

# Masked Contrastive Learning for Anomaly Detection

Hyunsoo Cho, Jinseok Seol and Sang-goo Lee

Seoul National University

{johyunsoo, jamie, sglee}@europa.snu.ac.kr

## Abstract

Detecting anomalies is one fundamental aspect of a safety-critical software system, however, it remains a long-standing problem. Numerous branches of works have been proposed to alleviate the complication and have demonstrated their efficiencies. In particular, self-supervised learning based methods are spurring interest due to their capability of learning diverse representations without additional labels. Among self-supervised learning tactics, *contrastive learning* is one specific framework validating their superiority in various fields, including anomaly detection. However, the primary objective of contrastive learning is to learn *task-agnostic* features without any labels, which is not entirely suited to discern anomalies. In this paper, we propose a *task-specific* variant of contrastive learning named *masked contrastive learning*, which is more befitted for anomaly detection. Moreover, we propose a new inference method dubbed *self-ensemble inference* that further boosts performance by leveraging the ability learned through auxiliary self-supervision tasks. By combining our models, we can outperform previous state-of-the-art methods by a significant margin on various benchmark datasets.

## 1 Introduction

Over the past few years, machine learning has achieved immense success surpassing human-level performance in many tasks, such as classification, segmentation, and object detection [Tan and Le, 2019; Tan *et al.*, 2020; Chen *et al.*, 2020a]. However, such a well-trained model assigns arbitrary high probability [Hein *et al.*, 2019] on the unfamiliar test samples, since most machine learning systems generally depend on the closed-set assumption (*i.e.*, i.i.d. assumption). This phenomenon may lead to a fatal accident in safety-critical applications like medical-diagnosis or autonomous driving. *Anomaly detection*<sup>1</sup> is a research area that aims to circumvent such symptoms by identifying whether the test samples

<sup>1</sup>also termed *out-of-distribution detection*, *novelty detection*, or *outlier detection* in the contemporary machine learning context.




Figure 1: *t*-SNE visualization of CIFAR-10 trained representation. Each color denotes 10 respective class labels in CIFAR-10 dataset. Stochastic Positive Attraction (SPA) indicates component in our model (MCL) and SupCLR is another *task-specific* variant of SimCLR. While other models show blurry decision boundaries between some class pairs, MCL forms dense cluster for each class label while preserving their unique individual representation of respective data point in each cluster. See Section 2 for further details.

come from in-distribution or not. A flurry of recent deep-learning based models, including reconstruction based [Oza and Patel, 2019; Li *et al.*, 2018], density estimation based [Malinin and Gales, 2018], post-processing methods [Lee *et al.*, 2018; Liang *et al.*, 2017], and self-supervised learning methods [Golan and El-Yaniv, 2018; Hendrycks *et al.*, 2019; Tack *et al.*, 2020; Winkens *et al.*, 2020], have been proposed for the task and have shown noticeable progress.

Among the numerous approaches mentioned, self-supervised learning (SSL) is in the limelight and validating its superiority over previous methods in various research areas

[Devlin *et al.*, 2018; Chen *et al.*, 2020a]. Since it is unfeasible to access *out-of-distribution* (OOD) data in most real-world scenarios, the ability of SSL to learn complex and diverse representations without additional labels is receiving much attention from anomaly detection lately. [Golan and El-Yaniv, 2018] is one of the earlier works to identify the potential of SSL and has proposed a simple yet effective technique that aims to learn intrinsic features within in-distribution (IND) samples via auxiliary tasks (*e.g.*, predicting flip, rotation, or translation of input data). Furthermore, [Hendrycks *et al.*, 2019] confirmed that using such auxiliary tasks not only helps to determine anomalous samples but also helps to defend against adversarial attacks. More recent works [Tack *et al.*, 2020; Winkens *et al.*, 2020] exploit contrastive learning (CL), especially SimCLR [Chen *et al.*, 2020a], that learns individual data representations in a *task-agnostic* way by maximizing the agreement between differently augmented views of the same image while repelling others in the batch.

