¹,
Pei Wang¹, Yuanmeng Yan¹, Weiran Xu^{1*}

¹Beijing University of Posts and Telecommunications, Beijing, China

²Meituan, Beijing, China

{yanan.wu, zengzhiyuan, myt, wangpei,
yanyuanmeng, xuweiran}@bupt.edu.cn
{hekeqing}@meituan.com

Abstract

Detecting Out-of-Domain (OOD) or unknown intents from user queries is essential in a task-oriented dialog system. Traditional softmax-based confidence scores are susceptible to the overconfidence issue. In this paper, we propose a simple but strong energy-based score function to detect OOD where the energy scores of OOD samples are higher than IND samples. Further, given a small set of labeled OOD samples, we introduce an energy-based margin objective for supervised OOD detection to explicitly distinguish OOD samples from INDs. Comprehensive experiments and analysis prove our method helps disentangle confidence score distributions of IND and OOD data.¹

1 Introduction

Detecting Out-of-Domain (OOD) or unknown intents from user queries is crucial to a task-oriented dialog system (Akasaki and Kaji, 2017; Tulshan and Dhage, 2018; Shum et al., 2018; Lin and Xu, 2019; Xu et al., 2020; Zeng et al., 2021a; Wu et al., 2022b). It can avoid performing wrong operations and provide potential directions of future development when an input query falls outside the range of predefined intents. Since the exact number of unknown intents in practical scenarios is hard to know and annotate, the lack of real OOD examples makes it challenging to identify these samples in dialog systems.

Depending on whether labeled OOD samples are available, previous OOD detection work can be generally classified into two types: unsupervised (Bendale and Boulton, 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;

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

¹Our code is available at https://github.com/pris-nlp/EMNLP2022-energy_for_OOD/.

Case	Ground Truth	Softmax Pred
1. Can you give me a meal suggestion from the south	meal_suggestion	meal_suggestion
2. Give me a suggestion for roofers	OOD	meal_suggestion

Figure 1: IND (case 1) vs OOD sample (case 2). Softmax score recognizes OOD sample as IND intent type because of overconfidence issue.




Figure 2: Softmax score from MSP vs energy score from our method. Softmax score are similar for IND and OOD (both > 0.85) but energy score are more distinguished.

Zeng et al., 2021a,b; Wu et al., 2022a) and supervised (Fei and Liu, 2016; Kim and Kim, 2018; Larson et al., 2019a; Zheng et al., 2020). The former firstly learn an in-domain (IND) intent classifier only using labeled IND data and then estimates the confidence score of a test query. For example, Maximum Softmax Probability (MSP) (Hendrycks and Gimpel, 2017) uses maximum softmax probability as the confidence score and regards an intent as OOD if the score is below a fixed threshold. The assumption is that OOD intents should produce a lower softmax probability than INDs. However, neural networks can produce arbitrarily high softmax confidence even for such abnormal OOD samples (Guo et al., 2017; Liang et al., 2018), as shown in Fig 1&2, which we call *overconfidence*. Further, another distance-based method, Gaussian discriminant analysis (GDA) (Xu et al., 2020), is proposed to use the maximum Mahalanobis distance (Mahalanobis, 1936) to all in-domain classes centroids as the confidence score. Compared to MSP, GDA gets better OOD performance but requires expensive computation for complex Mahalanobis distance.




Figure 3: The overall architecture of our proposed method.

In this paper, we aim to use simple softmax confidence scores for both higher performance and efficiency. For supervised OOD detection, Fei and Liu (2016); Larson et al. (2019a), form a $(N+1)$ -class classification problem where the $(N+1)$ -th class represents the OOD intents. Further, Zheng et al. (2020) uses labeled OOD data to generate an entropy regularization term. But these methods require numerous labeled OOD intents to get superior results. We focus on using fewer labeled OOD data (like 20 or 30) to achieve comparable even better performance.

