

# How Good Are LLMs at Out-of-Distribution Detection?

Bo Liu<sup>1\*</sup>, Li-Ming Zhan<sup>1\*</sup>, Zexin Lu<sup>1\*</sup>, Yujie Feng<sup>1</sup>, Lei Xue<sup>2</sup>, Xiao-Ming Wu<sup>1†</sup>

Department of Computing, The Hong Kong Polytechnic University, Hong Kong S.A.R.<sup>1</sup>

School of Cyber Science and Technology, Sun Yat-Sen University, Shenzhen, China.<sup>2</sup>

{bokelvin.liu, lmzhan.zhan, zexin.lu, yujie.feng}@connect.polyu.edu.hk

xuelei3@mail.sysu.edu.cn

xiao-ming.wu@polyu.edu.hk

## Abstract

Out-of-distribution (OOD) detection plays a vital role in enhancing the reliability of machine learning models. As large language models (LLMs) become more prevalent, the applicability of prior research on OOD detection that utilized smaller-scale Transformers such as BERT, RoBERTa, and GPT-2 may be challenged, due to the significant differences in the scale of these models, their pre-training objectives, and the paradigms used for inference. This paper initiates a pioneering empirical investigation into the OOD detection capabilities of LLMs, focusing on the LLaMA series ranging from 7B to 65B in size. We thoroughly evaluate commonly used OOD detectors, examining their performance in both zero-shot and fine-tuning scenarios. Notably, we alter previous discriminative in-distribution fine-tuning into generative fine-tuning, aligning the pre-training objective of LLMs with downstream tasks. Our findings unveil that a simple cosine distance OOD detector demonstrates superior efficacy, outperforming other OOD detectors. We provide an intriguing explanation for this phenomenon by highlighting the isotropic nature of the embedding spaces of LLMs, which distinctly contrasts with the anisotropic property observed in smaller BERT family models. The new insight enhances our understanding of how LLMs detect OOD data, thereby enhancing their adaptability and reliability in dynamic environments. We have released the source code at <https://github.com/Awenbocc/LLM-OOD> for other researchers to reproduce our results.

**Keywords:** Out-of-distribution detection, large language models, performance evaluation

## 1. Introduction

Out-of-distribution (OOD) detection has attracted significant attention due to its crucial role in ensuring AI safety (Salehi et al., 2022). The objective is to identify and raise an alarm for inputs that exhibit distributional shifts compared to the in-distribution (ID) training data. Given that the test distribution can dynamically change over time, OOD detection has become indispensable in high-stakes applications, such as healthcare and self-driving cars. Its ability to detect anomalous inputs and adapt to evolving scenarios makes it a vital component in ensuring the reliability and robustness of AI systems in real-world, dynamic environments.

Utilizing sentence representations yielded by pre-trained language models (PLMs) to derive OOD confidence scores has been the *de facto* method for textual OOD detection. Specifically, PLMs are first fine-tuned on the ID data and then OOD detectors are applied on the sentence representations generated by PLMs. Compared to ID data, there are two types of OOD instances: far-OOD where ID and OOD data come from different domains and near-OOD where ID and OOD data come from the same domain but with different classes, as shown in Figure 1. Typically, near-OOD samples are harder to recognize. A body of works (Hendrycks et al.,

In-Distribution (ID)	Out-of-Distribution (OOD)
<b>(I) ID: Sentiment Analysis</b> <i>Nelson's intentions are good.</i> SST-2	<b>(I) Far-OOD: Question Classification</b> <i>Who was the first American in space ?</i> TREC
<b>(II) ID: Transfer-related Intent</b> <i>I would like to make a transfer.</i> CLINC150-BANK (Transfer)	<b>(II) Near-OOD: Balance-related Intent</b> <i>What is in my bank accounts?</i> CLINC150-BANK (Balance)

Figure 1: Illustration of two types of OOD instances compared to ID samples: far-OOD where ID and OOD data come from different domains and near-OOD where ID and OOD data come from the same domain but with different classes.

2020; Podolskiy et al., 2021; Uppaal et al., 2023) have shown that Transformer-based models can produce better sentence representations for OOD detection. However, these studies have mainly focused on evaluating the OOD detection performance of small-scale encoder-based Transformers, such as RoBERTa and BERT.

Recently, large language models (LLMs) have made significant strides in various cognitive tasks, yet their capabilities on OOD detection remain largely unexplored. Unlike relatively small-scale PLMs used by prior studies, LLMs often display

\* Equal contribution.

† Corresponding author.

notable differences. In particular, the majority of previously prominent PLMs utilized for OOD detection adopt the encoder-based architecture, such as BERT and RoBERTa. These models are predominantly designed with a pre-training objective that focuses on sentence classification. However, recent LLMs (Touvron et al., 2023a; Zeng et al., 2022; Du et al., 2022; Chowdhery et al., 2022; Chung et al., 2022) exclusively adopt an autoregressive training objective during pre-training. Consequently, the hidden states of LLMs are specialized for next token prediction, which could influence their performance in OOD detection. Moreover, previous works test changes in OOD detection when adapting PLMs to downstream tasks through discriminative fine-tuning, even for decoder-based models (Cho et al., 2023). However, a more intuitive approach is to probe the pre-training knowledge of LLMs through generative fine-tuning, which better aligns LLMs’ pre-training objective with downstream tasks. Thus, it is imperative to extensively investigate the OOD detection capabilities of LLMs to gain deeper insights into their potential and limitations.

This paper aims to fill this gap by offering a comprehensive and structured assessment of OOD detection with LLMs across varying scales (ranging from 7B to 65B). Notably, our evaluation process is specifically designed to consider the scaling laws of LLMs with commonly utilized OOD detection detectors, ensuring broader and more generalized findings. In summary, our analysis has revealed the following new insights:

### 1. Discriminative vs. generative fine-tuning.

We have observed that generative fine-tuning demonstrates greater resilience to the issue of ID overfitting when compared to discriminative fine-tuning. As highlighted by Uppaal et al. (2023), there exists a trade-off between achieving higher accuracy on ID tasks and ensuring effective OOD detection. It has been shown that OOD detectors progressively lose efficacy as the training of ID tasks continues. However, our findings indicate that adopting a generative approach to fine-tuning LLMs can effectively mitigate this issue, potentially resulting in stable OOD performance even as training progresses and ID accuracy improves.