SimCLR obtains effective individual representation for each data point, as well as clustered representations for each class, even without any human label or supervision. (See Fig. 1a.) However, its *task-agnostic* feature results in blurry boundaries between each cluster, so it requires a fine-tuning process used for some downstream tasks (*e.g.*, multi-class classification). Such process undermines expression ability well-learned through SimCLR, given that most of the fine-tuning procedure leverages *cross-entropy loss* and it solely considers class labels of the data without taking into account unique characteristics of the data or similarity between them. Consequently, the fine-tuned model often assigns high confidence probabilities to OOD input, reducing the distributional discrepancy between IND and OOD [Hein *et al.*, 2019].

Our foremost insight is that forming dense clusters without fine-tuning while preserving individual representations by inheriting the advantages of SimCLR will shape a more meaningful visual representation contrary to the *pre-train then tune* paradigm, thus contributing to the effective detection of anomalous data. To this end, we propose a *task-specific* variant of contrastive learning called *masked contrastive learning* (MCL), which can shape more clear boundaries between each class. (See Fig.1d) The core idea of MCL is to generate a mask that can adjust the repelling-ratio properly by considering class labels in the batch. Experimental results show that MCL is more befitted to anomaly detection than SimCLR or its other *task-specific* variant (*i.e.*, SupCLR), which still exhibits blurry decision boundaries (See Figure 1.1b).

Moreover, contrary to the previous belief that the auxiliary self-supervision task (*e.g.*, predicting flip, rotation, or translation of input data) does not substantially improve label classification accuracy [Hendrycks *et al.*, 2019], we observe that it is possible to considerably improve both IND and OOD performance with a proper inference method. To this end, we propose *self-ensemble inference* (SEI) that fully exploits ability learned from simple auxiliary self-supervision task in the inference phase. SEI enhances model performance in all situations without losing generality and can be used in any classifier. By combining our models, we can outperform previous state-of-the-art methods.

Our main contributions are summarized as below:

- We propose a novel extension to contrastive learning dubbed *masked contrastive learning* which can shape dense class-conditional clusters.
- We also propose an inference method called *self-ensemble inference* that fully leverages ability learned from auxiliary self-supervision tasks in test time. *Self-ensemble inference* can further boost both IND and OOD performance.
- We validate our approaches on various image benchmark datasets, where we obtain significant performance gain over the previous state-of-the-art.

The source code for our model is available online.<sup>2</sup>

## 2 Masked Contrastive Learning

As the name implies, our method adopts contrastive learning, particularly SimCLR, with two additional components: *class-conditional mask* and *stochastic positive attraction*; see Fig. 2. In this section, we provide detailed explanations of each component in MCL. (See Section 5 for further details regarding contrastive learning.)

### 2.1 Background: Contrastive Learning

Recent contrastive learning algorithms (*e.g.*, SimCLR) learn representations by maximizing the agreement between differently augmented views of the same image while repelling others in the batch. Specifically, each image  $x_k$  from randomly sampled batch  $\mathcal{B} = \{(x_k, y_k)\}_{k=1}^N$  is augmented twice, generating an independent pair of views  $(\bar{x}_{2k-1}, \bar{x}_{2k})$  and augmented batch  $\bar{\mathcal{B}} = \{(\bar{x}_k, \bar{y}_k)\}_{k=1}^{2N}$ , where labels of augmented views  $\bar{y}_{2k-1}, \bar{y}_{2k}$  are equal to original label  $y_k$ . The augmented pair of views,  $\bar{x}_{2k-1} = t(x_k)$  and  $\bar{x}_{2k} = t'(x_k)$ , are generated via independent transformation instance  $t$  and  $t'$ , drawn from pre-defined augmentation function family  $\mathcal{T}$ .  $(\bar{x}_{2k-1}, \bar{x}_{2k})$ , then are passed sequentially through encoder network  $f_\theta$  and projection head  $g_\phi$ , yielding latent vectors  $(z_{2k-1}, z_{2k})$  that are utilized for the contrastive loss (*i.e.*, NT-Xent):

$$\ell(i, j) = -\log \frac{\exp(\text{sim}_{i,j}/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(\text{sim}_{i,k}/\tau)}, \quad (1)$$

where  $\text{sim}_{i,j} = z_i^\top z_j / (\|z_i\| \|z_j\|)$  denotes cosine similarity between pair of latent vectors in  $(z_i, z_j)$  and  $\tau$  stands for temperature hyper-parameter. The final objective is to minimize Eq. 1 over positive pairs, which maps the input into effective individual representation in a *task-agnostic* way:

$$\mathcal{L}_{\text{SimCLR}} = \frac{1}{2N} \sum_{k=1}^N [\ell(2k-1, 2k) + \ell(2k, 2k-1)]. \quad (2)$$

