

TagOOD: A Novel Approach to Out-of-Distribution Detection via Vision-Language Representations and Class Center Learning

Anonymous Authors

ABSTRACT

Multimodal fusion, leveraging data like vision and language, is rapidly gaining traction. This enriched data representation improves performance across various tasks. Existing methods for out-of-distribution (OOD) detection, a critical area where AI models encounter unseen data in real-world scenarios, rely heavily on whole-image features. These image-level features can include irrelevant information that hinders the detection of OOD samples, ultimately limiting overall performance. In this paper, we propose **TagOOD**, a novel approach for OOD detection that leverages vision-language representations to achieve label-free object feature decoupling from whole images. This decomposition enables a more focused analysis of object semantics, enhancing OOD detection performance. Subsequently, TagOOD trains a lightweight network on the extracted object features to learn representative class centers. These centers capture the central tendencies of IND object classes, minimizing the influence of irrelevant image features during OOD detection. Finally, our approach efficiently detects OOD samples by calculating distance-based metrics as OOD scores between learned centers and test samples. We conduct extensive experiments to evaluate TagOOD on several benchmark datasets and demonstrate its superior performance compared to existing OOD detection methods. This work presents a novel perspective for further exploration of multimodal information utilization in OOD detection, with potential applications across various tasks. Code will be available.

CCS CONCEPTS

• **Computing methodologies** → *Object recognition*.

KEYWORDS

Out-of-distribution detection, Vision-Language representations, Representative class centers

1 INTRODUCTION

Most modern deep neural networks [7, 14, 15, 27, 32, 47] are validated using test data from the same distribution as the training data. Nevertheless, encountering out-of-distribution (OOD) samples is inevitable when deploying deep learning models in real-world scenarios. This phenomenon highlights the importance of an ideal recognition model, which should not only deliver accurate predictions on the training distribution (referred to as in-distribution data)

Unpublished working draft. Not for distribution.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted by ACM, provided that the copies are not made for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM MM, 2024, Melbourne, Australia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/nmmmmmmmmmmmm>

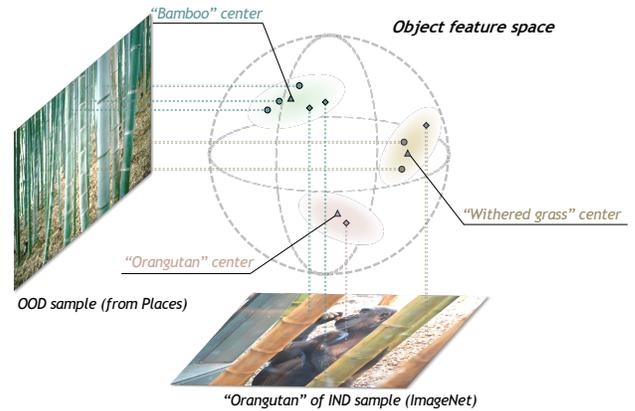


Figure 1: This toy sample illustrates a challenge in OOD detection using a 3D feature space. The features associated with "Bamboo" and "Withered grass" in the IND image are spatially close to the OOD image's features in this space. This proximity can lead the model to misclassify the IND image, mistaking it for containing similar objects to the OOD image.

but also raise alerts to humans when encountering unknown samples. OOD detection is the task of determining whether a test sample falls within the in-distribution (IND) or not. This task finds broad applications in autonomous driving [3, 4], fraud detection [37], medical image analysis [12, 41, 51], etc.

Current OOD detection methods [18, 19, 25, 28, 31, 42, 43, 52] often rely on image-level features to construct a score function for identifying OOD samples. In these approaches, a test sample is classified as OOD if its score surpasses a threshold, or vice versa. A major limitation of these approaches is that a single label in the training data fails to fully describe the content of an image, especially when the image contains multiple objects. In such cases, the model can learn "OOD-like" features from the IND data, leading to misclassifications. In essence, the model can learn some "OOD-like" characteristics even without encountering any actual OOD samples during training. A toy example is illustrated in Figure 1. We utilize an image tagging model to generate tags for both an IND image labeled as "orangutan" and an OOD image. Splitting the image based on these tags exposes a critical insight for OOD detection: a part of the IND image shares surprising feature similarity with the OOD image. For example, tags like "bamboo" and "withered grass" might be present in both. This observation highlights a crucial disconnect between the model-captured visual features and the semantic meaning of the image. While the "orangutan" label might be accurate, the presence of "bamboo" and "withered grass" suggests the image includes background elements that could be commonly found in OOD data. This disconnect motivates our exploration of vision-language representations for OOD detection. Our approach

integrates semantic information with visual features and learns the IND area of each image that truly corresponds to the label, resulting in a more refined IND distribution for enhanced OOD detection.

In this paper, we propose **TagOOD**, a novel approach for Out-of-Distribution detection that leverages vision-language representations. TagOOD specifically addresses the challenge of confusing OOD samples, which may contain objects visually similar to objects found in the in-distribution data. These OOD objects are particularly problematic because they not be explicitly labeled within the IND data. TagOOD tackles this challenge through two key mechanisms: 1) decoupling image features using a tagging model to focus on semantic content beyond the object, and 2) generating object-level class centers to capture the central tendencies of IND objects, minimizing the influence of irrelevant background features.

TagOOD leverages a pre-trained tagging model, typically a large vision-language model trained on extensive image-text datasets. This model goes beyond single-label classification by identifying and assigning multiple semantic tags to objects within an image. By incorporating the tagging model, TagOOD focuses on object-level features for OOD detection. Specifically, for a given in-distribution image, the tagging model generates a set of object tags along with their corresponding features. The object features representing IND objects are then fed into a lightweight network for projection into a common feature space. This allows for the subsequent generation of trainable class centers, representing the central tendencies of each object category within the IND data. Finally, TagOOD calculates an OOD score based on the cosine similarity between these learned class centers and a test sample. A high cosine similarity indicates the test sample is likely IND data, as it closely resembles the central tendencies learned from the training data. Conversely, a low cosine similarity suggests the test sample is likely OOD.

The contributions of this paper are summarized as follows:

- Decoupling image features using a vision-language model: allows TagOOD to focus on the semantic content of the object, providing a deeper understanding of the IND data and mitigating the influence of irrelevant background features.
- Generating object-level class centers: by capturing the central tendencies of objects within the in-distribution data, TagOOD can effectively distinguish between IND and OOD samples, even if they contain similar objects.
- To investigate the effectiveness of TagOOD, we conducted comprehensive experiments and ablation studies on the popular ImageNet benchmark. These studies demonstrate competitive performance and reveal the importance of both decoupling features and generating object-level class centers for TagOOD's superior performance.