In this paper, we propose an energy-based score function to detect OOD in an unsupervised manner. The energy-based score function maps each query to a single energy scalar which is lower for IND samples and higher for OOD samples based on the energy theory (LeCun et al., 2006). We first train an in-domain intent classifier via IND data, then replace the original softmax layer with the energy-based score function. Our method can not only mitigate the issue of overconfident softmax probability but also reduce expensive post-processing computation. Further, given a small portion of labeled OOD samples, we propose an energy-based margin objective to explicitly distinguish OOD samples from IND samples. Our contributions are three-fold: (1) We propose an energy-based learning method for OOD intent detection to achieve higher performance and efficiency. (2) We propose an energy-based margin objective to distinguish energy distributions of OOD and IND samples. (3) Extensive experiments and analysis on two benchmarks demonstrate the effectiveness of our method.

2 Methodology

Overall Architecture Fig 3 (a) shows the overall architecture of our proposed method. We first train

an in-domain intent classifier using IND data in training stage. Then in the test stage, we extract the intent feature of a test query and employ the detection algorithms MSP (Hendrycks and Gimpel, 2017) or Energy to detect OOD. Fig 3 (b) demonstrates the effectiveness of our method distinguishing OOD distributions from IND².

Energy-based Score Function To mitigate the issue of overconfident softmax probability in MSP, we propose an energy-based score function to push apart score distributions of OOD and IND samples. We first briefly review the energy theory (LeCun et al., 2006) then explain our proposed energy-based score function for OOD detection. The previous energy work (LeCun et al., 2006; Zhai et al., 2016; Grathwohl et al., 2020; Liu et al., 2020b; Kaur et al., 2021) aims to build a function $E(\mathbf{x}) : R^D \rightarrow R$ which maps a sample \mathbf{x} to a single scalar called the *energy*. Given a data point $\mathbf{x} \in R^D$, the energy function can be defined as follows:

$$E(\mathbf{x}) = -T \cdot \log \int_{y'} e^{-E(\mathbf{x}, y')/T} \quad (1)$$

where T is the temperature parameter and $E(\mathbf{x}, y')$ is the marginal energy over label y' . Essentially, energy scores can be transferred to the likelihood probability:

$$p(y | \mathbf{x}) = \frac{e^{-E(\mathbf{x}, y)/T}}{\int_{y'} e^{-E(\mathbf{x}, y')/T}} = \frac{e^{-E(\mathbf{x}, y)/T}}{e^{-E(\mathbf{x})/T}} \quad (2)$$

For OOD detection, since we focus on the detection algorithms for the test stage in this paper, we train the same BiLSTM in-domain intent classifier

²Because the max softmax score is higher for IND samples and lower for OOD samples, we use the negative energy score to align with the conventional definition where positive (IND) samples get higher scores.

Models		CLINC-Full				CLINC-Small			
		IND		OOD		IND		OOD	
		Acc	F1	Recall	F1	Acc	F1	Recall	F1
Unsupervised OOD	MSP (Hendrycks and Gimpel, 2017)	87.16	87.64	41.40	44.86	85.02	85.18	35.81	36.60
	LOF (Lin and Xu, 2019)	85.87	86.08	58.32	59.28	82.83	82.98	53.96	54.63
	GDA (Xu et al., 2020)	86.83	87.90	64.14	65.79	84.46	84.87	60.72	61.89
	SCL (Zeng et al., 2021a)	87.01	88.28	66.80	67.68	85.73	86.61	63.96	64.44
	Energy (Ours)	88.71	89.17	68.10	69.64	86.42	86.48	65.78	66.52
Supervised OOD	N+1	91.24	85.29	24.51	31.08	90.13	83.23	21.50	29.17
	MSP+Entropy (Zheng et al., 2020)	87.48	87.81	49.90	53.93	85.24	85.31	45.90	48.57
	MSP+Bound (Liu et al., 2020a)	88.03	87.26	45.21	56.86	86.16	83.04	42.38	51.43
	MSP+Margin (Ours)	88.31	87.98	57.27	59.96	85.33	85.37	54.90	55.37
	LOF+Entropy	85.98	86.37	61.10	61.13	83.49	83.86	57.70	57.79
	LOF+Bound	86.36	85.66	57.83	60.15	81.36	82.88	64.41	59.30
	LOF+Margin (Ours)	86.13	86.59	65.70	65.59	83.57	83.97	63.60	63.18
	GDA+Entropy	87.27	88.14	68.53	68.82	85.01	85.53	65.22	65.65
	GDA+Bound	87.09	86.86	67.32	66.41	84.44	84.75	65.19	64.14
	GDA+Margin (Ours)	87.54	88.23	68.42	68.73	85.51	85.81	65.13	65.68
Energy+Margin (Ours, Full Model)	89.75	89.46	73.92	74.06	87.84	87.53	72.76	72.98	

Table 1: Performance comparison on CLINC-Full and CLINC-Small datasets ($p < 0.01$ under t-test).