### 2.2 Class-Conditional Mask

The benefit of CL in anomaly detection has been reported recently [Winkens *et al.*, 2020; Tack *et al.*, 2020]. Nonetheless, we found that well-formed representations from CL, which

<sup>2</sup><https://github.com/HarveyCho/MCL>

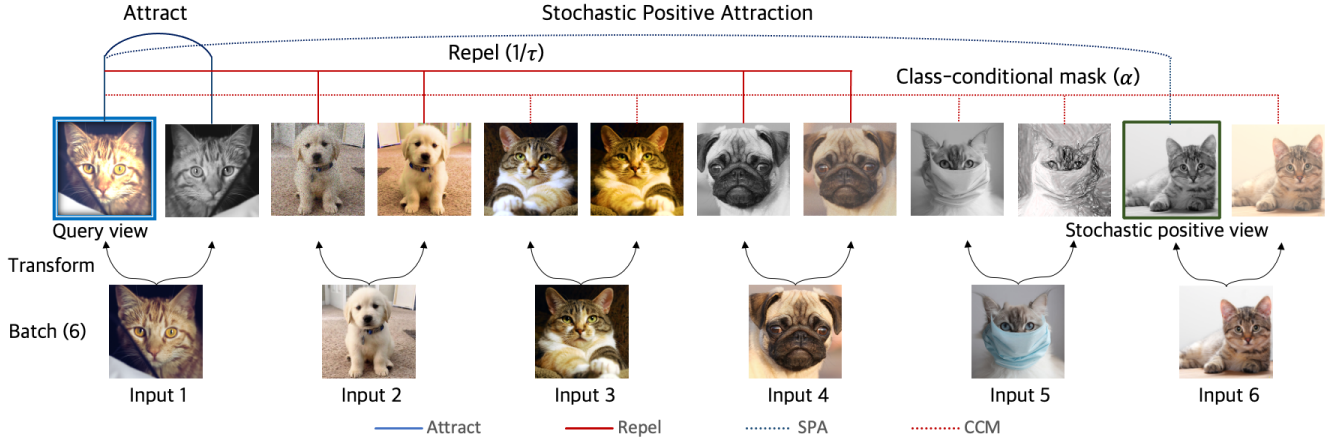


Figure 2: Description of MCL framework. Query view (gray-framed) attracts views connected with a blue-colored line and repels views connected with a red-colored line.

facilitate distinguishing anomalous data, are lost during the fine-tuning procedure. Due to the *task-agnostic* nature of CL, however, the fine-tuning steps are essential, making it difficult to avoid the aforementioned phenomenon. MCL mitigates such symptoms by injecting *task-specific* characteristics to existing CL, resulting in fine-tuning procedure inessential. One key component in MCL is *class-conditional mask* (CCM) which is a simple yet effective masking technique that adaptively determines the repelling-ratio considering the label information in each  $\bar{\mathcal{B}}$ . CCM can be defined as follows:

$$\text{CCM}(i, j) = \begin{cases} \alpha & \text{if } \bar{y}_i = \bar{y}_j \\ 1/\tau & \text{if } \bar{y}_i \neq \bar{y}_j, \end{cases} \quad (3)$$

where  $0 < \alpha < 1/\tau$ . CCM alters the temperature for the same label views to a smaller value  $\alpha$ , so that query view repels views with the same label relatively small amount compared to views with different labels. The generated CCM is then multiplied to the similarity score in Eq. 1, modifying the previous SimCLR loss to the following equation.

$$p_{\text{ccm}}(i, j) = \frac{\exp(\text{sim}_{i,j}/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(\text{sim}_{i,k} \text{CCM}(i, k))}, \quad (4)$$

$$\ell_{\text{ccm}}(i, j) = -\log p_{\text{ccm}}(i, j), \quad (5)$$

$$\mathcal{L}_{\text{ccm}} = \frac{1}{2N} \sum_{k=1}^N [\ell_{\text{ccm}}(2k-1, 2k) + \ell_{\text{ccm}}(2k, 2k-1)]. \quad (6)$$

Penalizing a small ratio  $\alpha$  to positive views restrains respective representation in the same cluster from being too similar to each other, making individual data representation more distinctive.