2 RELATED WORK

2.1 Vision-Language Models for Image Tagging

Recent advancements [1, 10, 44] in large language models (LLMs) have revolutionized natural language processing (NLP). In computer vision (CV), vision-language models (VLMs) [29, 56] bridge the gap between visual and textual data, leveraging the strengths of both modalities for tasks like image tagging. VLMs fall into two main categories: generation-based and alignment-based. Generation models create captions for images, with recent advancements

in language models [7, 35] making this approach dominant. Unlike earlier methods [9] that relied solely on recognizing tags and then composing captions, these models [5, 26, 50] leverage powerful language models for text generation conditioned on visual information. Alignment models, on the other hand, aim to determine if an image and its description match. Many works [21, 22, 35] typically rely on aligning features from both modalities, often using a dual-encoder or fusion-encoder architecture. Image Tagging, a crucial computer vision task that involves assigning multiple relevant labels to an image, plays a vital role in VLM performance. Traditionally, classifiers and loss functions were used for image tagging. Recent work [30, 39] utilizes transformers and robust loss functions to address challenges like missing data and unbalanced classes. Our work utilizes this powerful VLM technology to enhance OOD detection performance.

2.2 Out-of-Distribution Detection

Out-of-distribution detection aims to distinguish OOD samples from IND data. Numerous methods have been proposed for OOD detection. Maximum softmax probability (MSP) [16] serves as a common baseline, utilizing the highest score across all classes as an OOD indicator. ODIN [28] builds upon MSP by introducing input perturbations and adjusting logits through rescaling. Gaussian discriminant analysis (GDA) has also been employed for OOD detection in works by [25, 52]. ReAct [42] leverages rectified activation to mitigate model overconfidence in OOD data. In recent developments, several multi-modal methods [8, 33, 34, 48] have emerged to address the challenge of OOD detection, utilizing the generalized representations acquired through CLIP [11]. CLIPN [48] introduces an extra text encoder to CLIP, aligning entire images with two types of textual information and enabling the model to respond with "no" when faced with OOD samples. MCM [34] aligns image-level features with their corresponding textual description features, enhancing the capabilities of IND estimates and consequently improving OOD detection performance. While CLIP demonstrates remarkable zero-shot OOD detection capabilities, these OOD detection methods utilizing CLIP still neglect to address the challenge of similar object confusion between IND and OOD images. Different from these CLIP-based methods, we attempt to exploit an image tagging model to generate multiple tags in terms of mitigating the impact of confused objects.

3 METHOD

3.1 Preliminary

Out-of-distribution(OOD) detection aims to distinguish between data the model has seen during training and unseen data from a different distribution. Formally, we can represent the training set as $\mathcal{D}_{IND}^{train} = \{(x_i, y_i)\}_{i=1}^n$, where $x_i \in \mathbb{R}^{3 \times H \times W}$ is the input image typically a 3-channel image with height H and width W , $y_i \in \{1, 2, \dots, K\}$, one of K possible IND categories, is the corresponding class label, and n is the number of samples in training set $\mathcal{D}_{IND}^{train}$. When a recognition model F is presented with a test set $\mathcal{D}^{test} = \{x'_i\}_{i=1}^p$, the goal is to split the test data into IND and OOD categories, denoted as \mathcal{D}_{IND}^{test} and \mathcal{D}_{OOD}^{test} respectively. Unlike the training data with labeled categories, OOD data typically lacks

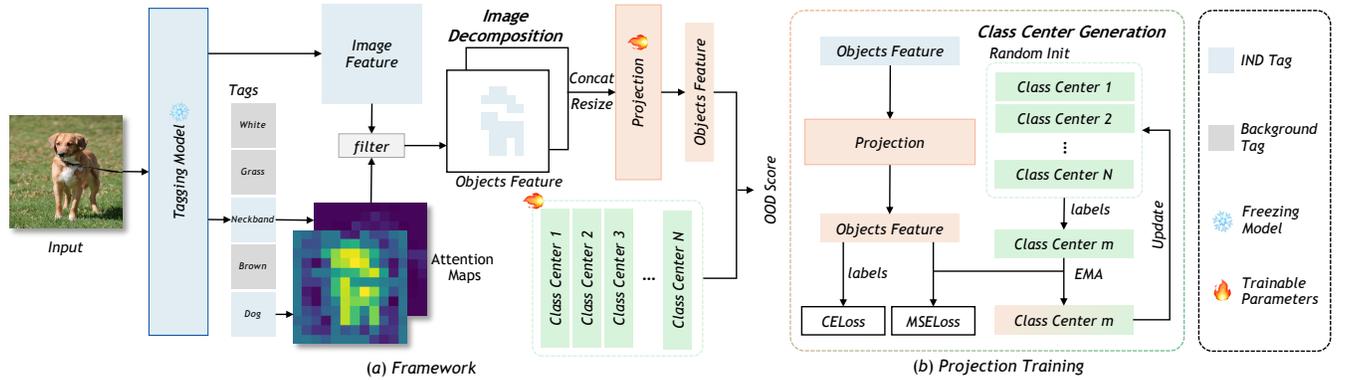


Figure 2: The TagOOD pipeline for OOD detection consists of two main stages, illustrated in (a) and (b). First, image feature decomposition (see (a) on the left) leverages a vision-language model to generate multiple tags. The model then identifies tags belonging to the IND category and creates corresponding attention masks within the image. Next, as shown in (b), the extracted IND object features are used to train a lightweight network that produces a set of IND class centers. During inference (referring back to (a)), TagOOD computes a distance-based metric between IND class centers and test sample features as the OOD score.

labels entirely and does not fall within the set of K possible IND categories. This separation is often achieved by defining an OOD score function $h(x)$ which assigns a score to a test image x :

$$h(x) = \begin{cases} 0, & \text{if } x \in \mathcal{D}_{OOD}^{test}, \\ 1, & \text{if } x \in \mathcal{D}_{IND}^{test}, \end{cases} \quad (1)$$

As shown in Equation (1), the function outputs 0 if the image likely belongs to the OOD data, and 1 if likely belongs to the IND data. In practice, the OOD score function provides a continuous value between 0 and 1, with a higher score indicating a higher likelihood of the data point being IND. This score is then used to classify the test data into IND and OOD categories.