$f(\mathbf{x})$ via IND data as Lin and Xu (2019) in the training stage. Then given a test query, we simply use the logits from the intent classifier to represent $E(\mathbf{x}, y')$. Therefore, the energy score function Eq 1 can be formulated as:

$$E(\mathbf{x}; f) = -T \cdot \log \sum_i^K e^{f_i(\mathbf{x})/T} \quad (3)$$

where K is the size of IND intent classes and $f_i(\mathbf{x})$ is the logit of \mathbf{x} belonging to i -th class. We simply use a threshold on the energy score to consider whether a test query belongs to OOD. Intuitively, the reason why the energy score works for OOD detection is that higher energy represents a lower likelihood of occurrence according to LeCun et al. (2006). Therefore, unobserved OOD samples in the training stage should get lower likelihoods as well as higher energy scores than observed IND samples. In Appendix C, we provide a detailed theoretical derivation of why the energy function can alleviate the overconfidence problem. Besides, Experiment 4.1 also proves energy scores better distinguish confidence distribution of OOD data from IND data than softmax probabilities.

Energy-guided Margin Objective To further distinguish OOD from IND, we propose an energy-guided margin objective for few-shot supervised OOD detection. Different from Liu et al. (2020a), our approach directly models the energy boundary by pushing apart the samples from IND and OOD, which helps recognize OOD intents near the decision boundary and is easier to tune and less sensitive to the noise. Specifically, we use an energy-based max-margin loss as well as the standard cross-entropy loss to explicitly set an energy

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 2: Statistics of the CLINC datasets.

gap between OOD and IND. We aim to learn more discriminative representations for energy score distributions in the training stage. The energy margin loss is formulated as:

$$\mathcal{L} = \mathbb{E}_{(\mathbf{x}_{\text{ind}}, \mathbf{x}_{\text{ood}}) \sim \mathcal{D}} \max(0, m + E(\mathbf{x}_{\text{ind}}) - E(\mathbf{x}_{\text{ood}})) \quad (4)$$

where m is the energy margin and E is the energy score of IND or OOD samples in train set. Then in the test stage, we still use the energy score to detect OOD. Analysis 4.1 displays the effectiveness of the margin loss over unsupervised OOD.

3 Experiments

3.1 Datasets

We use two public benchmark OOD datasets³, CLINC-Full and CLINC-Small (Larson et al., 2019b). We show the detailed statistic of these

³<https://github.com/clinc/oos-eval>




Figure 4: Distribution of softmax scores vs energy scores.

datasets in Table 2. They both contain 150 in-domain intents across 10 domains. The difference is that CLINC-Small has fewer in-domain training examples than CLINC-Full. Note that all the datasets we used have a fixed set of labeled OOD data but we don’t use it for training.

3.2 Metrics

We report both OOD metrics: Recall and F1-score (F1) and in-domain metrics: F1-score (F1) and Accuracy (ACC). Since we aim to improve the performance of detecting out-of-domain intents from user queries, OOD Recall and F1 are the main evaluation metrics in this paper.

3.3 Baselines

For detection algorithms, we use MSP, LOF and GDA as baselines. For training objectives, we use N+1, entropy and bound as baselines. We present dataset statistics, baselines and implementation details in the appendix. We will release our code after blind review.