### 2.3 Stochastic Positive Attraction

As can be seen in Fig. 1c, CCM promotes a more label-wise cluster compared to SimCLR. Even in CCM, however, the core operating principle is still identical to SimCLR in that

it only attracts the view from the same image while it repels remaining views within the batch. Due to this repulsive nature, each data representation gets more distant as training continues, and it leads to the formation of unsatisfactory scattered clusters. To alleviate this phenomenon, we add another component named *stochastic positive attraction* (SPA), an additional attraction with the stochastically sampled view in the batch. Specifically, SPA attracts query  $\bar{x}_i$  with stochastic positive sample  $(\bar{x}_j, \bar{y}_j) \sim \mathcal{U}(\bar{\mathcal{B}}_i^+)$  in the positive augmented batch  $\bar{\mathcal{B}}_i^+$  for query  $\bar{x}_i$ , where  $\mathcal{U}$  refers to the discrete uniform distribution. The positive augmented batch  $\bar{\mathcal{B}}_i^+$  for query  $\bar{x}_i$  contains views with the same label except views from its parent image  $\mathbf{x}_{(i-1) \setminus 2}$ , where the symbol  $\setminus$  denotes inter quotient operator:

$$\bar{\mathcal{B}}_i^+ = \{(\bar{x}_k, \bar{y}_k) \in \bar{\mathcal{B}} \mid \bar{y}_k = \bar{y}_i \text{ and } (k-1) \setminus 2 \neq (i-1) \setminus 2\}. \quad (7)$$

CCM is also used for negative views with the additional constraint which excludes views from its parent image. SPA for query view  $\bar{x}_i$  now can be defined as follows:

$$p_{\text{spa}}(i, j) = \frac{\exp(\text{sim}_{i,j}/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[(k-1) \setminus 2 \neq (i-1) \setminus 2]} \exp(\text{sim}_{i,k} \text{CCM}(i, k))}, \quad (8)$$

$$\ell_{\text{spa}}(i) = E_{(\bar{x}_j, \bar{y}_j) \sim \mathcal{U}(\bar{\mathcal{B}}_i^+)} [-\log p_{\text{spa}}(i, j)]. \quad (9)$$

The complete version of MCL is acquired by combining CCM and SPA, where the overall loss term being as follows:

$$\mathcal{L}_{\text{MCL}} = \mathcal{L}_{\text{ccm}} + \frac{\lambda}{2N} \sum_{k=1}^{2N} \ell_{\text{spa}}(k), \quad (10)$$

where  $\lambda$  denotes weight hyper-parameter for SPA loss.

### 2.4 Training Auxiliary Task in MCL

Training simple auxiliary self-supervision task along with the main downstream task is possible in MCL by adding

constraint in CCM. Let  $T_{\text{main}}$  be the main task with  $C_{\text{main}}$  number of classes,  $T_{\text{aux}}$  be an auxiliary task with  $C_{\text{aux}}$  number of classes and corresponding augmented batch be  $\tilde{\mathcal{B}} = \{(\tilde{\mathbf{x}}_i, \tilde{y}_i^{\text{main}}, \tilde{y}_i^{\text{aux}})\}_{i=1}^{2N}$  with additional auxiliary task label. Then CCM with auxiliary self-supervision task can be defined as follows:

$$\text{CCM}_{\text{aux}}(i, j) = \begin{cases} \alpha & \text{if } \tilde{y}_i^{\text{main}} = \tilde{y}_j^{\text{main}} \text{ and } \tilde{y}_i^{\text{aux}} = \tilde{y}_j^{\text{aux}} \\ \beta & \text{if } \tilde{y}_i^{\text{main}} = \tilde{y}_j^{\text{main}} \text{ and } \tilde{y}_i^{\text{aux}} \neq \tilde{y}_j^{\text{aux}} \\ 1/\tau & \text{otherwise.} \end{cases} \quad (11)$$

By simply setting  $\beta = 1/\tau$ , it is possible to train  $C_{\text{main}} \times C_{\text{aux}}$  distinctive clusters for each  $(\tilde{y}_j^{\text{main}}, \tilde{y}_j^{\text{aux}})$  pairs. Since the auxiliary task plays a complementary role to the main task, it is more plausible to form grouped clusters for respective main labels and to have distinctive clusters for each auxiliary label inside them. With appropriate constraint (*i.e.*,  $0 < \alpha < \beta < 1/\tau$ ), MCL forms hierarchical clusters by dint of CCM.