3.2 Overview of TagOOD

Our proposed TagOOD approach aims to address a key challenge in OOD detection: distinguish the confusing OOD samples, which contain objects similar to objects found in the training data. These objects in test samples can be particularly problematic because they are non-labeled but occur in in-distribution data. Figure 2 illustrates our framework, which consists of two main steps: **1) Image Decomposition.** In Figure 2(a), TagOOD first leverages a pre-trained tagging model to generate multiple semantic tags for each object within an image. The model then identifies tags belonging to the IND category and creates corresponding attention masks within the image. By analyzing the attention maps produced by the tagging model, TagOOD identifies the relevant regions within the image feature extracted by the backbone network of the pre-trained tagging model. This allows our approach to obtain object features for each identified category object (See Sec 3.3 for details). **2) Class Center Generation.** We generate object-level class centers to further accurately construct the IND distribution and eliminate the error impact caused by the tagging model. Once IND object features are extracted, TagOOD employs a lightweight projection model. This model learns to project these object features, which might have different shapes due to the varying nature of objects, into a common feature space. Simultaneously, as Figure 2(b), it updates class centers (see Sec 3.4 for details) for each object category.

During the inference phase, referring back to Figure (a), TagOOD leverages a distance-based metric, such as Euclidean distance or cosine similarity, for OOD detection. This metric directly compares the features extracted from a test sample with the learned class centers representing the IND data. A larger distance or lower cosine similarity between the test features and the nearest class center suggests a higher likelihood of the sample being OOD. This strategy effectively utilizes the learned class centers that capture the central tendencies of objects within the training data. By identifying deviations from these expected patterns through distance metrics, TagOOD achieves robust OOD detection.

3.3 Vision-Language Approach for Image Feature Decomposition

TagOOD leverages a pre-trained vision-language model to identify and label objects within an image, enabling feature decomposition. In practice, we use RAM [56], which combines three components for generating image features, tags, and corresponding attention maps required for feature decomposition: **1) Image Encoder f** captures features from the input image, capturing visual information about all of the objects within the image. **2) Text Encoder g** based on the CLIP model [38] encodes potential object tags from a pre-defined vocabulary T into a common feature space. This vocabulary $T = \{t_i\}_{i=1}^M$, provided by RAM, containing M object used for more detailed image labeling. **3) Tagging Head ϕ** combines the image features and tag information to predict the final result. The result contains the tags $T^x = \{t\}$ for the objects within the image and corresponding attention maps $A^x = \{a\}$, where $a \in \mathbb{R}^{H \times W}$. In practice, we utilize the last cross-attention module within the Tagging Head to generate attention maps. Eq. 2 summarizes how the tagging model operates:

$$T^x, A^x = \phi(f(x), g(T)), \quad (2)$$

When the vision-language model processes an image x , we will obtain an image feature $f(x)$, predicted tags T^x for the objects within the image, and corresponding attention maps A^x for each

349 predicted tag. We then filter T^x to retain only the tags belonging
350 to our IND vocabulary T^{in} .

351 TagOOD leverages the attention maps A^x to achieve image fea-
352 ture decomposition by selecting IND object features from the full
353 image feature $f(x)$. In Figure 2(a), Our filter applies a threshold
354 τ to the attention map values corresponding to the object. This
355 essentially creates a binary mask (values above τ become 1, others
356 become 0), highlighting the most influential image regions of each
357 object. Instead of element-wise multiplication, TagOOD directly se-
358 lects features from the original image representation corresponding
359 to locations with 1 in the mask. This effectively chooses the parts
360 of the image feature that are most relevant to the objects based on
361 the predicted tag. When multiple IND objects are detected, their
362 corresponding object features are combined into a single feature
363 representation, denoted as z . This combined feature vector then
364 undergoes reshaping to prepare it for the next step of the model.
365 Through the process described above, we can obtain a training
366 dataset consisting of object features and their corresponding labels.
367 This dataset is denoted as $\mathcal{X}_{IND}^{train} = \{(z_i, y_i)\}_{i=1}^n$, where z_i
368 represents the object features for the i -th sample, y_i represents the
369 corresponding label for the i -th sample, indicating its IND category.

370 Since the original vocabulary T might not encompass all IND
371 labels, we construct an IND vocabulary which is a subset of T ,
372 denoted by T^{in} . Before training TagOOD, we go through a process
373 to select the most relevant tags T^{in} representing the IND data from
374 T . T^{in} are used during training to establish the central tendencies of
375 objects within each category. We collect all the tags associated with
376 the objects within our training data. To ensure the selected tags
377 accurately represent the ground truth, human experts review the
378 results after identifying the most frequent tag within each category.
379 Through this process, we obtain a set of IND tags T^{in} from the pre-
380 defined vocabulary T that best corresponds to the object categories
381 present in the training data.

3.4 Object-Level Class Center Generation

384 TagOOD employs a lightweight projection model \mathbf{p} with dataset
385 $\mathcal{X}_{IND}^{train}$ to learn a set of object-level IND class centers, denoted as
386 $U = \{\mu^i\}_{i=1}^K$. These class centers represent the typical characteris-
387 tics of objects belonging to IND categories and serve as effective
388 indicators for OOD detection. Initially, we randomly initialize a set
389 of tensors $\{\mu^i\}_{i=1}^K$ to represent the class centers for each of the K
390 IND categories. During training, as illustrated in Figure 2(b), the
391 model receives the object features z , which can have varying sizes
392 depending on the specific objects. The projection model transforms
393 these features z into a common feature space. This ensures consis-
394 tency when comparing the projected features with μ . The object
395 feature z^c and its label y^c of category c are used to calculate a
396 combined loss function:

$$398 \mathcal{L} = \alpha \cdot \text{CE}(\mathbf{p}(z^c), y^c) + \beta \cdot \text{MSE}(\mathbf{p}(z^c), \mu^c), \quad (3)$$

400 where $\text{CE}(\cdot)$ represents the cross-entropy loss, $\text{MSE}(\cdot)$ stands for
401 mean squared error loss. The hyperparameters α and β control the
402 relative importance of each loss term during training. The loss term
403 $\text{CE}(\mathbf{p}(z^c), y^c)$ helps the model learn to differentiate between object
404 features belonging to different IND categories. $\text{MSE}(\mathbf{p}(z^c), \mu^c)$ term
405 encourages the model to refine the projected features to better