### 2. LLM-based far- vs. near-OOO detection.

Our results consistently demonstrate that LLMs are natural far-OOO detectors. Remarkably, LLMs of all scales achieve near-perfect OOD performance in far-OOO scenarios without requiring any fine-tuning. However, when it comes to near-OOO detection, only the 65B model is able to achieve satisfactory performance without any fine-tuning. Despite that, we discover that fine-tuning significantly improves the near-OOO detection capability of LLMs.

### 3. Anisotropy vs. isotropy.

Our experimental results suggest that the cosine distance function, when used as a straightforward OOD detector, performs exceptionally well. This observation leads to an intriguing discovery: the embedding spaces of LLMs exhibit a desirable isotropic property, which is not possessed by the BERT family models. The sentence embeddings produced by the BERT family models have been noted to possess an undesirable characteristic of being concentrated within a narrow cone, a phenomenon referred to as anisotropic representations (Ethayarajh, 2019), which negatively affects tasks involving semantic relationships and is commonly known as representation degeneration (Gao et al., 2019). The issue is resolved through the isotropic representations generated by LLMs, which allow the cosine distance to excel in OOD detection and may potentially benefit a broad spectrum of tasks.

## 2. Related Work

### 2.1. Out-of-Distribution Detection

Out-of-Distribution (OOD) detection has a long history in machine learning and is highly related to research topics like outlier detection, anomaly detection and novelty detection (Hendrycks and Mazeika, 2022). In the task setting of OOD detection, the *in-distribution* is characterized by the labeled training dataset and the *out-of-distribution* refers to anything else that possesses distributional shifts. OOD detection distinguishes from other related topics in the case that it requires both of the ID and OOD classification accuracy (Yang et al., 2021).

### 2.2. Textual OOD Detection with PLMs

The significance of textual OOD detection in ensuring the robustness of NLP applications, such as dialogue systems, has led to a surge in research interest. Pre-trained Transformers have shown intrinsic superiority in handling OOD detection (Hendrycks et al., 2020; Zhan et al., 2021).

Several works have further evaluated the OOD performance of PLMs with respect to commonly used OOD detectors including MSP (Hendrycks and Gimpel, 2017), Mahalanobis distance (MD) (Lee et al., 2018), and Energy score (Liu et al., 2020). For example, Podolskiy et al. (2021) show that the Gaussian distribution assumption of MD better matches the representation space of BERT and can yield the best OOD performance in intent OOD detection benchmarks. Zhou et al. (2021) show that a contrastive regularizer can further improve the sentence representation of Transformers for OOD detection.

More recently, Uppaal et al. (2023) present a thorough analysis on the fine-tuning strategies

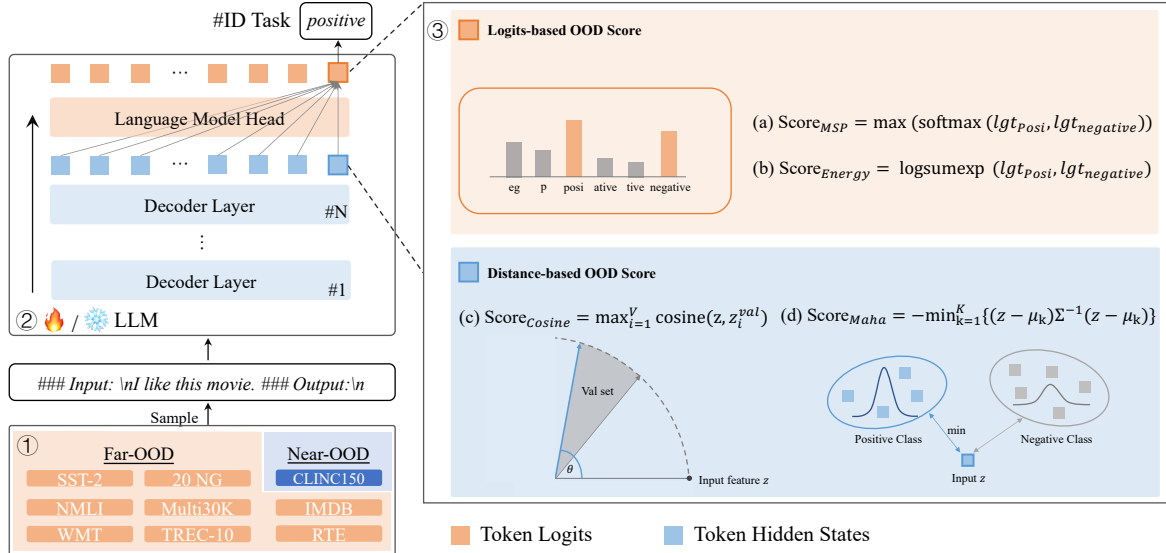


Figure 2: Our proposed evaluation framework for LLMs at OOD detection, taking three aspects into consideration: ① distribution of OOD samples (near or far), ② impact of model tuning on OOD detection, and ③ diverse OOD score functions.

for OOD detection with RoBERTa and show that RoBERTa (Liu et al., 2019) without fine-tuning can achieve near-perfect far-OOD detection performance. Similarly, we find that LLMs can also achieve perfect far-OOD detection performance without fine-tuning. Cho et al. (2023) explore the OOD detection capability of medium-sized PLMs (such as GPT-2 (Radford et al., 2019)), as well as the impact of various ID fine-tuning techniques. While they also assess decoder-based models, the models they evaluated are not extensive as this work and they neglect to undertake generative ID tuning, a crucial step to fine-tune decoder-based models for downstream ID tasks. Furthermore, the models they examine remain at relatively moderate scales, and an exploration of the possible data-efficient characteristics of the model is lacking.

Recently, large language models (LLMs) have been leading a paradigm shift in the field of natural language processing (NLP) (Touvron et al., 2023a; Wei et al., 2022; Zeng et al., 2022; Du et al., 2022; Chowdhery et al., 2022; Chung et al., 2022; Feng et al., 2023). The use of LLMs to solve NLP tasks in a generative way has become widespread. These LLMs commonly adopt the decoder-based architecture and are trained with the autoregressive objective. In this paper, we focus on the OOD performance of open-source LLMs and anticipate our work can provide useful insights for OOD detection under this paradigm.