3.4 Main Results

Table 1 shows the main results. (1) For unsupervised OOD detection, using the energy function achieves 24.78, 10.36, 3.85, 1.96 OOD F1 improvements over MSP, LOF, GDA and SCL on CLINC-Full. The results prove the effectiveness of energy score function for OOD detection. Besides, for IND metrics, energy function also outperforms SCL by 0.89% (F1), which reflects energy scores can better distinguish OOD from IND samples without sacrificing IND performance. (2) For supervised OOD detection, we compare different pre-training losses under the same detection score function. We find our Margin approach achieve consistent improvements under different detection functions on both datasets. It demonstrates that

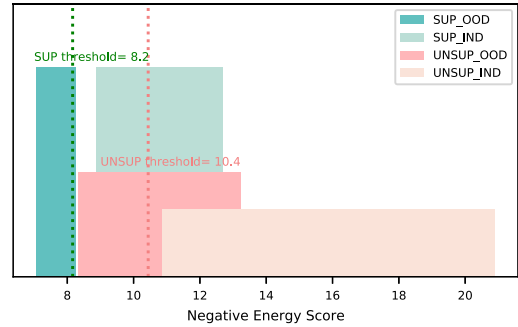


Figure 5: Unsupervised vs supervised OOD detection.

Margin objective can stably improve the representation space by directly pushing apart the samples from IND and OOD. We also observe under MSP, our proposed Margin objective outperforms Entropy by 6.03% and Bound by 3.10% on CLINC-Full. But on GDA we find no significant performance difference. We argue the energy-based learning may not always fit in generative distance-based detection methods like GDA. Overall, combining energy score function and margin objective achieve the best performance over the previous state-of-the-art by 5.24%.

4 Analysis

4.1 Distribution of softmax scores vs energy scores

To figure out why energy scores outperform softmax scores, we compare the score histogram distributions for IND and OOD data in Fig 4. We use the same pre-trained intent classifier to compute scores on the test set. We find softmax scores for both IND and OOD data concentrate on high values, resulting in severe overconfidence. By contrast, energy scores better distinguish score distribution of OOD data from IND data. And energy distribu-




Figure 6: Effect of number of labeled OOD samples.

tions are smoother than softmax score distributions. Overall, our proposed energy-based score function can disentangle confidence score distributions for IND and OOD data.

4.2 Unsupervised vs supervised OOD detection

To verify the effectiveness of our proposed energy-based margin objective, we compare the energy score statistics of unsupervised (Energy) and supervised (Margin+Energy) OOD detection in Fig 5. Each rectangle in Fig 5 represents the energy distribution of IND or OOD data, where the middle of the rectangle is energy mean and the width of the rectangle is energy variance. Results show that compared to Energy, Margin+Energy makes negative energy scores of both OOD and IND data smaller. Further, the supervised Margin objective can significantly decrease the variance of both OOD (1.86 ↓) and IND (3.11 ↓) data. Therefore, Margin can push apart energy score distributions for OOD detection by shrinking its variance to avoid overlapping. Besides, combined with the energy threshold (dot line in Fig 5), unsupervised (Energy) still gets a portion of OOD samples above the threshold which are misclassified into IND, but supervised (Margin+Energy) on the opposite. It proves that Margin can further mitigate the issue of overconfidence.

4.3 Effect of number of labeled OOD samples

Fig 6 shows the effect of labeled OOD training data size for supervised OOD detection. We find Margin+Energy consistently outperforms Entropy+Energy, especially in the few-shot supervised OOD scenario, which demonstrates strong robustness and generalization of our proposed energy-based margin objective for OOD detection.




Figure 7: Effect of energy temperature T




Figure 8: Effect of energy margin m

4.4 Effect of Parameters

Temperature T . Fig 7 shows the effect of different energy temperature T . We conduct the experiments on the CLINC-Full dataset, using Energy for unsupervised OOD. The X-axis denotes the value of temperature T . In general, $T \in (0.5, 1.0)$ achieves relatively better performances and has a broad range.

Margin m . Fig 8 shows the effect of different energy margin m . We conduct the experiments on the CLINC-Full dataset, using Margin+Energy for supervised OOD. The X-axis denotes the value of margin m . Results show that $m = 19.0$ achieves the best performance and is robust to minor changes.

5 Conclusion

Traditional softmax-based OOD detection methods are susceptible to the overconfidence issue. Therefore, we propose a novel energy-based score function to mitigate the issue of softmax overconfidence. To use labeled OOD data, we further introduce an energy-based margin objective to explicitly distinguish energy score distributions of OOD from IND. Experiments and analysis confirm the effectiveness of our energy-based method for OOD detection. For future work, we hope to explore theoretical concepts of energy and provide new guidance.