407 align with the central tendencies represented by the class centers.
408 Before the next training iteration, the IND class centers are updated.
409 TagOOD utilizes an exponential moving average (EMA) [13] with
410 $\mathbf{p}(z)$ to achieve this update. EMA helps smooth out fluctuations in
411 the projected features and provides a more stable estimate of the
412 central tendencies for each IND category. According to the process
413 we described above, the parameter of projection model θ^p and the
414 class centers $\{\mu^i\}_{i=1}^K$ are updated by the following equations:

$$415 \theta_{t+1}^p = \theta_t^p - \gamma_1 \frac{\partial \mathcal{L}}{\partial \theta_t^p}, \quad (4)$$

$$416 \mu_{t+1}^c = (1 - \gamma_2)\mu_t^c + \gamma_2 \mathbf{p}(z^c), \quad (5)$$

417 where γ_1 , and γ_2 are learning rates controlling the update speed.
418

3.5 OOD Detection Based on the Class Centers

420 When the training of our projection model \mathbf{p} finishes, we will get
421 a set of IND class centers U that capture the central tendencies of
422 the IND object accurately. To assess the likelihood of a test sample
423 x' belonging to the IND data, we compute the cosine similarity
424 between its projected features z' and each class center within U .
425 The maximum cosine similarity score across U is used as the OOD
426 score for the test sample, denoted by $\mathbf{h}(x')$. Intuitively, a high cosine
427 similarity score indicates that the test sample feature closely resem-
428 bles one of the learned IND class centers μ^i , suggesting it's likely
429 IND data. Conversely, a low OOD score suggests the test sample
430 deviates significantly from the expected IND patterns. The provided
431 mathematical formula accurately represents this calculation:

$$432 \mathbf{h}(x') = \max_c \left(\frac{z' \cdot \mu^c}{\|z'\| \|\mu^c\|} \right). \quad (6)$$

433 If the tagging model predicts no tags in T^{in} for a given image,
434 TagOOD directly classifies the sample as OOD. Overall, by compar-
435 ing the test sample feature (\mathbf{p} supports both image feature or
436 object feature) to the learned class centers using cosine similar-
437 ity, TagOOD can effectively distinguish between in-distribution
438 and out-of-distribution data, enhancing the model's robustness in
439 real-world applications.

4 EXPERIMENTS

4.1 Experimental Setup

440 **In-distribution Dataset.** We adopt the well-established ImageNet-
441 1K [40] dataset, a standard dataset for OOD detection, as our in-
442 distribution data. ImageNet-1K is a large-scale visual recognition
443 dataset containing 1000 object categories and 1281167 images.

444 **Out-of-Distribution Dataset.** Based on ImageNet as IND datasets,
445 we follow Huang *et al.* [20] to test our TagOOD with iNaturalist [46],
446 SUN [53], Places365 [57], and Texture [6]. To further explore the
447 generalization ability of our approach, we follow Wang *et al.* [49]
448 and employ another two OOD datasets, OpenImage-O [24] and
449 ImageNet-O [17] to evaluate our method.

450 **Evaluation Metrics.** We evaluate our approach using common
451 OOD detection metrics [20, 34, 42, 54, 55]: **Area Under the ROC
452 Curve (AUROC)**. This metric summarizes the model's ability to
453 distinguish between IND and OOD samples. Higher AUROC indi-
454 cates better overall performance. **False Positive Rate at 95%
455 True Positive Rate (FPR95)**. This metric measures the model's

Table 1: OOD detection performance comparison of TagOOD and existing methods. The best and second-best results are indicated in Bold and Underline. The method marked with * indicates that it utilizes a more expansive backbone in comparison to the others. \uparrow indicates larger values are better and \downarrow indicates the opposite. All values are expressed in percentages.

Method	OOD Datasets								Average	
	iNatrualist		SUN		Place		Texture		AUROC \uparrow	FPR95 \downarrow
	AUROC \uparrow	FPR95 \downarrow								
MSP[16]	95.60	18.89	86.88	49.18	85.64	51.73	<u>85.17</u>	49.66	88.32	42.37
ODIN[28]	71.04	48.53	60.47	69.00	57.01	72.69	63.03	68.88	62.89	64.78
ENERGY[31]	94.87	14.47	83.44	45.34	82.25	48.55	80.64	52.93	85.30	40.32
ReAct[42]	<u>97.61</u>	<u>7.32</u>	86.71	40.30	85.14	44.49	82.99	<u>48.97</u>	88.11	<u>35.27</u>
MCM(CLIP-L)*[34]	94.95	28.38	94.12	29.00	92.00	35.42	84.88	59.88	<u>91.49</u>	38.17
LHAct[55]	63.60	58.53	65.10	92.11	64.84	92.69	81.06	85.35	68.65	82.17
FeatureNorm[54]	66.42	89.70	60.77	92.22	61.44	91.80	55.86	95.78	61.12	92.38
TagOOD(Ours)	98.97	5.00	<u>92.22</u>	<u>29.70</u>	<u>87.81</u>	<u>40.40</u>	90.60	36.31	92.40	27.85

Table 2: Evaluation on more challenging datasets. The method marked with * indicates that it utilizes a more expansive backbone in comparison to the others. The best and second-best results are indicated in Bold and Underline. \uparrow indicates larger values are better and \downarrow indicates the opposite. All values are expressed in percentages.

Method	OOD Datasets				Average	
	ImageNet-O		OpenImage		AUROC \uparrow	FPR95 \downarrow
	AUROC \uparrow	FPR95 \downarrow	AUROC \uparrow	FPR95 \downarrow		
MSP[16]	85.22	59.25	92.57	29.89	88.90	44.57
ODIN[28]	72.85	57.15	69.21	52.73	71.03	54.94
ENERGY[31]	87.33	<u>37.05</u>	86.64	36.83	86.99	36.94
React[42]	87.85	34.45	<u>93.39</u>	<u>20.81</u>	90.62	27.63
MCM(CLIP-L)*[34]	82.46	68.75	92.92	35.84	87.69	52.30
LHAct[55]	57.72	91.75	59.27	70.66	58.50	81.21
FeatureNorm[54]	54.22	94.84	61.00	92.83	57.61	93.84
TagOOD(Ours)	<u>87.48</u>	51.91	96.28	20.44	91.88	<u>36.18</u>

tendency to misclassify OOD samples as IND samples when the true positive rate (correctly identified IND samples) is 95%. Lower FPR95 suggests the model effectively identifies OOD data even with a high true positive rate for IND data.

