Probing Out-of-Distribution Robustness of Language Models with Parameter-Efficient Transfer Learning

Hyunsoo Cho[†], **Choonghyun Park**[†], **Junyeop Kim**[†], **Hyuhng Joon Kim**[†], **Kang Min Yoo**[†‡], **Sang-goo Lee**[†]
[†] Seoul National University, [‡] NAVER
{johyunsoo,pch330,juny116,heyjoonkim,sglee}@europa.snu.ac.kr
{kangmin.yoo}@navercorp.com

Abstract

As the size of the pre-trained language model (PLM) continues to increase, numerous parameter-efficient transfer learning methods have been proposed recently to compensate for the tremendous cost of fine-tuning. Despite the impressive results achieved by large pre-trained language models (PLMs) and various parameter-efficient transfer learning (PETL) methods on sundry benchmarks, it remains unclear if they can handle inputs that have been distributionally shifted effectively. In this study, we systematically explore how the ability to detect out-of-distribution (OOD) changes as the size of the PLM grows or the transfer methods are altered. Specifically, we evaluated various PETL techniques, including fine-tuning, Adapter, LoRA, and prefix-tuning, on three different intention classification tasks, each utilizing various language models with different scales.

1 Introduction

Pre-trained language models (PLM), which are pre-trained on large-scale corpora using transformer-based architectures (Vaswani et al., 2017), have achieved groundbreaking success on sundry benchmarks (Wang et al., 2019b; Rajpurkar et al., 2016; Wang et al., 2019a), establishing themselves as the standard neural model in countless applications. Moreover, language models pre-trained with larger parameters on a rich volume of corpora tend to exhibit more intriguing potentials, such as the ability to capture world knowledge (Petroni et al., 2019), generate codes (Poesia et al., 2022), and even solve mathematical problems (Henighan et al., 2020), on top of understanding linguistic knowledge (e.g., semantic or syntactic). To explore the apex of pre-trained language models (PLMs), the size of PLMs is growing exponentially and has reached billions to a trillion (Brown et al., 2020; Chowdhery et al., 2022; Fedus et al., 2022; Hoffmann et al., 2022).

Under these circumstances, the conventional method for transferring PLMs to a target task (i.e., fine-tuning) is now infeasible as it entails prohibitive costs to train and store the entire parameters of large PLMs for every desired task. To mitigate this issue, several recent parameter-efficient transfer learning (PETL) methods have been proposed to improve task scalability. For instance, adapter-based (Houlsby et al., 2019; Hu et al., 2022) approaches insert small neural modules into each layer of the PLM and update those lightweight modules in the training phase. Inspired by the recent success of textual prompts (Brown et al., 2020), prompt-based methods (Li and Liang, 2021; Lester et al., 2021; Shin et al., 2020) concatenate extra tunable tokens to the front of the input or hidden layers and update prepended soft prompts in the training phase.

Despite these breakthroughs in NLP, even very recent anomaly detection studies (Cho et al., 2022; Shen et al., 2021) are still limited to relatively small bi-directional PLMs (e.g., BERT, RoBERTa). Thus, *how large-scale PLMs or auto-regressive PLMs cope with outliers* is uncharted territory, naturally begging the following questions:

- **Q1**: Does increasing model size improve OOD detection performance without model parameters?
- **Q2**: If so, does scaling the size of PLM makes the model robust enough to utilize them without any additional process?
- **Q3**: Do fine-tuning and various PETL methodologies display differences in OOD detection performance according to the size of PLMs?
- **Q4**: Can the OOD detection methods from previous works (usually for the bi-directional PLMs) be transferred to auto-regressive PLMs (GPT)?

To resolve these questions, this paper investigates the capability of large PLMs as outlier detectors from various perspectives. Specifically, we compare the robustness to outliers with various transfer learning techniques on several OOD bench-

*Proceedings of the The 12th Joint Conference on Lexical and Computational Semantics (*SEM 2023)*, pages 225–235
July 13-14, 2023 ©2023 Association for Computational Linguistics

marks: Full fine-tuning, LoRA (Hu et al., 2022), Adapter (Houlsby et al., 2019), and prefix-tuning (Li and Liang, 2021) on various auto-regressive PLMs with different sizes, i.e., GPT2-S, M, L, XL (Radford et al., 2019), GPT-Neo (Black et al., 2021) and GPT-J (Wang and Komatsuzaki, 2021). From in-depth investigations, we share several intriguing observations: (1) As the size of the PLM increases, the performance improves without any update of model parameters. However, it is still challenging to use it without supervision since their performances still lag far behind compared to the fine-tuned small PLM (i.e., BERT-base). (2) PETLs outperform fine-tuning with sufficiently large PLMs in both IND and OOD metrics. (3) Lastly, leveraging the information of the last hidden representation, which is the most prevailing method for bi-directional PLM in recent OOD detection, does not transfer well in auto-regressive PLM, requiring a novel representation extracting technique. We believe that these findings will help future anomaly detection studies.