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.
- 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.
- Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. 2020. Your classifier is secretly an energy based model and you should treat it like one. *ArXiv*, abs/1912.03263.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR.
- Dan Hendrycks and Kevin Gimpel. 2017. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *ArXiv*, abs/1610.02136.
- Ramneet Kaur, Susmit Jha, Anirban Roy, O. Sokolsky, and Insup Lee. 2021. Are all outliers alike? on understanding the diversity of outliers for detecting oods. *ArXiv*, abs/2103.12628.
- 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. *ArXiv*, abs/1807.00072.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Stefan Larson, Anish Mahendran, Joseph Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael Laurenzano, Lingjia Tang, and Jason Mars. 2019a. An evaluation dataset for intent classification and out-of-scope prediction. In *EMNLP/IJCNLP*.
- 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 *EMNLP-IJCNLP*.
- Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and F Huang. 2006. A tutorial on energy-based learning. volume 1.
- 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.
- Shiyu Liang, Yixuan Li, and R. Srikant. 2018. Enhancing the reliability of out-of-distribution image detection in neural networks. *arXiv: Learning*.
- 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.
- Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. 2020a. Energy-based out-of-distribution detection. *Advances in Neural Information Processing Systems*.
- Weitang Liu, Xiaoyun Wang, John Douglas Owens, and Yixuan Li. 2020b. Energy-based out-of-distribution detection. *ArXiv*, abs/2010.03759.
- Prasanta Chandra Mahalanobis. 1936. On the generalized distance in statistics. National Institute of Science of India.
- 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.
- 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.
- 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 Sudhir Namdeorao Dhage. 2018. Survey on virtual assistant: Google assistant, siri, cortana, alexa. In *International symposium on signal processing and intelligent recognition systems*, pages 190–201.
- Yanan Wu, Keqing He, Yuanmeng Yan, QiXiang Gao, Zhiyuan Zeng, Fujia Zheng, Lulu Zhao, Huixing Jiang, Wei Wu, and Weiran Xu. 2022a. Revisit overconfidence for ood detection: Reassigned contrastive learning with adaptive class-dependent threshold. In *NAACL*.
- Yanan Wu, Zhiyuan Zeng, Keqing He, Yutao Mou, Pei Wang, and Weiran Xu. 2022b. Distribution calibration for out-of-domain detection with bayesian approximation. In *COLING*.

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.

Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Zijun Liu, Yanan Wu, Hong Xu, Huixing Jiang, and Weiran Xu. 2021a. [Modeling discriminative representations for out-of-domain detection with supervised contrastive learning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 870–878, Online. Association for Computational Linguistics.

Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Hong Xu, and Weiran Xu. 2021b. Adversarial self-supervised learning for out-of-domain detection. In *NAACL*.

Shuangfei Zhai, Yu Cheng, W. Lu, and Zhongfei Zhang. 2016. Deep structured energy based models for anomaly detection. In *ICML*.

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 Baseline Details

We perform main experiments based on two different settings, unsupervised OOD and supervised OOD detection. For unsupervised OOD detection, we compare our proposed energy detection algorithm with other methods, MSP (Maximum Softmax Probability) (Hendrycks and Gimpel, 2017), LOF (Local Outlier Factor) (Lin and Xu, 2019), GDA (Gaussian Discriminant Analysis) (Xu et al., 2020). For supervised OOD detection, we also compare our proposed energy-based margin objective with entropy (Zheng et al., 2020) and N+1 (Fei and Liu, 2016; Larson et al., 2019a). Note that margin and entropy objectives are used in the training stage, we still need detection algorithms MSP, GDA or Energy to detect in the test stage. We supplement the relevant baseline details as follows: **MSP** (Maximum Softmax Probability) (Hendrycks and Gimpel, 2017) uses maximum softmax probability as the confidence score and regards an intent as OOD if the score is below a fixed threshold.

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 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 method till now, our proposed energy score still significantly outperforms GDA.

Note that LOF and GDA both require additional post-processing modules to estimate density or distance, which induces expensive computation. We conduct a performance comparison for inference time in Table 3. Since SCL only adds a pre-training loss along with CE and also uses GDA for detection, the inference time is equal to GDA.