Training Details. We leverage RAM [56] as the underlying vision-language model for OOD sample detection. For consistency across all compared methods, the feature extractor within RAM, swin-large [32], is used for all approaches. Our lightweight projection model employs two serially connected self-attention block layers. The projected feature space has a dimensionality of 512. ImageNet-1k is used for training, with a total of 100 epochs, and the class centers are initialized with Gaussian noise for the 1000 ImageNet categories. The hyperparameters α , β , γ_1 , γ_2 are set to 1, 0.1, 1, 1×10^{-4} , respectively. These values were chosen based on our experimental results. During training, The Adam optimizer [23] is used during training with an initial learning rate of 0.01, a Cosine Annealing learning rate scheduler, and a batch size of 256. All the experiments are performed using PyTorch [36] with default parameters on two NVIDIA V100 GPUs.

4.2 Main results

Standard evaluation on ImageNet. We compare the performance of our TagOOD against seven popular OOD detection approaches, including MSP [16], ODIN [28], Energy [31], ReAct [42], MCM [34], LHAct [55], FeatureNorm [54]. All methods, except MCM, utilize Swin-L as the backbone extractor for a fair comparison. MCM employs a larger and more complex CLIP-based ViT-L model. As shown in Table 1, our method achieves better AUROC and FPR95 metrics. We highlight that TagOOD achieves 27.85% on FPR95, which outperforms the previous best method ReAct [42] by 7.42% across various OOD datasets. While MCM [34] exhibits higher AUROC and lower FPR95 on SUN and Places datasets due to its significantly larger **CLIP ViT-L backbone**, which is 1.7 times the size of Swin-L, TagOOD maintains the best average performance across all datasets. Our analysis suggests that LHAct [55] and FeatureNorm [54] might suffer limitations due to hyperparameter settings optimized for a different backbone network. This mismatch between hyperparameters and the chosen Swin-L model can lead to suboptimal performance on these methods.

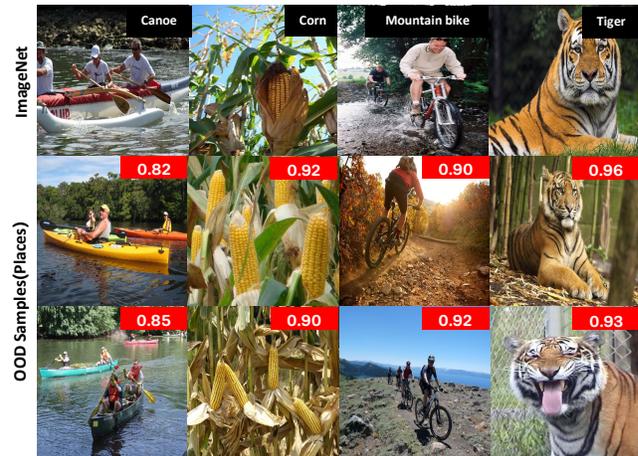


Figure 3: Overlap between IND and OOD samples. the black-tagged images are from ImageNet for reference and represent a standard IND class sample. The samples with red tags containing IND class objects are selected from the SUN and Places datasets. Red tags indicate the corresponding OOD scores assigned by our method.

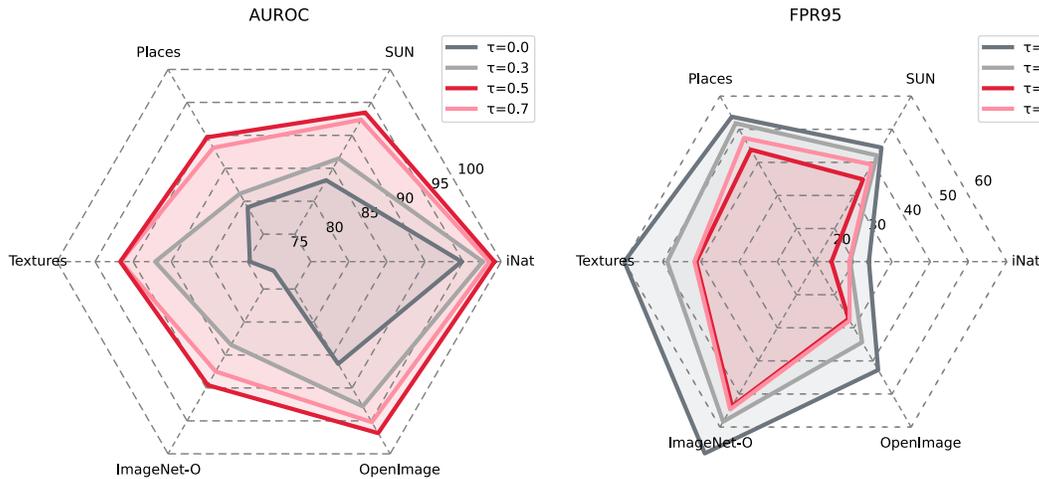


Figure 4: The performance of TagOOD training with varying values of the hyperparameter τ .

The study by Bitterwolf *et al.* [2] highlights a crucial aspect of OOD detection: the overlap between training and OOD data. The commonly used datasets in our evaluation (detailed in Section 4.1) exhibit varying degrees of overlap with the ImageNet-1k training data. This overlap can potentially influence the evaluation of a model to identify OOD samples. In Figure 3, the black-tagged images are from ImageNet for reference and represent a true IND class sample. The samples with red tags, selected from the SUN and Places datasets, contain IND class objects. These red tags indicate the corresponding OOD scores assigned by our method. TagOOD correctly assigns them OOD scores with high confidence, indicating it is reliable to accurately distinguish IND objects. This highlights the robust OOD detection performance of TagOOD, leveraging the power of the vision-language model.

Evaluation on more challenging OOD datasets. To further assess the limitations and robustness of our TagOOD, we conducted experiments on two challenging datasets: OpenImage-O [24] and

ImageNet-O [17]. OpenImage-O contains a diverse set of OOD samples, while ImageNet-O specifically includes adversarial examples designed to fool detection models. As shown in Table 2, TagOOD demonstrates strong performance on OpenImage-O, indicating its effectiveness in handling general OOD data. However, on ImageNet-O, which contains specifically crafted adversarial examples, the performance of TagOOD shows some limitations while still achieving the best AUROC score. This suggests that further exploration is needed to enhance TagOOD’s resilience against adversarial attacks.