### 3. Method

**Problem statement.** The objective of OOD detection is to effectively differentiate between instances

that belong to a specific distribution (in-distribution  $\mathcal{D}_{ID}$ ) and those falling outside of that distribution (out-of-distribution  $\mathcal{D}_{OOD}$ ). To better and fairly evaluate the capabilities of LLMs for OOD detection compared to prior smaller models (e.g., RoBERTa (Liu et al., 2019)) (Uppaal et al., 2023), we utilize the same sentence classification task as the ID training task. In practical application scenarios, undesired inputs (e.g., a severe distribution shift towards ID data) may occur, and an OOD confidence scoring function  $f_{OOD}$  can be used to reject whether outputting results for such inputs or not.

#### 3.1. ID Generative Fine-tuning with LLMs

For the ID sentence classification task, we align with the nature of LLMs and adopt a generative approach (referred to as open-ended classification) (Radford et al., 2018). Concretely, given an input sentence  $\mathbf{X}_s$ , we first expand it with a simple template: “### Input:\n $\mathbf{X}_s$  ### Output:\n”, to facilitate the extraction of outputs by identifying the section following the “Output” symbol. Subsequently, we maximize the probability of generating the target label  $\mathbf{X}_a$  with  $L$  tokens by:

$$\max p(\mathbf{X}_a|\mathbf{X}_s) = \prod_{i=1}^L p_\theta(x_i|\mathbf{X}_s, \mathbf{X}_{a,<i}), \quad (1)$$

where  $\theta$  represents the model parameters and  $\mathbf{X}_{a,<i}$  are partial label tokens that come before the current prediction token  $x_i$ .

**Parameter-efficient fine-tuning.** To improve the performance of LLMs in the in-distribution sentence

classification task, we employ a parameter-efficient fine-tuning (PEFT) approach, to minimize the usage of additional parameters. Specifically, we utilize the low-rank adaptation (LoRA) (Hu et al., 2021) technique which freezes the pre-trained LLMs’ weights and inserts trainable rank decomposition matrices into each Transformer layer. We perform PEFT with answer predictions, i.e., only class label tokens are utilized to compute the auto-regressive loss. During the test stage, we use *strict matching* to determine whether the generated labels are identical to the ground truth.

### 3.2. OOD Detection with LLMs

The overview of our OOD detection framework is illustrated in Figure 2. Our primary focus is on decoder-like LLMs, such as LLaMA (Touvron et al., 2023a), as they have demonstrated excellent performance when their model size scales up (OpenAI, 2023; Brown et al., 2020). To obtain a comprehensive observation, we conduct OOD detection experiments on two different semantic distribution settings (Ming et al., 2022; Lin et al., 2021): far-ODD and near-ODD (*cf.* Figure 1). Regarding OOD detection methods, we focus on the prevailing post-hoc paradigm (Yang et al., 2021). In the following, we elaborate on how to integrate the post-hoc OOD detectors into decoder-like LLMs, which has never been addressed in existing literature.

**Customized post-hoc methods.** According to prior studies (Hendrycks et al., 2020; Zhou et al., 2021), there mainly exist two categories of post-hoc methods: logits-based OOD score functions and distance-based ones. Since previous works used these methods for language models accompanied by a classifier, we here customize them for decoder-type LLMs with only a language model head (as shown in Figure 2) in the following:

**Logits-based OOD score functions** operate on the final class-related logits. In generative classification, the generated class name is usually composed of several tokens, e.g., “positive” consists of “posi” and “tive”. Instead of calculating the probability (logits) for the complete ID class name, we simplify the process by considering the probability assigned to the first token of its class name. For instance, in a sentiment analysis task with classes like “positive” and “negative” as depicted in Figure 2, we only need to identify the probability corresponding to the tokens “posi” and “negative” respectively. Considering that different class names may have common prefixes, such as “positive” and “position”, we will rephrase the conflicting class names at the beginning of ID training, such as replacing “position” with “location”. In practice, we observe this re-translation has no impact on the ID task. Overall,

there are mainly two logits-based functions:

- Maximum softmax probability (MSP) (Lee et al., 2018) utilizes the maximum softmax probability corresponding to each class, i.e, score  $\mathcal{S}(x) = \max\{p(y_i|x)\}_i^K$ , where  $K$  is the number of classes, and ID samples always exhibit higher probability scores while OOD ones correspond to lower scores.
- Energy score (Energy) (Liu et al., 2020; LeCun et al., 2006) computes confidence score  $\mathcal{S}(x) = \log \sum_i^K e^{(w^T \cdot z)_i}$  where  $w^T$  is the weight of the language model head and  $z$  is all word embeddings. Note that for both MSP and Energy, we only select the probability and logits corresponding to the first token of each class name, as mentioned before.

**Distance-based OOD score functions** apply to sentence representations. Prior studies using encoder-based PLMs treated the embeddings of special token <cls> as sentence representations. For LLMs, we employ the embeddings of the last token as the representation. There are mainly two functions considered for evaluation: Mahalanobis distance (Maha) (Lee et al., 2018) and Cosine distance (Cosine) (Zhou et al., 2021)<sup>1</sup>.

## 4. Experimental Setup

### 4.1. Datasets

To draw universal conclusions, we conduct a comprehensive evaluation of two kinds of dataset distribution settings (Arora et al., 2021) as illustrated in Figure 1 and Figure 2.

**Far-ODD.** In this paradigm, ID and OOD samples come from different distributions (datasets), exhibiting significant semantic differences. Following Hendrycks et al. (2020) and Zhou et al. (2021), we evaluate 8 datasets, including 20 News-groups (20NG) (Lang, 1995) for topic classification, RTE (Wang et al., 2018) and MNLI (Williams et al., 2017) for nature language inference, TREC-10 (Li and Roth, 2002) for question classification, SST-2 (Socher et al., 2013) and IMDB (Maas et al., 2011) for sentiment analysis, and the English side of Multi30K (Elliott et al., 2016) and WMT16 (Bojar et al., 2016) for machine translation. Among them, we choose 20NG and SST-2 as two separate *in-distribution* tasks and the remaining ones are recognized as *out-distribution*. Note that when SST-2 is used as the ID, we do not consider IMDB as an OOD dataset since both of them are sentiment analysis tasks.