2 Probing OOD Robustness

2.1 Backbones and Models

To investigate the trend of OOD performance under varying scales of PLM, we consider three factors during backbone selection. They should be (1) publicly available, (2) reasonably large, and (3) share identical structures to eliminate factors other than size. Since recent large PLMs utilize auto-regressive objectives due to their computational complexity, we adopt six auto-regressive PLMs as the backbone of our experiments accordingly: **GPT2 (S,M,L,XL)**, **GPT-Neo**, and **GPT-J**.

For the parameter-efficient transfer methods, we selected two methods: two adapter-based and one prompt engineering-based. Namely, **Adapter** (Houlsby et al., 2019), **LoRA** (Hu et al., 2022), and **Prefix-tuning** (Li and Liang, 2021) are selected for the adapter approach, which is compatible with classification tasks, for the prompt approach. We also report the performance of linear evaluation, i.e., single layer perceptron (SLP) on top of PLMs, and fine-tuning, which act like a lower-bound and upper-bound, respectively.

2.2 Dataset and Metrics

Dataset. We evaluate our model on two datasets, CLINC150 and Banking77, widely used in OOD detection. CLINC150 dataset (Larson et al., 2019) contains 150 class labels (15 intents for 10 domains), while Banking77 dataset (Casanueva et al., 2020) consists of fine-grained 77 bank-related intents. Following the experimental settings from previous works (Cho et al., 2022; Zhang et al., 2022; Shu et al., 2017; Fei and Liu, 2016; Lin and Xu, 2019), we validate our models in two different scenarios: far-OOD setting and close-OOD setting. For CLINC dataset, we train our model with the whole training dataset and test with an independent OOD test split from CLINC dataset, which does not overlap with 150 classes in the training dataset. Outliers in CLINC OOD split are distributionally far from the training distribution (Zhang et al., 2022), so it is relatively easy to discern. For Banking77, we partition the dataset into 2 disjoint datasets (i.e., IND / OOD dataset) based on the class label. Since both IND and OOD datasets originated from the equivalent dataset, they share similar distributions and properties, making the task more demanding. Thus, we refer to a CLINC OOD setting as far-OOD and split settings in Banking as close-OOD settings, respectively.

Metrics. To evaluate IND performance, we measured the classification accuracy. And for OOD performance, we adopt two metrics commonly used in recent OOD detection literature:

• **FPR@95.** The false-positive rate at the true-positive rate of 95% (FPR@95) measures the probability of classifying OOD input as IND input when the true-positive rate is 95%.

• **AUROC.** The area under the receiver operating characteristic curve (AUROC) is a threshold-free metric that indicates the ability of the model to discriminate outliers from IND samples.

2.3 OOD Evaluation Methods

Evaluation in OOD detection is done via a scoring function, which outputs the appropriateness of the input into a single scalar value (p). Then we compare p with the pre-set threshold δ to determine whether the input is an outlier or not:

$$I_\delta(\boldsymbol{x}) = \begin{cases} \text{IND} & p(\boldsymbol{x}) \geq \delta \\ \text{OOD} & p(\boldsymbol{x}) < \delta, \end{cases} \quad (1)$$

In this paper, we evaluate the performance of our method in 4 different evaluation methods, which can be categorized into 2 higher branches: representation-based and logit-based.