4.3 A closer look of TagOOD

How does image feature decomposition affect the performance? We investigate the influence of the hyperparameter τ on image feature decomposition for our projection model training. Figure 4 visualizes the effects of varying τ values on model performance. Specifically, we set $\tau = \{0, 0.3, 0.5, 0.7\}$. Higher τ values correspond to a more selective feature subset, prioritizing those

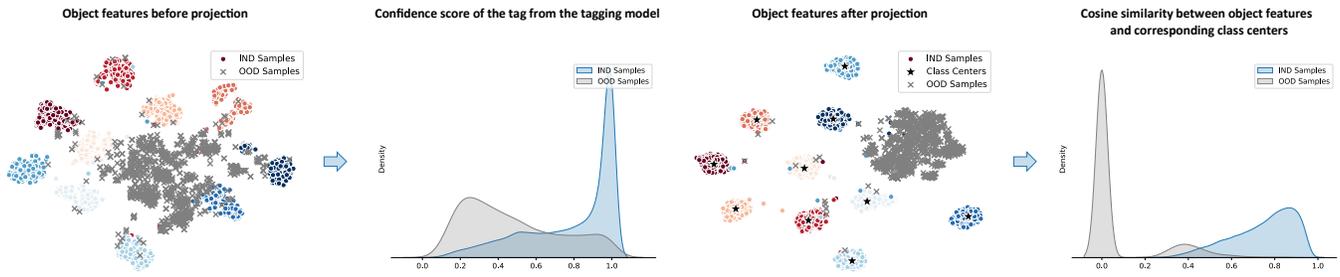


Figure 5: Illustration of object features visualization and OOD score distribution. Both object features are visualized by T-SNE [45]. Data points represent object features, and colors encode their corresponding IND class labels. Gray X marks indicate OOD data points. The features presented on the left are directly extracted from the tagging model before projection. Following the projection process, the object features become more condensed, as demonstrated on the right.

with strong alignment to IND object classes. Conversely, $\tau = 0$ represents a baseline where all image features are directly used for object model training. We observe that the optimal performance across all the OOD datasets used in our evaluation is achieved at $\tau = 0.5$. This finding suggests a critical balance between feature selectivity and information retention for effective OOD detection. Interestingly, the performance suffers more significantly when τ decreases from 0.5. In contrast, increasing τ beyond 0.5 leads to only a slight performance degradation. We hypothesize that this occurs because higher τ values eliminate some redundant information without discarding critically important features for accurate OOD classification. In essence, our model demonstrates the ability to retain the most crucial information from image samples while discarding irrelevant features at $\tau = 0.5$. This selectivity contributes to the reliable performance of TagOOD in identifying OOD data.

Table 3: A set of ablation results for TagOOD, averaged across four standard OOD datasets. CCG stands for class center generation, Proj. represents the projection model training, CS means Cosine Similarity distance and CE and MSE represent Cross Entropy loss and Mean Square Error loss.

	CCG	Proj.	AUROC \uparrow	FPR95 \downarrow
Tag Score			85.62	43.44
CS	✓		83.28	39.94
$p(\text{CE})+\text{CS}$		✓	91.58	32.13
$p(\text{CE}+\text{MSE})+\text{CS}$	✓	✓	92.40	27.85

What is the effect of the projection model and class centers?

Table 3 and Figure 5 systematically explore the influence of the projection model and class centers on the performance of TagOOD. This analysis helps us understand how each step contributes to the final OOD detection effectiveness. Tag Score in Table 3 utilizes the confidence score directly from the tagging model as the OOD score. CS leverages object features obtained after the image feature decomposition. In this case, we compute the average feature of each IND class as the class center. Cosine similarity is then computed as the OOD score. $p(\text{CE})+\text{CS}$ uses the projection model trained solely with a CE loss to project the object features. Subsequent operations are the same as CS. The bottom line represents the full

TagOOD approach. It incorporates all the previous steps: image feature decomposition, a projection model trained with a combined loss function which includes both CE loss and MSE loss, and class center generation. As observed, each stage progressively improves performance, culminating in superior OOD detection capabilities. This highlights the importance of both the projection model and class centers in achieving reliable OOD detection.

To gain deeper insights into the operation of TagOOD, Figure 5 utilizes T-SNE [45] for visualization. The data points represent object features, and colors encode their corresponding IND class labels. The gray 'X' markers represent OOD samples. A critical observation is the improved feature separation achieved through projection. Features that are directly procured from the tagging model, shown on the left, are likely to be scattered and overlapping. On the right side, we see the features after the projection operation in TagOOD. These features form tighter and more distinct clusters, with clearer boundaries between different IND classes. This enhanced separation in the projected feature space allows TagOOD to assign more distinguishable OOD scores. This facilitates the identification of OOD samples because their features deviate significantly from the learned representations of known object classes.

Table 4: Performance of TagOOD using various distance metrics. Results are averaged across four standard OOD datasets.

Metrics	AUROC \uparrow	FPR95 \downarrow
KL Divergence	90.66	32.53
Euclidean Distance	92.03	28.74
Cosine Similarity	92.40	27.85

Is TagOOD robust with various distance metrics?

This ablation experiment investigates the sensitivity of TagOOD to various distance metrics commonly used in OOD detection: KL divergence, Euclidean distance, and cosine similarity. The results in Table 4 demonstrate that TagOOD exhibits robustness across these metrics, maintaining strong performance. While KL divergence and Euclidean distance show slight performance degradation compared to cosine similarity, the effectiveness remains consistent. This suggests that TagOOD is not overly reliant on a specific distance metric for accurate OOD detection.

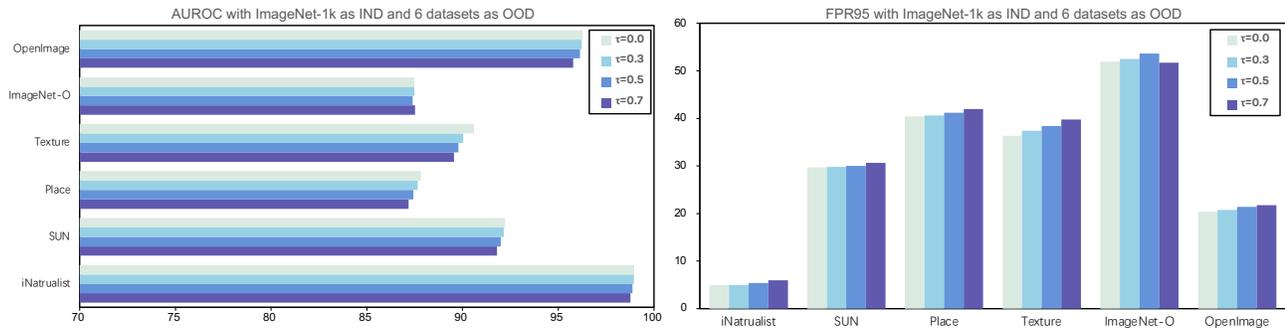


Figure 6: Results of evaluation on varying levels of features selected during image feature decomposition.