<sup>1</sup>We refer authors to original papers for more details.

ID Dataset	LLM	Maha			Cosine			MSP			Energy			
		AUROC $\uparrow$	FAR@95 $\downarrow$	AUPR $\uparrow$	AUROC $\uparrow$	FAR@95 $\downarrow$	AUPR $\uparrow$	AUROC $\uparrow$	FAR@95 $\downarrow$	AUPR $\uparrow$	AUROC $\uparrow$	FAR@95 $\downarrow$	AUPR $\uparrow$	
Far-OOD	SST-2	LLaMA-7B	0.991	0	0.993	0.990	0.006	0.990	0.905	0.318	0.811	0.368	0.930	0.380
		LLaMA-13B	0.992	0	0.993	0.990	0.005	0.989	0.939	0.213	0.818	0.571	0.778	0.478
		LLaMA-30B	0.994	0.003	0.993	0.991	0.009	0.990	0.881	0.361	0.742	0.651	0.738	0.540
	LLaMA-65B	0.991	0.007	0.991	0.990	0.007	0.992	0.776	0.621	0.646	0.544	0.821	0.485	
	20NG	LLaMA-7B	0.997	0	0.995	0.998	0	0.996	0.441	0.929	0.391	0.571	0.784	0.417
		LLaMA-13B	0.996	0.006	0.989	0.993	0.004	0.990	0.622	0.754	0.482	0.491	0.932	0.362
LLaMA-30B		0.995	0.005	0.987	0.995	0.002	0.993	0.533	0.847	0.424	0.491	0.906	0.362	
LLaMA-65B	1	0	0.998	0.999	0	0.997	0.616	0.764	0.421	0.508	0.925	0.369		
Near-OOD	CLINC-Banking	LLaMA-7B	0.896	0.568	0.921	0.891	0.587	0.916	0.720	0.814	0.763	0.722	0.818	0.758
		LLaMA-13B	0.905	0.408	0.922	0.903	0.514	0.922	0.739	0.769	0.760	0.713	0.831	0.743
		LLaMA-30B	0.895	0.472	0.913	0.910	0.424	0.923	0.733	0.813	0.746	0.724	0.795	0.735
	LLaMA-65B	0.951	0.255	0.964	0.956	0.200	0.964	0.823	0.604	0.834	0.826	0.614	0.834	
	CLINC-Travel	LLaMA-7B	0.895	0.680	0.932	0.887	0.738	0.927	0.584	0.921	0.640	0.637	0.912	0.674
		LLaMA-13B	0.942	0.485	0.964	0.922	0.730	0.955	0.639	0.834	0.696	0.633	0.909	0.695
LLaMA-30B		0.926	0.458	0.950	0.928	0.523	0.950	0.650	0.911	0.697	0.653	0.888	0.698	
LLaMA-65B	0.959	0.182	0.971	0.976	0.076	0.986	0.739	0.745	0.753	0.755	0.681	0.768		

Table 1: OOD detection performance of zero-grad LLaMA models. We use the full validation set to calculate each OOD score. The results are averaged over five seeds.

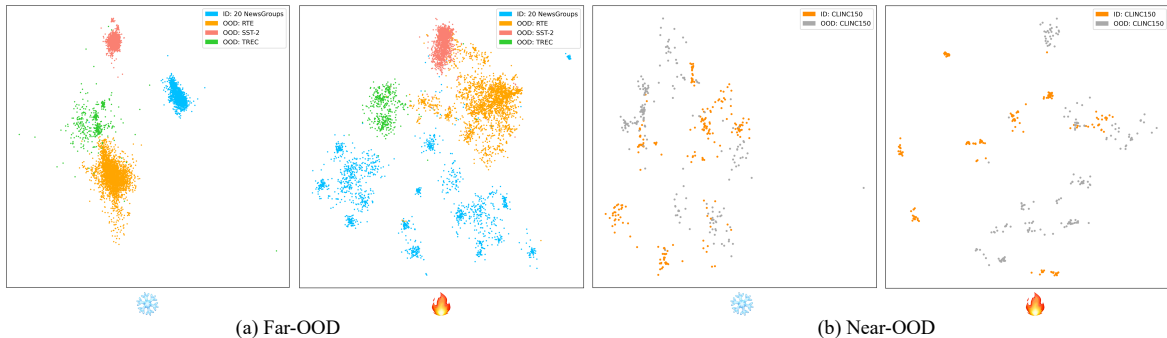


Figure 3: UMAP (McInnes et al., 2018) visualization of representations generated from the penultimate layer of the zero-grad (❄️) and fine-tuned (🔥) LLaMA-7B models. (a) Far-OOD: 20NG is treated as ID while SST-2, RTE, and TREC are treated as OOD. (b) Near-OOD: the banking domain of CLINC150 is selected, of which 50% of the classes are treated as ID, and the rest are treated as OOD.

**Near-OOD.** We also test on a more challenging scenario, where ID and OOD samples come from the same domain but with disjoint label sets. A well-researched domain is OOD intent detection (Larson et al., 2019). Specifically, we use CLINC150 dataset and choose *Banking* and *Travel* domains. Within each domain, 50% of the classes are chosen as ID, and the remaining classes as OOD.

Dataset	Full-shot	10-shot	5-shot	1-shot
SST-2	16	8	4	2
20NG	8	8	8	8
CLINC150 (Banking or Travel)	16	16	16	8

Table 2: Batch size configuration for each dataset.

## 4.2. Evaluation Metrics

We employ three commonly used metrics for OOD detection: (1) AUROC (area under the receiver operating characteristic curve). (2) FAR@95 (false alarm rate at 95% recall). It represents the probability of incorrectly classifying a negative sample as positive when the Recall or True Positive Rate (TPR) is 95%. We treat the OOD class as negative. (3) AUPR (area under the precision-recall curve). Additionally, we use accuracy as a metric for ID classification task.