Logit-based approaches exploit the PLM's prediction result extracted from the classification layer as

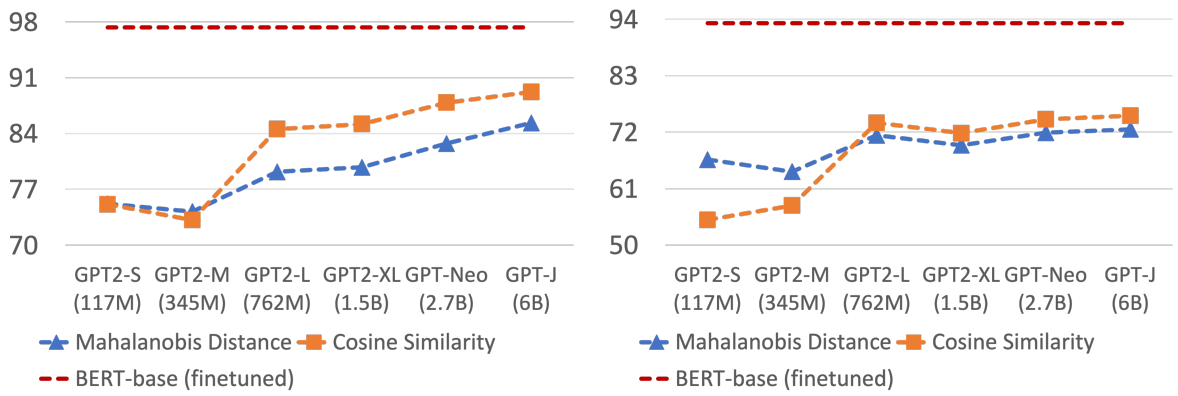

(a) Performance on far-OOD setting.

(b) Performance on close-OOD setting.

Figure 1: OOD detection performance of PLMs without updating the model parameters.

their primary information to discern outliers. Logit-based approaches are simple and have their own dominance in computational cost since it pursues OOD detection and general classification nigh simultaneously.

• **MSP** is a baseline method in this branch that employs the maximum softmax probability to score the appropriateness of the given input, based on the idea that the model will output more certain output (higher probability) to a normal sample (Hendrycks and Gimpel, 2017):

$$p(\boldsymbol{x}) = \frac{e^{f_i(\boldsymbol{x})}}{\Sigma_{j=1}^{N} e^{f_j(\boldsymbol{x})}}, \tag{2}$$

where $f_i(\boldsymbol{x})$ refer to as max value from the classification layer (max logit value).

• **Energy** is a variant of MSP, which calibrates logit value based on energy function (Liu et al., 2020):

$$p(\boldsymbol{x}) = -E(\boldsymbol{x}; f) = T \cdot \log \Sigma_i^N e^{f(\boldsymbol{x})/T}. \tag{3}$$

Representation-based approaches, on the other hand, employ the hidden representation from PLM as their primary source. Since the size of the hidden representation is larger and inheres more copious information, they generally yield a more precise decision than logit-based approaches. However, they require more inference time to derive a final score. We employed Mahalanobis distance-based and cosine similarity-based methods in this branch.

• **Mahalanobis distance** refers to the distance between the specific distribution and the input. In OOD detection, we estimate the gaussian distribution of the training dataset and utilize the minimum Mahalanobis distance to score the input suitability (Lee et al., 2018):

$$p(\boldsymbol{x}) = (\boldsymbol{h} - \boldsymbol{\mu}_k)^\top \Sigma^{-1} (\boldsymbol{h} - \boldsymbol{\mu}_k), \tag{4}$$

where training distribution is ($\mathcal{N}(\boldsymbol{\mu}_i, \Sigma)$ for $i \in i = \{1, 2, \cdots, |C|\}$), and k refers to a index of minimum mahalanobis distance.

• **Cosine Similarity** method utilizes the cosine distance between the representation of the given input ($z(x)$) and the nearest neighbor $z(x_{nn})$ (Tack et al., 2020):

$$p(\boldsymbol{x}) = \text{sim}(\boldsymbol{z}(\boldsymbol{x}), \boldsymbol{z}(\boldsymbol{x}_{nn})) \tag{5}$$

3 Analysis

In this section, we share several intriguing findings and insights from various settings.

3.1 OOD Robustness of PLMs without Supervision.

In this experiment, we investigate the OOD detection capability of PLMs without parameter tuning. Precisely, we extract the final layer representation from each frozen PLM and evaluate their performance via representation-based evaluation methods. (Logit-based evaluation methods are not used as they require additional training of the classification layer.) Figure 1 summarizes the results in two scenarios (i.e., far-OOD and close-OOD). We verified the correlation between the size of PLMs and their OOD detection ability, but utilizing them without parameter supervision is roughly impossible since they still lag far behind the small supervised methods (i.e., BERT-base with Mahalanobis evaluation) in a barebone setting. Moreover, performance improvement from the scaling saturates in a more harsh setting (i.e., close-OOD), displaying an unbridgeable gap with the fine-tuned model.