Is the performance of TagOOD enhanced when training utilizes a single ground truth tag as opposed to multiple predicted tags? This ablation experiment explores the influence of tagging strategies on object feature selection for OOD detection. We evaluate two strategies: utilize a single tag corresponding to the ground truth for feature decomposition or leverage the predictions of the tagging model, generating multiple tags for each image. The decomposition is then guided by the attention masks constructed based on these predicted tags. As shown in Table 5, the multiple tags strategy achieves superior performance compared to using the single ground-truth tag. We hypothesize that relying solely on the ground truth might lead to inaccurate feature decomposition due to potential tagging model errors. These errors could introduce "missing information" or "offset issues" in the decomposed features, hindering the ability to capture essential characteristics for effective OOD detection. In contrast, the multiple tags strategy incorporates the predictions of the tagging model, potentially offering a richer representation of the image content. This can lead to more accurate feature decomposition and the extraction of more informative features, ultimately enhancing OOD detection capabilities.

Table 5: Ablation experiment about impact of tag selection strategies on TagOOD Performance.

Tag for Training	AUROC \uparrow	FPR95 \downarrow
One tag	91.71	31.41
Multiple tags	92.40	27.85

Does TagOOD Maintain Robustness with Feature Variations?

This section investigates the ability of TagOOD to distinguish OOD data at inference despite variations introduced (different values of τ) during image feature decomposition. As mentioned earlier, our projection model utilizes both image and object features. To explore the robustness of TagOOD against such variations, we evaluate TagOOD on all six OOD datasets while adjusting the hyperparameter τ in the image feature decomposition stage. This manipulation effectively controls the level of detail extracted from image features. The results visualized in Figure 6 are encouraging. They demonstrate that TagOOD consistently maintains strong performance in OOD detection, regardless of the utilization of either image features or object features, and irrespective of how the hyperparameter τ impacts

image feature decomposition. This suggests that the projection model in TagOOD effectively learns discriminative features and exhibits robustness to variations in feature representation extracted during image feature decomposition.

5 LIMITATIONS

While TagOOD demonstrates promising results, there are areas where further exploration could enhance its capabilities. The performance of the tagging model can influence TagOOD's ability to accurately detect OOD data. Currently, TagOOD solely leverages the backbone network of the tagging model for feature extraction. This limits the potential for exploring alternative feature extractors that might be better suited for OOD detection tasks. Opening doors for investigating alternative feature extraction techniques will be an exciting future direction for this research.

6 CONCLUSION

Existing OOD detection methods often rely on whole-image features, which can incorporate irrelevant information beyond IND objects. This paper introduces TagOOD, a novel approach that leverages the power of vision-language models (VLMs) to achieve more focused OOD detection. TagOOD tackles the challenge of OOD detection by decoupling images from their corresponding classes and generating class centers within a common feature space. This innovative approach offers many advantages: TagOOD employs a VLM to decouple images from their associated classes. This allows the model to focus on the essential semantic information and learn the IND area of each image that truly corresponds to the label, resulting in a more refined IND distribution for enhanced OOD detection. TagOOD generates object-level class centers for each category within the IND data. These class centers serve as reference points in the common feature space, enabling the model to distinguish OOD samples effectively. A distance-based metric is used between the test sample feature and class centers for OOD detection. This simple yet effective approach allows the model to identify samples that fall outside the expected distribution of the IND data. This work presents a novel perspective for further exploration of multimodal information utilization in OOD detection, we hope our work contributes to the advancement of robust and reliable AI systems.