## 4.3. Implementation Details

All experiments are conducted on a workstation with 4 NVIDIA A100 80G GPUs. For zero-grad

OOD detection, LLaMA-7B, -13B, -30B, and -65B are deployed on 1, 1, 2, and 4 A100 GPUs, respectively. When further fine-tuning LLMs on ID tasks, the LoRA configurations (Section 3.1) are that rank  $r$  is 16, scaling  $\alpha$  is 16, and query/key/value/output projection matrices  $\{W_q, W_k, W_v, W_o\}$  in each self-attention module need to be updated. We train the network for 50 epochs with early stop criteria that if the model’s performance on the validation set continuously drops for 6 consecutive epochs and the current epoch number exceeds 15, training will be terminated. We use AdamW optimizer with learning rate  $1 \times 10^{-4}$ , further decayed by linear schedule. Due to the varying lengths of sentences in different ID datasets, we configure different batch sizes shown in Table 2. All experiments are conducted over five seeds (1, 2, 3, 4, 5).

ID Dataset	Shot	ID ACC $\uparrow$	Maha			Cosine			MSP			Energy			
			AUROC $\uparrow$	FAR@95 $\downarrow$	AUPR $\uparrow$	AUROC $\uparrow$	FAR@95 $\downarrow$	AUPR $\uparrow$	AUROC $\uparrow$	FAR@95 $\downarrow$	AUPR $\uparrow$	AUROC $\uparrow$	FAR@95 $\downarrow$	AUPR $\uparrow$	
Far-OOD	SST-2	1	0.535	0.5	1.0	0.422	0.954	0.250	0.934	0.664	0.581	0.587	0.716	0.589	0.637
		5	0.664	0.878	0.625	0.843	0.973	0.045	0.971	0.768	0.493	0.674	0.885	0.408	0.794
		10	0.857	0.967	0.204	0.962	0.991	0.009	0.987	0.771	0.514	0.693	0.896	0.379	0.803
	20NG	Full	0.976	0.993	0.004	0.992	0.993	0.005	0.991	0.947	0.298	0.888	0.961	0.189	0.907
		1	0.463	0.5	1	0.380	0.991	0.047	0.985	0.756	0.779	0.670	0.850	0.681	0.824
		5	0.713	0.983	0.074	0.975	0.991	0.023	0.989	0.868	0.503	0.799	0.947	0.283	0.918
Near-OOD	CLINC-Banking	10	0.796	0.992	0.042	0.987	0.996	0.013	0.991	0.893	0.438	0.840	0.951	0.215	0.924
		Full	0.944	0.995	0.003	0.991	0.993	0.007	0.991	0.959	0.207	0.939	0.968	0.114	0.945
		1	0.589	0.5	1	0.533	0.905	0.510	0.926	0.846	0.696	0.860	0.870	0.658	0.897
	CLINC-Travel	5	0.882	0.863	0.614	0.879	0.962	0.255	0.968	0.873	0.556	0.878	0.903	0.463	0.916
		10	0.949	0.937	0.424	0.956	0.968	0.157	0.974	0.902	0.422	0.902	0.919	0.346	0.929
		Full	0.973	0.958	0.231	0.969	0.964	0.147	0.970	0.936	0.269	0.945	0.930	0.225	0.931
CLINC-Travel	1	0.526	0.5	1	0.533	0.910	0.481	0.925	0.767	0.756	0.771	0.780	0.733	0.793	
	5	0.964	0.897	0.644	0.925	0.974	0.148	0.983	0.886	0.415	0.886	0.875	0.420	0.872	
	10	0.984	0.975	0.137	0.983	0.982	0.078	0.988	0.930	0.3	0.931	0.933	0.231	0.933	
Full	0.991	0.980	0.045	0.988	0.978	0.049	0.987	0.942	0.121	0.933	0.948	0.112	0.953		

Table 3: The performance of the fine-tuned LLaMA-7B model for OOD detection and ID classification. “Shot” denotes the number of examples in the ID training or validation set. We report the average results of five seeds.

## 5. Findings

### 5.1. Zero-grad OOD Detection with LLMs

In this section, we evaluate the *zero-grad* OOD performance of LLMs. The objective is to examine how well OOD detection performs when utilizing the knowledge acquired by LLMs during pre-training. The results are summarized in Table 1 and all LLMs are frozen in this setting. Note that we use full-shot validation set to calculate each OOD score.

**LLMs are natural far-OOD detectors.** As shown in Table 1, when applying distance-based OOD detection methods, such as Maha and Cosine, all LLMs can achieve near-perfect results (e.g., AUROC and AUPR approach 1 while FAR@95 approaches 0). To better understand why distance-based OOD detectors are so effective, we visualize the corresponding sentence representations yielded by the penultimate layer (before the top head layer), as shown in Figure 3 (a). It can be found that representations from the same dataset are tighter, while ID and OOD sentences have clear boundaries, indicating the profound semantic discrimination prowess exhibited by LLMs. However, both MSP and Energy generate poor results. This is foreseeable, as both of them condition on the first token generated from the input sentence. When the model has not been fine-tuned, it often struggles to accurately output class names, leading to inferior OOD performance. Moreover, from the probability density of Figure 4 (a), it can be found that there is a significant overlap between ID and OOD, leading to a decrease in OOD detection performance.

**The capability of LLMs for near-OOD detection improves with their scale.** We present the zero-grad near-OOD results in Table 1 (CLINC-Banking and CLINC-Travel). For the near-OOD setting, as the number of model parameters increases, the OOD detection performance will also be improved.

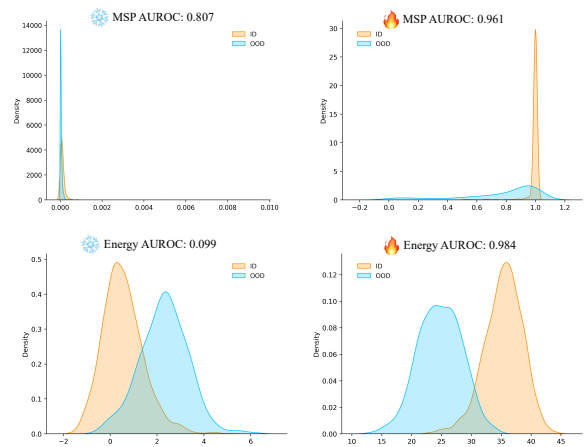


Figure 4: Impact of fine-tuning on logits-based OOD scores (MSP at the top row and Energy at the bottom row). We plot SST-2 (ID) vs. TREC-10 (OOD) for visualization.