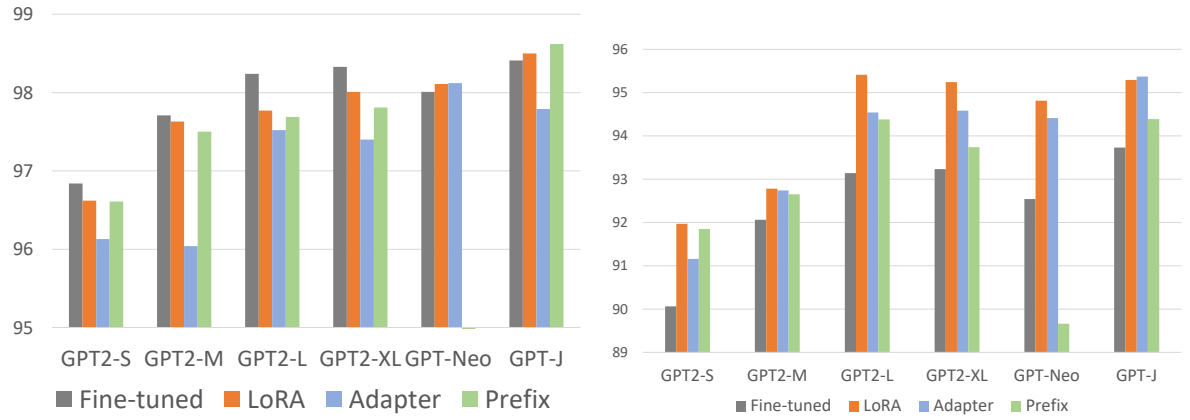

(a) Performance on far-OOD setting.

(b) Performance on close-OOD setting.

Figure 2: OOD detection performance of PLMs without updating the model parameters.

Setting	Backbone	Evaluation Method			
		MSP	Energy	Mahal.	Cosine
CLINC Setting	GPT2-S	93.22	**95.79**	77.63	76.34
	GPT2-M	95.41	**97.63**	82.42	79.82
	GPT2-L	96.21	**97.77**	96.93	97.57
	GPT2-XL	96.48	**97.99**	97.28	97.66
	GPT-Neo	96.04	**97.72**	96.59	97.64
	GPT-J	97.34	**98.50**	97.91	98.20
Banking Split 25%	GPT2-S	90.12	**91.32**	75.32	73.11
	GPT2-M	91.74	**92.78**	78.03	76.56
	GPT2-L	93.02	**93.45**	92.44	93.41
	GPT2-XL	94.29	**94.95**	93.24	94.10
	GPT-Neo	93.83	**94.85**	92.79	93.88
	GPT-J	94.11	**95.10**	93.66	94.80

Table 1: AUROC of each PLMs trained with LoRA. Energy function consistently outperforms other methods .

3.2 Evaluation methods for auto-regressive PLMs.

Many recent OOD works (Zhou et al., 2021; Shen et al., 2021) leverage hidden representation-based evaluation, as they generally surpass logit-based evaluations (Podolskiy et al., 2021). The reasonable conjecture behind their success is that hidden representations have more copious information than the logit value. However, in auto-regressive PLMs, logit-based evaluations (i.e., MSP and Energy) outperform representation-based methods (i.e., Mahalanobis distance and cosine similarity), as shown in Table 1. The reasonable conjecture for this phenomenon is due to the characteristic of the language model. Unlike bi-directional models (e.g., BERT, RoBERTa, DeBERTa), decoder models (e.g., GPT and its variants) do not have [CLS] embedding, which assembles the token embeddings to capture holistic information (Devlin et al., 2019; Kim et al., 2021). Therefore, auto-regressive PLMs generally utilize the last token embedding as a final feature embedding replacing [CLS] embedding of encoder-based models. While the last token of GPT is befitted for predicting the next token, however, it cannot extract the holistic semantics of the sentence suitably, unlike [CLS] embedding. We believe extracting a better representation through various pooling (Wang and Kuo, 2020) methods might be a possible avenue for auto-regressive models to improve the OOD robustness further.