REFERENCES

- [1] Bard. 2023. A Large Language Model from Google AI. <https://ai.googleblog.com/2022/01/lambda-language-model-for-dialogue-and.html>
- [2] Julian Bitterwolf, Maximilian Mueller, and Matthias Hein. 2023. In or out? fixing imagenet out-of-distribution detection evaluation. *arXiv preprint arXiv:2306.00826* (2023).
- [3] Hermann Blum, Paul-Edouard Sarlin, Juan Nieto, Roland Siegwart, and Cesar Cadena. 2019. Fishyscapes: A benchmark for safe semantic segmentation in autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*. 0–0.
- [4] Daniel Bogdoll, Jasmin Breitenstein, Florian Heidecker, Maarten Bieshaar, Bernhard Sick, Tim Fingscheidt, and Marius Zollner. 2021. Description of Corner Cases in Automated Driving: Goals and Challenges. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 1023–1028.
- [5] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. 2022. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794* (2022).
- [6] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. 2014. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3606–3613.
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [8] Sepideh Esmailpour, Bing Liu, Eric Robertson, and Lei Shu. 2022. Zero-shot out-of-distribution detection based on the pre-trained model clip. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 36. 6568–6576.
- [9] Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al. 2015. From captions to visual concepts and back. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1473–1482.
- [10] Luciano Floridi and Massimo Chiriatti. 2020. GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines* 30 (2020), 681–694.
- [11] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. 2023. Clip-adapter: Better vision-language models with feature adapters. *International Journal of Computer Vision* (2023), 1–15.
- [12] Xiaoyuan Guo, Judy Wawira Gichoya, Saptarshi Purkayastha, and Imon Banerjee. 2021. CVAD: A generic medical anomaly detector based on Cascade VAE. *arXiv preprint arXiv:2110.15811* (2021).
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16000–16009.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [15] Shuai He, Anlong Ming, Shuntian Zheng, Haobin Zhong, and Huadong Ma. 2023. EAT: An Enhancer for Aesthetics-Oriented Transformers. In *Proceedings of the 31st ACM International Conference on Multimedia*. 1023–1032.
- [16] Dan Hendrycks and Kevin Gimpel. 2016. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136* (2016).
- [17] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. 2021. Natural adversarial examples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 15262–15271.
- [18] Yen-Chang Hsu, Yilin Shen, Hongxia Jin, and Zsolt Kira. 2020. Generalized ODIN: Detecting Out-of-Distribution Image Without Learning From Out-of-Distribution Data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [19] Rui Huang, Andrew Geng, and Yixuan Li. 2021. On the importance of gradients for detecting distributional shifts in the wild. *Advances in Neural Information Processing Systems* 34 (2021), 677–689.
- [20] Rui Huang and Yixuan Li. 2021. MOS: Towards Scaling Out-of-distribution Detection for Large Semantic Space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [21] Xinyu Huang, Youcai Zhang, Ying Cheng, Weiwei Tian, Ruiwei Zhao, Rui Feng, Yuejie Zhang, Yaqian Li, Yandong Guo, and Xiaobo Zhang. 2022. Idea: Increasing text diversity via online multi-label recognition for vision-language pre-training. In *Proceedings of the 30th ACM International Conference on Multimedia*. 4573–4583.
- [22] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*. PMLR, 4904–4916.
- [23] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [24] Ivan Krasin, Tom Duerig, Neil Aldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, et al. 2017. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from https://github.com/openimages* 2, 3 (2017), 18.
- [25] 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*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2018/file/abdeb6f575ac5c6676b747bca8d09cc2-Paper.pdf>
- [26] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*. PMLR, 12888–12900.
- [27] Xujin Li, Wei Wei, Shuang Qiu, and Huiguang He. 2022. TFF-Former: Temporal-frequency fusion transformer for zero-training decoding of two BCI tasks. In *Proceedings of the 30th ACM international conference on multimedia*. 51–59.
- [28] Shiyu Liang, Yixuan Li, and R. Srikant. 2018. Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=H1VGkIXRZ>
- [29] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024. Visual instruction tuning. *Advances in neural information processing systems* 36 (2024).
- [30] Shilong Liu, Lei Zhang, Xiao Yang, Hang Su, and Jun Zhu. 2021. Query2label: A simple transformer way to multi-label classification. *arXiv preprint arXiv:2107.10834* (2021).
- [31] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. 2020. Energy-based Out-of-distribution Detection. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 21464–21475. <https://proceedings.neurips.cc/paper/2020/file/f5496252609c43eb8a3d147ab9b9c006-Paper.pdf>
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*. 10012–10022.
- [33] Felix Michels, Nikolas Adaloglou, Tim Kaiser, and Markus Kollmann. 2023. Contrastive Language-Image Pretrained (CLIP) Models are Powerful Out-of-Distribution Detectors. *arXiv preprint arXiv:2303.05828* (2023).
- [34] Yifei Ming, Ziyang Cai, Jiuxiang Gu, Yiyun Sun, Wei Li, and Yixuan Li. 2022. Delving into out-of-distribution detection with vision-language representations. *Advances in neural information processing systems* 35 (2022), 35087–35102.
- [35] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems* 35 (2022), 27730–27744.
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).
- [37] Clifton Phua, Vincent Lee, Kate Smith, and Ross Gayler. 2010. A comprehensive survey of data mining-based fraud detection research. *arXiv preprint arXiv:1009.6119* (2010).
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- [39] Tal Ridnik, Gilad Sharir, Avi Ben-Cohen, Emanuel Ben-Baruch, and Asaf Noy. 2023. MI-decoder: Scalable and versatile classification head. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 32–41.
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115 (2015), 211–252.
- [41] Nina Shvetsova, Bart Bakker, Irina Fedulova, Heinrich Schulz, and Dmitry V. Dylow. 2021. Anomaly Detection in Medical Imaging With Deep Perceptual Autoencoders. *IEEE Access* 9 (2021), 118571–118583. <https://doi.org/10.1109/ACCESS.2021.3107163>
- [42] Yiyun Sun, Chuan Guo, and Yixuan Li. 2021. ReAct: Out-of-distribution Detection With Rectified Activations. In *Advances in Neural Information Processing Systems*.
- [43] Yiyun Sun, Yifei Ming, Xiaojin Zhu, and Yixuan Li. 2022. Out-of-distribution Detection with Deep Nearest Neighbors. *ICML* (2022).
- [44] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [45] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [46] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. 2018. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 8769–8778.

929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044

1045	[47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. <i>Advances in neural information processing systems</i> 30 (2017).	1103
1046		
1047	[48] Hualiang Wang, Yi Li, Huifeng Yao, and Xiaomeng Li. 2023. CLIPN for Zero-Shot OOD Detection: Teaching CLIP to Say No. <i>arXiv preprint arXiv:2308.12213</i> (2023).	1104
1048		1105
1049	[49] Haoqi Wang, Zhizhong Li, Litong Feng, and Wayne Zhang. 2022. Vim: Out-of-distribution with virtual-logit matching. In <i>Proceedings of the IEEE/CVF conference on computer vision and pattern recognition</i> . 4921–4930.	1106
1050		1107
1051	[50] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. 2021. Simvln: Simple visual language model pretraining with weak supervision. <i>arXiv preprint arXiv:2108.10904</i> (2021).	1108
1052		1109
1053	[51] Qi Wei, Yin hao Ren, Rui Hou, Bibo Shi, Joseph Y. Lo, and Lawrence Carin. 2018. Anomaly detection for medical images based on a one-class classification. In <i>Medical Imaging 2018: Computer-Aided Diagnosis</i> , Nicholas Petrick and Kensaku Mori (Eds.), Vol. 10575. International Society for Optics and Photonics, SPIE, 375 – 380. https://doi.org/10.1117/12.2293408	1110
1054		1111
1055	[52] Jim Winkens, Rudy Bunel, Abhijit Guha Roy, Robert Stanforth, Vivek Natarajan, Joseph R. Ledsam, Patricia MacWilliams, Pushmeet Kohli, Alan Karthikesalingam, Simon Kohl, Taylan Cemgil, S. M. Ali Eslami, and Olaf Ronneberger. 2020. Contrastive Training for Improved Out-of-Distribution Detection. <i>arXiv preprint</i>	1112
1056		1113
1057		1114
1058		1115
1059		1116
1060		1117
1061		1118
1062		1119
1063		1120
1064		1121
1065		1122
1066		1123
1067		1124
1068		1125
1069		1126
1070		1127
1071		1128
1072		1129
1073		1130
1074		1131
1075		1132
1076		1133
1077		1134
1078		1135
1079		1136
1080		1137
1081		1138
1082		1139
1083		1140
1084		1141
1085		1142
1086		1143
1087		1144
1088		1145
1089		1146
1090		1147
1091		1148
1092		1149
1093		1150
1094		1151
1095		1152
1096		1153
1097		1154
1098		1155
1099		1156
1100		1157
1101		1158
1102		1159
		1160