Remarkably, when the model has an exceedingly large number of parameters (i.e., LLaMA-65B), we can observe a dramatic performance surge (Wei et al., 2022) to detect OOD inputs, especially with distance-based OOD methods. In particular, the AUROC values for Maha and Cosine both surpass 95%, and FAR95 is enhanced by **at least 30%** in comparison to the 7B model.

Furthermore, it is evident that the near-OOD performance of LLMs is notably inferior compared to their performance on far-OOD instances. To understand this, we provide a visualization for this setting as illustrated in Figure 3 (b). The embeddings of ID and OOD samples are mixed up since their labels come from the same domain (i.e., travel or banking). Consequently, detecting near-OOD instances becomes notably more challenging than far-OOD instances.

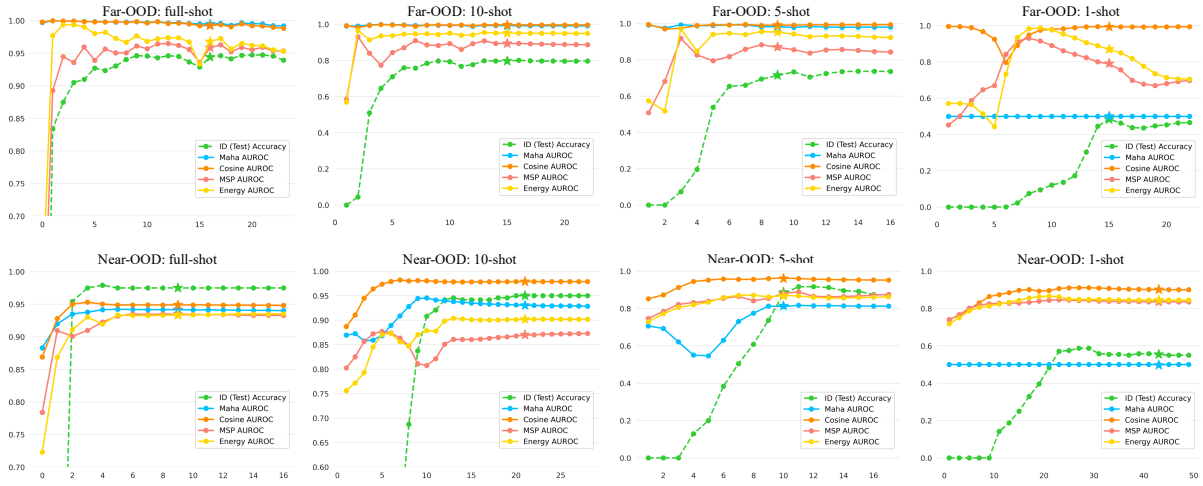


Figure 5: Performance changes for ID classification and OOD detection as training progresses with the different number of training shots. Top row: 20NG is ID training task; Bottom row: banking domain of CLINC150 is selected where 50% classes are used as ID training task and the rest are OOD samples. The star (\*) on each line indicates the selected results whose epoch corresponds to the best ID performance on the validation set.

## 5.2. OOD Detection with Generatively Fine-tuned LLMs

In this subsection, we study the influence of fine-tuning LLMs on OOD detection. Specifically, we conduct an in-depth examination of how the OOD detection performance evolves with the progression of ID task training.

**ID fine-tuning can boost OOD detection.** We fine-tune LLMs in a generative manner in both few-shot and full-shot scenarios. The results are summarized in Table 3. Likewise, we present both far- and near-OOD results comparable to the zero-grad configuration. Clearly, fine-tuning LLMs on in-distribution tasks can notably augment the models’ capacity to detect OOD instances, surpassing the performance of the zero-grad setting by a significant margin in most cases like in near-OOD setting and with logits-based functions (in both full-shot scenarios with LLaMA-7B model).

In Figure 5, we present the fine-tuning curves. It can be observed that as the ID accuracy increases, almost all OOD detectors are improved accordingly. To study how fine-tuning impacts the ID vs. OOD separability, we plot their density distributions in Figure 4. Clearly, fine-tuning can improve the separability between ID and OOD instances. A similar effect can be cross-validated in Figure 3 (b) in which the embedding of different classes within the ID becomes more compact, while the separation between ID and OOD becomes clearer. However, it is important to highlight that as the training continues, there is a possibility of encountering overfitting, which could result in inferior OOD performance, es-




Figure 6: Impact of different ID training objectives, discriminative vs. generative. SST-2 dataset with full data is used as the ID training task.

pecially for logits-based methods as illustrated in Figure 5 for both full-shot and 1-shot far-OOD scenarios. This observation is similar to the findings in (Uppaal et al., 2023).

**Generative fine-tuning generalizes better.** In addition to generative fine-tuning, we also explore discriminative fine-tuning by appending a classifier after LLMs (replacing the language model head<sup>2</sup>) to conduct ID task. The comparison of the trend charts presented in Figure 6 reveals that generative fine-tuning tends to be less overfit on the ID task and all OOD detectors consistently perform better than discriminative fine-tuning, especially for distance-based OOD detectors. To better understand this effect, based on the transformations of embeddings illustrated in Figure 3 (a), it becomes evident that throughout the generative training pro-

<sup>2</sup>We use `LlamaForSequenceClassification` provided by Huggingface (Wolf et al., 2020)

ID Dataset	PTM	Maha		Cosine		MSP		Energy	
		AUROC ↑	FAR@95 ↓	AUROC ↑	FAR@95 ↓	AUROC ↑	FAR@95 ↓	AUROC ↑	FAR@95 ↓
SST-2	RoBERTa-L <sub>†</sub> (Uppaal et al., 2023)	0.971	0.152	0.919	0.414	-	-	-	-
		<b>0.991</b>	<b>0</b>	<b>0.990</b>	<b>0.006</b>	0.905	0.318	0.368	0.930
	LLaMA-7B	<i>Fine-tuned</i>							
		0.969	0.183	0.962	0.236	0.889	0.613	0.877	0.632
	LLaMA-7B	<b>0.993</b>	<b>0.004</b>	<b>0.993</b>	<b>0.005</b>	<b>0.947</b>	<b>0.298</b>	<b>0.961</b>	<b>0.189</b>
20NG	RoBERTa-L <sub>†</sub> (Uppaal et al., 2023)	0.998	0.002	0.998	0.002	-	-	-	-
		0.997	<b>0</b>	<b>0.998</b>	<b>0</b>	0.441	0.929	0.571	0.784
	LLaMA-7B	<i>Fine-tuned</i>							
		0.983	0.073	0.978	0.107	0.946	0.305	0.965	0.158
	LLaMA-7B	<b>0.995</b>	<b>0.003</b>	<b>0.993</b>	<b>0.007</b>	<b>0.959</b>	<b>0.207</b>	<b>0.968</b>	<b>0.114</b>

Table 4: Comparison of large and small PTMs under zero-grad and fine-tuned settings for OOD detection. † denotes the results we reproduce due to different calculating methods, while ‡ indicates results cited from the original paper.