Detect Method	Inference time
MSP	1.00x
Energy (Ours)	1.00x
GDA/SCL	30.63x
LOF	30.89x

Table 3: Inference time comparison between different methods.

SCL (Zeng et al., 2021a) uses 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. Note that SCL still needs a confidence score function. To keep fair comparison, we follow the original paper using GDA detection method.

N+1 (Fei and Liu (2016); Larson et al. (2019a)) is an N+1 classification model which simply considers OOD samples as a new class.

Entropy (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:

$$\mathcal{L} = \mathbb{E}_{(x_{ood}) \sim \mathcal{D}}[-H(p_{\theta}(y|x_{ood}))] \quad (5)$$

where H is the Shannon entropy of the predicted distribution. $p_{\theta}(y|x_{ood})$ is the predicted distribution of the input OOD utterance x_{ood} .

Bound (Liu et al. (2020b)) uses a regularization loss defined in terms of energy to further widen the energy gap:

$$\mathcal{L} = \mathbb{E}_{(\mathbf{x}_{\text{ind}}) \sim \mathcal{D}} \max(0, E(\mathbf{x}_{\text{ind}}) - m_{\text{ind}})^2 + \mathbb{E}_{(\mathbf{x}_{\text{ood}}) \sim \mathcal{D}} \max(0, m_{\text{ood}} - E(\mathbf{x}_{\text{ood}}))^2 \quad (6)$$

where E is the energy score of IND or OOD samples in the train set. This learning objective using two squared hinge loss with two hyper-parameters m_{ind} and m_{ood} . Note that **Bond** aims at OOD image classification and relies on two independent energy bounds. Instead, our proposed Margin constructs a contrastive energy margin between IND intents and OOD intents to better disentangle energy distributions.

B Implementation Details

We use the public pre-trained 300 dimensions GloVe embeddings (Pennington et al., 2014)⁴ to embed tokens. We use a two-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 the learning rate to 1E-03. In the training stage, we use standard cross-entropy loss for unsupervised OOD and cross-entropy+energy-guided margin loss for supervised OOD. Besides, in supervised OOD scenario, we employ restriction-oriented random sampling. Specifically, we guarantee that IND and OOD samples are both included in each batch to facilitate calculation of margin loss. We both set the training epoch up to 200 with a early stop of patience 15. For our proposed energy-guided margin loss, we set the margin m to 19.0 and the temperature T to 0.8. We use the best OOD F1 scores on the validation set to calculate the threshold adaptively. Each result of the experiments is tested 5 times under the same setting and gets the average value. The training stage of our models lasts about 2 minutes for unsupervised OOD and 4 minutes for supervised OOD both on a single Tesla T4 GPU (16 GB of memory). The average value of the trainable model parameters is 3.05M. We will release our code after blind review.

C A Theoretical Proof of Energy Score vs Softmax Score

In this section, we give a theoretical proof of why energy score outperforms softmax score. Supposing we get the output logits from the intent classi-

fier, we represent MSP as follows:

$$\begin{aligned} \log \text{MSP}(\text{logits}) &= \log \max \text{softmax}(\text{logits}) \\ &= \log \max \frac{\exp(\text{logits}_i)}{\sum_i \exp(\text{logits}_i)} \\ &= \log \frac{\exp \max(\text{logits})}{\sum_i \exp(\text{logits}_i)} \\ &= \max(\text{logits}) - \log \text{sum} \exp(\text{logits}) \end{aligned} \quad (7)$$

where logits_i represents the i -th value in the vector logits . Recap the energy definition:

$$E(\mathbf{x}; f) = -T \cdot \log \sum_i^K e^{f_i(\mathbf{x})/T} \quad (8)$$

Here we set T to 1. Therefore, we get the following equation:

$$\log \text{MSP}(\text{logits}) = \underbrace{\max(\text{logits})}_{\text{regularization item}} + \text{Energy}(\text{logits}) \quad (9)$$

If the output logits get a high max value, then $\max(\text{logits})$ performs as a regularization item to avoid energy score increasing. Therefore, energy score can better mitigate the overconfidence issue than softmax score.

⁴<https://github.com/stanfordnlp/GloVe>