3.3 PETLs VS. Fine-tuning

In this experiment, we investigate the performance gap between various PETL methods (i.e., Adapter, LoRA, prefix-tuning) and model fine-tuning. To compare the performance of each method under similar circumstances, we set every PETL method to utilize a similar number of parameters sufficient enough to reach maximum accuracy. Moreover, we utilized the energy function to evaluate each method as they displayed the best performance among other evaluation methods, i.e., cosine, Mahalanobis, and MSP, in the previous experiments. Table 2 summarizes the results.

From this experiment, we observed that PETL methods are more robust than fine-tuning with reasonably large PLMs (i.e., GPT-J). Specifically, most PELT methods on GPT-J outperform fine-tuning with proper tunable parameters. Nevertheless, size is not the ultimate answer. While it is clear that the scale of a model is an essential factor in OOD robustness, larger models are still vulnerable to close-OOD inputs. The capability to detect

Setting	Method	# Params.	Backbone					
			GPT2 (S)	GPT2 (M)	GPT2 (L)	GPT2 (XL)	GPT Neo	GPT-J
CLINC (far-ood)	Linear (SLP)	0%	83.03	87.39	88.47	89.55	89.44	91.94
	Fine-tuning	100%	96.84	97.71	98.24	98.33	98.01	98.41
	LoRA	0.1%	95.00	96.54	97.66	97.72	98.14	97.79
		0.5%	96.41	96.04	97.52	97.45	98.12	97.89
		1%	96.13	95.89	97.61	97.40	98.11	98.50
	Adapter	0.1%	96.62	97.52	97.74	97.71	97.81	96.80
		0.5%	95.64	97.07	97.86	96.94	97.98	98.37
		1%	95.79	97.63	97.77	97.99	98.12	98.50
	Prefix	0.1%	95.53	96.93	96.38	97.88	90.25	98.55
		0.5%	96.91	96.96	97.78	97.88	89.81	97.92
		1%	96.97	97.50	97.69	97.81	88.98	98.62
Banking split 25% (close-ood)	Linear (SLP)	0%	72.97	75.17	80.46	77.59	86.55	89.12
	Fine-tuning	100%	90.06	92.06	93.14	93.23	92.54	93.73
	LoRA	0.1%	91.18	91.74	94.65	94.58	94.29	95.82
		0.5%	91.16	92.98	94.54	94.04	94.55	94.65
		1%	91.39	92.39	93.45	93.59	94.81	95.29
	Adapter	0.1%	91.97	93.24	94.90	94.69	93.26	95.59
		0.5%	92.90	92.63	95.18	95.24	93.61	95.83
		1%	91.32	92.78	95.41	94.95	94.41	95.37
	Prefix	0.1%	91.22	91.92	93.96	93.48	81.9	94.93
		0.5%	91.85	92.55	93.84	93.34	80.82	93.99
		1%	92.09	92.65	94.38	93.74	89.66	94.39

Table 2: AUROC of various PETL methods with various number of parameters evaluated by the energy function.

far-OOD inputs (far from the training distribution) improves proportionally as the size grows, while the ability to identify close-OOD input improves rather trivially. PLM's vulnerability to close-OOD has already been reported in other studies (Zhang et al., 2022), and this may be related to shortcut learning (Geirhos et al., 2020) that predicts with high probability by looking at specific words. Generating OOD data with particular keywords or utilizing another pretext task, such as (Moon et al., 2021), can be worthy approaches to alleviate such phenomena. A suitable OOD approach is necessary to alleviate the aforementioned issue, as it can further boost the robustness. We conduct additional experiments with PETLs on three different numbers of tunable parameters: 0.1%, 0.5%, and 1% of the PLM parameters. Figure 2 summarizes the results. With sufficient parameters to reach maximum performance, there is no meaningful difference or improvement within each methodology. Also, empirically, we confirmed that LoRA is the most stable during learning and that prefix-tuning fluctuates severely according to learning.

4 Conclusion and Future Work

In this study, we showed that the scale of the language model is an important factor in OOD robustness. Moreover, we also showed that various methodologies outperform fine-tuning when applied to sufficiently large PLM. Our follow-up work seeks to create a methodology that allows large PLMs to be more robust to OOD input. The performance improvement that can be achieved by the size of PLM and OOD technique is orthogonal. In line with the growing size of PLM, the OOD technique needs to be developed in a more parameter-efficient way. As such, developing a proper OOD technique compatible with the parameter-efficient transfer methods is our proper goal.