PTMs	Data Corpus		
	CLINC150	CLINC150-Banking	CLINC150-Travel
LLaMA-7B	<b>0.4731</b>	<b>0.5529</b>	<b>0.5312</b>
RoBERTa-L	0.9991	0.9992	0.9989

Table 5: Average sentence anisotropy of model’s last layer.

cess, while the ID’s distribution shifts into class-specific clusters, a distinct separation continues to exist between these clusters and the OOD samples. This preserves the effectiveness of distance-based OOD detection methods. Prior study (Uppaal et al., 2023) pointed out that discriminative tuning the small models (e.g., RoBERTa (Liu et al., 2019)) negatively impacts the performance of distance-based OOD detection methods. This issue also exists in discriminative tuning LLMs but has been solved in the generative tuning.

Besides, in Table 4, we compare encoder-based and decoder-based Transformers and observe impressive improvement on decoder-based LLMs.

**Cosine distance is a data-efficient OOD detector.** To further investigate whether LLMs possess data-efficient OOD detection capabilities, we configure the training samples of the ID as few-shot instances (e.g., 1, 5, and 10). Please note that we also set the number of validation sets to be the same shot, since all OOD detection methods rely on the validation set. Results presented in Table 3 convey that as the number of shots increases, the OOD detection capability of the LLMs also improves. Moreover, distance-based OOD detection methods are superior to logits-based ones, and they can achieve good performance even with only 10-shot samples. Particularly, cosine distance is a data-efficient OOD detector that can provide effective detection by requiring only **1-shot** instance. For example, it achieves AUROC of **99.1%** (near-perfect) on 20NG (ID) and over 90% on others. Besides, in the 1-shot setting, the Mahalanobis distance loses its efficacy since it’s unfeasible to

model the necessary Gaussian distribution when there’s only a single sample for each class.

**Isotropy vs. anisotropy.** By examining Table 1 and Table 3, it becomes evident that Cosine distance, as a simple OOD detector, consistently delivers superior performance and ranks among the top performers in both the zero-grad and generative fine-tuning settings. We provide an explanation of this phenomenon from the perspective of representation learning. In the past few years, the anisotropic issue, also known as the representation degeneration problem, of BERT family models has garnered considerable attention (Ethayarajh, 2019; Gao et al., 2019). Researchers have highlighted that BERT’s sentence representations are concentrated within a narrow cone, resulting in substantial challenges for tasks involving semantic matching. Nevertheless, we discover that this concern does not apply to LLMs. The representations generated by off-the-shelf LLMs inherently exhibit isotropy, enabling Cosine distance to excel in OOD detection. To quantify anisotropy, we adopt the methodology introduced by Ethayarajh (2019) to measure sentence-level anisotropy. Let  $\mathbf{X}_i$  be a sentence that appears in the corpus. The anisotropy value can be calculated by:

$$\text{anisotropy} = \frac{1}{n^2 - n} \left| \sum_i \sum_{j \neq i} \cos(z(\mathbf{X}_i), z(\mathbf{X}_j)) \right|, \quad (2)$$

where  $\cos$  is the cosine similarity and  $z(\cdot)$  is the sentence embedding from the last layer. A higher anisotropy value suggests that the sentence embeddings are less distinguishable by Cosine distance. The quantitative results presented in Table 5 show that the anisotropy values of LLMs are considerably lower in comparison to those of RoBERTa.



## 6. Analysis

### 6.1. Performance of More LLMs

ID Dataset	PTMs	Maha	Cosine	MSP	Energy
<b>Zero-grad</b>					
SST-2	OPT-6.7B	0.982	0.983	0.413	<b>0.571</b>
	LLaMA-7B	<b>0.991</b>	0.990	0.905	0.368
	LLaMA2-7B	<b>0.991</b>	<b>0.997</b>	<b>0.917</b>	0.516
<b>Fine-tuned</b>					
CLINC-Banking (Full)	OPT-6.7B	0.921	0.932	0.915	0.922
	LLaMA-7B	0.958	0.964	0.936	0.930
	LLaMA2-7B	<b>0.967</b>	<b>0.970</b>	<b>0.944</b>	<b>0.937</b>

Table 6: OOD detection performance of various LLMs. AUROC result is reported.

Additionally, we test OPT-6.7B (Zhang et al., 2022) and LLaMA2-7B (Touvron et al., 2023b) here with settings that zero-grad OOD performance for SST-2 (ID) and fine-tuned OOD performance for CLINC-Banking, as shown in Table 6. For simplicity, we report the AUROC result of each OOD score function. Overall, the OOD detection performance of LLaMA2-7B surpasses that of LLaMA-7B, and LLaMA-7B outperforms OPT-6.7B, which is consistent with the general performance trends observed in these models.

### 6.2. Impact of Quantization

Quantized	SST-2			
	Maha	Cosine	MSP	Energy
float32	0.991	0.990	0.905	0.368
float16	0.990	0.987	0.893	0.331
Int8	0.955	0.960	0.874	0.228

Table 7: Impact of quantization for OOD detection.

We test the zero-grad OOD detection performance of LLaMA-7B with different quantized levels (float32, float16, and Int8) for SST-2 as the ID task, shown in Table 7. It was observed that the float16 model can largely preserve the OOD detection capability of the model. However, 8-bit quantization leads to a degradation in its capability.

### 6.3. Error Analysis

We mainly analyze the fine-tuning OOD detection performance in the near-OOO setting since both RoBERTa (Liu et al., 2019) and LLaMA (Touvron et al., 2023a) can achieve near-perfect performance in the far-OOO setting. Here, we provide the OOD detection performance with fine-tuned RoBERTa in the following Table 8, as a complement to Table 3.