References

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.

Sid Black, Leo Gao, Phil Wang, Connor Leahy, and Stella Biderman. 2021. GPT-Neo: Large

Scale Autoregressive Language Modeling with Mesh-Tensorflow.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. Efficient intent detection with dual sentence encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*.

Hyunsoo Cho, Choonghyun Park, Jaewook Kang, Kang Min Yoo, Tae-uk Kim, and Sang-goo Lee. 2022. Enhancing out-of-distribution detection in natural language understanding via implicit layer ensemble. In *Findings of the Association for Computational Linguistics: EMNLP*.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics, NAACL*.

William Fedus, Barret Zoph, and Noam Shazeer. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39.

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, NAACL*.

Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. Making pre-trained language models better few-shot learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics, ACL*.

Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. 2020. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673.

Dan Hendrycks and Kevin Gimpel. 2017. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *5th International Conference on Learning Representations, ICLR*.

Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo Jun, Tom B Brown, Prafulla Dhariwal, Scott Gray, et al. 2020. Scaling laws for autoregressive generative modeling. *arXiv preprint arXiv:2010.14701*.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR*.

Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. How can we know what language models know. *Transactions of the Association for Computational Linguistics*.

Taeuk Kim, Kang Min Yoo, and Sang-goo Lee. 2021. Self-guided contrastive learning for BERT sentence representations. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics ACL*.

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 EMNLP*.

Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. 2018. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Processing Systems*, pages 7167–7177.

Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP*.

Tian Li, Xiang Chen, Shanghang Zhang, Zhen Dong, and Kurt Keutzer. 2021. Cross-domain sentiment classification with contrastive learning and mutual information maximization. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8203–8207. IEEE.

Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics, ACL*.

Ting-En Lin and Hua Xu. 2019. Deep unknown intent detection with margin loss. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL*.

Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. 2020. Energy-based out-of-distribution detection. *Advances in Neural Information Processing Systems*, 33:21464–21475.

Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021. Gpt understands, too. *arXiv preprint arXiv:2103.10385*.

Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR*.

Seung Jun Moon, Sangwoo Mo, Kimin Lee, Jaeho Lee, and Jinwoo Shin. 2021. MASKER: masked keyword regularization for reliable text classification. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021*.

Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. 2019. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing, EMNLP*.

Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. Adapterfusion: Non-destructive task composition for transfer learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics, EACL*.

Alexander Podolskiy, Dmitry Lipin, Andrey Bout, Ekaterina Artemova, and Irina Piontkovskaya. 2021. Revisiting mahalanobis distance for transformer-based out-of-domain detection. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021*.

Gabriel Poesia, Alex Polozov, Vu Le, Ashish Tiwari, Gustavo Soares, Christopher Meek, and Sumit Gulwani. 2022. Synchromesh: Reliable code generation from pre-trained language models. In *The Tenth International Conference on Learning Representations, ICLR*.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*.

Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–16.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100, 000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP*.

Timo Schick and Hinrich Schütze. 2021. Exploiting cloze-questions for few-shot text classification and natural language inference. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics, EACL*.

Yilin Shen, Yen-Chang Hsu, Avik Ray, and Hongxia Jin. 2021. Enhancing the generalization for intent classification and out-of-domain detection in SLU. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics ACL*.

Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP*.

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, EMNLP*.

Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. 2020. CSI: novelty detection via contrastive learning on distributionally shifted instances. In *Advances in Neural Information Processing Systems 33, NeurIPS 2020*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*.

Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019a. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019b. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations, ICLR*.

Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. https://github.com/kingoflolz/mesh-transformer-jax.

Bin Wang and C-C Jay Kuo. 2020. Sbert-wk: A sentence embedding method by dissecting bert-based word models. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2146–2157.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Zijun Liu, Yanan Wu, Hong Xu, Huixing Jiang, and Weiran Xu. 2021. 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, ACL*.

Jianguo Zhang, Kazuma Hashimoto, Yao Wan, Zhiwei Liu, Ye Liu, Caiming Xiong, and Philip Yu. 2022. Are pre-trained transformers robust in intent classification? a missing ingredient in evaluation of out-of-scope intent detection. In *Proceedings of the 4th Workshop on NLP for Conversational AI*.

Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *Proceedings of the 38th International Conference on Machine Learning, ICßML*.

Wenxuan Zhou, Fangyu Liu, and Muhao Chen. 2021. Contrastive out-of-distribution detection for pre-trained transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP*.

Appendix

A Related Work

Parameter-Efficient Transfer Learning is drawing considerable attention lately, emerging as an alternative strategy to fine-tuning. Compared to fine-tuning, parameter-efficient transfer methods show superiority in the number of trainable parameter usage while achieving performance analogous to fine-tuning. Depending on the characteristics of the methods, parameter-efficient transfer methods can be categorized into *Adapter-based* and *Prompt-Engineering* approaches.

Adapter (Houlsby et al., 2019; Pfeiffer et al., 2021) refers to a lightweight neural module injected within each layer of PLM. The structure of the adapter generally consists of a bottleneck layer (down-projection and up-projection), a nonlinear function, a normalization layer, and a residual connection. The adapter has many different variants due to numerous design choices, such as the order or specifics of each component (e.g., which normalization technique will be used) and where the adapter will be attached. For example, LoRA (Hu et al., 2022) inserts low-rank decomposition matrices in each weight in self-attention (Vaswani et al., 2017) (i.e., query, key, and value).

Another line of work, *prompt engineering*, casts the existing task as a text generation problem to fully leverage the capability of PLMs to predict the appropriate word in the given sentence. This approach requires an empirical endeavor of optimizing the prompt to maximize a PLM's performance. Earlier works exploit handcrafted manual prompts (Schick and Schütze, 2021; Jiang et al., 2020) or by providing demonstrations to PLM [1] (Brown et al., 2020; Raffel et al., 2020; Gao et al., 2021; Zhao et al., 2021). More recent work replaces the manual prompt with a soft prompt (Li and Liang, 2021; Lester et al., 2021; Shin et al., 2020; Liu et al., 2021), a machine trainable continuous vector. The soft prompt is a more modular and versatile method that evades additional latency in the inference phase because it detaches the additionally trained parameters and solely employs the final output of the trained parameters as the prompt.

While former parameter-efficient transfer methods showed noticeable achievements, their evaluations generally assume the train and test distributions are identical (i.e., i.i.d. assumption); however, this condition is rarely satisfied in real-world scenarios due to the diversity and volatility of user input. Consequently, if the model can not correctly handle distribution-shifted malicious input and misconceives it as an in-distribution (IND) example, it may lead to fatal accidents.

Despite its practical importance, how large PLMs or parameter-efficient transfer learning cope with unknown input is poorly understood. This work aims to understand language models' capabilities to detect outliers through parameter-efficient transfer learning methods.

B Parameter-Efficient Transfer Learning

Adapter The adapter approach inserts small trainable adapter modules between transformer layers while the parameters of the original network remain fixed. The adapter module uses a bottleneck architecture which projects the input dimension h to a lower-dimensional space specified by bottleneck dimension r, followed by a nonlinear activation function, and a up-projection to initial dimension h. In this work, we attach adapter modules in two places, i.e., after the projection following multi-head attention and after the two feed-forward layers, following original implementation in (Houlsby et al., 2019). Also, we use relu as a nonlinear function and layer normalization (Ba et al., 2016).

LoRA LoRA injects trainable low-rank matrices into transformer layers to approximate the weight updates. For a pre-trained weight matrix $W \in \mathbb{R}^{h\times k}$, LoRA decompose $\Delta W = W_{down}W_{up}$ where $W_{down} \in \mathbb{R}^{h\times r}, W_{up} \in \mathbb{R}^{r\times k}$ are trainable parameters. Specifically we attach LoRA in weight matrices in the self attention module. Specifically we attached LoRA to query and key vector following the original implementation.

Prefix-Tuning Prefix tuning prepends l tunable prefix vectors to the keys and values of the multi-head attention at every layer. Following the original implementation, we reparametrize the prefix matrix of dimension h by a smaller matrix of dimension r composed with a large feedforward neural network with tanh as a nonlinear function.

C Expanded Configuration Details

C.1 Common Environment

For the experiments, 4 Tesla V100 SXM2 32GB GPUs are used. The batch size is 8 per GPU. When the GPU is too small for the batch size, we set

[1] also termed as in-context learning.

dataset	#domain	#intent	#data (train/val/test/ood)
CLINC	10	15	15000/3000/4500/1000
Banking	1	77	7812 / 1520 / 3040

Table 3: Dataset statistics.