Overall, LLaMA’s detection capabilities are much stronger than RoBERTa’s. Through the

PTMs	CLINC-Banking (Full)			
	Maha	Cosine	MSP	Energy
LLaMA-7B	<b>0.958</b>	<b>0.964</b>	<b>0.936</b>	<b>0.930</b>
RoBERTa-L	0.821	0.793	0.670	0.717

Table 8: Performance comparison between LLaMA and RoBERTa-L in the near-OOO setting. AUROC result is reported for each score function.

analysis of specific error cases, we found that RoBERTa struggles to distinguish semantically similar OOD samples. For instance, when choosing the “bank\_balance” class (e.g., the sentence “what’s my account balance”) as the ID distribution, RoBERTa tends to incorrectly classify the majority of “bill\_balance” class inputs (such as “what are my bills this month”) as “bank\_balance”. In contrast, LLaMA generally makes correct judgments in most cases. We attribute this phenomenon to the anisotropy characteristic which is explained in Section 5.2), i.e., sentence embeddings produced by the BERT family models have been noted to possess an undesirable characteristic of being concentrated within a narrow cone, causing representation degeneration. Despite this, in the cases involving extremely similar semantics, LLaMA also makes errors in judgment, such as misclassifying “bill\_balance: what is the balance on my bills” as “bank\_balance”.

## 7. Conclusion

This paper has delved into the critical realm of OOD detection within the context of LLMs. The growing utilization of LLMs across various natural language processing tasks has underscored the need to understand their capabilities and limitations, especially in scenarios involving distribution shifts. Our work deepens the comprehension of OOD detection capabilities of LLMs. Through meticulous analysis, we have showcased the effectiveness of LLMs for OOD detection under various settings, including zero-grad and generative fine-tuning scenarios. Our findings reveal that a simple OOD detector utilizing the cosine similarity function outperforms other sophisticated OOD detectors, especially in the few-shot setting. Our work may serve as a foundational stepping stone for future advancements in effectively and responsibly harnessing the potential of LLMs in diverse environments.

## Acknowledgments

We thank the anonymous reviewers for their valuable feedback. This research was partially supported by the grant of HK ITF ITS/359/21FP.

## 8. Bibliographical References

- Udit Arora, William Huang, and He He. 2021. [Types of out-of-distribution texts and how to detect them](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10687–10701, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurélie Névéol, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. [Findings of the 2016 conference on machine translation](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 131–198, Berlin, Germany. Association for Computational Linguistics.
- 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.
- Hyunsoo Cho, Choonghyun Park, Junyeop Kim, Hyuhng Joon Kim, Kang Min Yoo, and Sang-goo Lee. 2023. Probing out-of-distribution robustness of language models with parameter-efficient transfer learning methods. *arXiv preprint arXiv:2301.11660*.
- 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*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. Glm: General language model pretraining with autoregressive blank infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 320–335.
- Desmond Elliott, Stella Frank, Khalil Sima'an, and Lucia Specia. 2016. Multi30k: Multilingual english-german image descriptions. In *Proceedings of the 5th Workshop on Vision and Language*, pages 70–74.
- Kawin Ethayarajh. 2019. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings. *arXiv preprint arXiv:1909.00512*.
- Yujie Feng, Zexin Lu, Bo Liu, Liming Zhan, and Xiao-Ming Wu. 2023. Towards llm-driven dialogue state tracking. *arXiv preprint arXiv:2310.14970*.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2019. Representation degeneration problem in training natural language generation models. *arXiv preprint arXiv:1907.12009*.
- 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 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzic, Rishabh Krishnan, and Dawn Song. 2020. Pretrained transformers improve out-of-distribution robustness. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2744–2751.
- Dan Hendrycks and Mantas Mazeika. 2022. X-risk analysis for ai research. *arXiv preprint arXiv:2206.05862*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Ken Lang. 1995. Newsweeder: Learning to filter netnews. In *Machine learning proceedings 1995*, pages 331–339. Elsevier.
- Stefan Larson, Anish Mahendran, Joseph J Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K Kummerfeld, Kevin Leach, Michael A Laurenzano, Lingjia Tang, et al. 2019. An evaluation dataset for intent classification and out-of-scope prediction. *arXiv preprint arXiv:1909.02027*.
- Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and Fuyie Huang. 2006. A tutorial on energy-based learning. *Predicting structured data*, 1(0).

- Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. 2018. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31.
- Xin Li and Dan Roth. 2002. Learning question classifiers. In *COLING 2002: The 19th International Conference on Computational Linguistics*.
- Ziqian Lin, Sreya Dutta Roy, and Yixuan Li. 2021. Mood: Multi-level out-of-distribution detection. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 15313–15323.
- 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.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pages 142–150.
- Leland McInnes, John Healy, and James Melville. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Yifei Ming, Hang Yin, and Yixuan Li. 2022. On the impact of spurious correlation for out-of-distribution detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10051–10059.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Alexander Podolskiy, Dmitry Lipin, Andrey Bout, Ekaterina Artemova, and Irina Piontkovskaya. 2021. Revisiting mahalanobis distance for transformer-based out-of-domain detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13675–13682.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training.
- 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.
- Mohammadreza Salehi, Hossein Mirzaei, Dan Hendrycks, Yixuan Li, Mohammad Hossein Rohban, and Mohammad Sabokrou. 2022. [A unified survey on anomaly, novelty, open-set, and out-of-distribution detection: Solutions and future challenges](#). *Trans. Mach. Learn. Res.*, 2022.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrusti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Rheeya Uppaal, Junjie Hu, and Yixuan Li. 2023. Is fine-tuning needed? pre-trained language models are near perfect for out-of-domain detection. *arXiv preprint arXiv:2305.13282*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent abilities of large language models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Adina Williams, Nikita Nangia, and Samuel R Bowman. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*.
- 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.
- Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. 2021. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b: An open bilingual pre-trained model. *arXiv preprint arXiv:2210.02414*.
- Li-Ming Zhan, Haowen Liang, Bo Liu, Lu Fan, Xiao-Ming Wu, and Albert Lam. 2021. Out-of-scope intent detection with self-supervision and discriminative training. *arXiv preprint arXiv:2106.08616*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Wenxuan Zhou, Fangyu Liu, and Muhao Chen. 2021. Contrastive out-of-distribution detection for pretrained transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1100–1111.