BERT-base	CLINC150 Full		
	ACC ↑	FPR-95 ↓	AUROC ↑
Shu et al. (2017)	94.51±0.45	23.33±1.27	95.92±0.05
Li et al. (2021)	96.1±0.37	10.6±0.26	97.72±0.03
Zeng et al. (2021)	94.19±0.28	23.4±1.97	95.75±0.2
Zhou et al. (2021)	95.79±0.13	10.7±0.95	97.6±0.11
Shen et al. (2021)	96.66	10.88	97.43
Cho et al. (2022)	**96.96**±0.39	**6.67** ±0.51	**98.27** ±0.16

Table 4: Results of each model trained on the CLINC150 dataset. The best performance in each metric is indicated in **bold**.

batch size to 4 and the number of gradient accumulation steps to 2. We implemented our model based on Transformers (Wolf et al., 2020) library by Huggingface. Additionally, we used deepspeed (Rajbhandari et al., 2020) to train models. Specifically, we used ZeRO2 with cpu offload on a 240GB RAM CPU. In this setting, fine-tuning GPT-J on CLINC150 full dataset takes about 7.1 GPU hours per epoch. We used AdamW (Loshchilov and Hutter, 2019) optimizer with epsilon 1e-6 and weight decay 0.1. Furthermore, we apply the cosine annealing scheduler. For GPT-neo, the minimum learning rate is 0. For GPT-J, the minimum learning rate is the one fifth of maximum learning rate.

C.2 Number of Trainable-Parameter

For each method, a feed-forward layer is added at the end of the model. In this section, we will calculate the number of additional trainable parameters of each training methods discussed in this paper. Biases are omitted for better readability.

Adapter Adapter method adds four feed-forward layers per transformer layer in the model. Two of them are down-projection layers, and the others are up-projection layers. When the original embedding size of the model is h, the bottleneck dimension is r, and the number of transformer layers is L, the number of the trainable parameters of these layers is calculated as $4Lhr$, excluding the bias of the added layers.

LoRA Similar to adapter, LoRA also adds feed-forward layers per transformer layer. Therefore, the number of the trainable parameters of $4Lhr$. However, the number of parameters are less than adapter if h and r is the same, since LoRA does not use bias of the feed-forward layers.

Prefix-Tuning There are two trainable elements in prefix tuning. The first one is the prefix embeddings. When the number of prefixes is l, and the embedding size is h, lh parameters are used by the prefixes. Second, the reparametrization matrix is also trained. The down-projection matrix has hr parameters, when the reduced dimension for reparametrization is r. The up-projection matrix has $2Lhr$ parameters. As a result, there are $h(2Lr + l)$ trainable parameters on prefix tuning approach.

C.3 Hyper-parameter Search

Tab 5 summarizes hyper parameters for each model.

D Selecting SOTA OOD Method.

The Tab.4 summarizes the results with recently proposed OOD approaches on BERT-base with CLINC dataset. The best performing model (Cho et al., 2022) is selected as the baseline.

Method	Parameters	Values
LoRA	Learning rate	2e-4 (GPT-Neo), 5e-5 (GPT-J)
	Bottleneck dim	8 (GPT-Neo / 0.1%), 80 (GPT-Neo / 1%), 12 (GPT-J /0.1%), 128 (GPT-J / 1%)
	Location	query, value
Adapter	Learning rate	8e-5 (GPT-Neo / 0.1%), 1e-4 (GPT-Neo / 1%), 5e-5 (GPT-J), 5e-4 (GPT-J / 0.1%), 1e-4 (GPT-J / 1%)
	Bottleneck dim	6 (GPT-Neo / 0.1%), 80 (GPT-Neo / 1%), 11 (GPT-J /0.1%), 128 (GPT-J / 1%)
	Location	after Multi-head, after Feed-forward,
Prefix-tuning	Learning rate	2E-4 (GPT-Neo), 5E-5 (GPT-J)
	Bottleneck dim	12 (GPT-Neo / 0.1%), 160 (GPT-Neo / 1%), 20 (GPT-J /0.1%), 256 (GPT-J / 1%)
	Prefix length	5, 10, 20

Table 5: Hyper-parameter search for each model.