### 3 Inference

#### 3.1 Scoring Function in MCL

Since there is no task-specific final layer in MCL, classification or anomaly detection are conducted via class-wise density estimation analogous to [Lee *et al.*, 2018], utilizing negative Mahalanobis distance  $-d_M$  as a scoring function  $s$ :

$$s_i(\mathbf{z}) = -d_M(\mathbf{z}, \mu_i; \Sigma_i) = (\mathbf{z} - \mu_i)^\top \Sigma_i^{-1} (\mathbf{z} - \mu_i), \quad (12)$$

$$S(\mathbf{x}) = [s_1(\mathbf{x}), s_2(\mathbf{x}), \dots, s_{C_{\text{main}}}(\mathbf{x})], \quad (13)$$

where  $\mathbf{z} = g_\phi(f_\theta(\mathbf{x}))$ , and  $\mu_i, \Sigma_i$  refer to mean and covariance matrix of  $n$ -dimensional multivariate normal distribution (MND)  $\mathcal{N}(\mu_i, \Sigma_i)$  for class  $i \in I = \{1, 2, \dots, C_{\text{main}}\}$ . Note that calculating MNDs for each class is a one-time operation acquired from training data. The vector  $S(\mathbf{x})$  contains scores of each label for image  $\mathbf{x}$  and the class label with highest score  $i^* = \arg\max_{n \in I} S_n(\mathbf{x})$  is selected as a predictive label, where  $S_n(\mathbf{x})$  denotes  $n$ -th element of  $S(\mathbf{x})$ . The corresponding IND score for predictive label  $s_{i^*}(\mathbf{x})$  measures the confidence for predictive label  $i^*$  which are used to distinguish OOD data, following the binary decision function  $h_\delta$  from below:

$$h_\delta(\mathbf{x}) = \begin{cases} \text{IND} & S_{i^*}(\mathbf{x}) \geq \delta \\ \text{OOD} & S_{i^*}(\mathbf{x}) < \delta, \end{cases} \quad (14)$$

where  $\delta$  denotes anomaly threshold.

#### 3.2 Self Ensemble Inference

The key idea of SEI is to exploit the model's ability to discriminate within IND, learned through an auxiliary self-supervision task, in the inference phase. For example, consider predicting 4-directional rotations (from  $0^\circ$ ,  $90^\circ$ ,  $180^\circ$ , to  $270^\circ$ ) is employed for an auxiliary task. Then SEI ensembles the results from corresponding the 4-rotated test images and derives calibrated index  $i^*$  and corresponding score  $s_{i^*}$ . Specifically, let  $i \in I = \{1, 2, \dots, C_{\text{main}}\}$  be the main task label, and  $j \in J = \{1, 2, \dots, C_{\text{aux}}\}$  be the auxiliary task label. Then,  $C_{\text{main}} \times C_{\text{aux}}$  number of MNDs  $\mathcal{N}(\mu_i^{(j)}, \Sigma_i^{(j)})$  are

calculated for every label combinations. The test image  $\mathbf{x}$  is augmented  $C_{\text{aux}}$  times, yielding  $\{\mathbf{x}^{(j)}\}_{j=1}^{C_{\text{aux}}}$ . Each augmented test image  $\mathbf{x}^{(m)}$  with the class label  $y^{\text{aux}} = m$  is fed into the corresponding MND  $\mathcal{N}(\mu_i^{(m)}, \Sigma_i^{(m)})$ , where  $i \in I$  and  $j = m$ , yielding a score vector  $S^{(m)}(\mathbf{x})$ :

$$s_i^{(j)}(\mathbf{x}) = -d_M(g_\phi(f_\theta(\mathbf{x})), \mu_i^{(j)}; \Sigma_i^{(j)}), \quad (15)$$

$$S^{(m)}(\mathbf{x}) = [s_1^{(m)}(\mathbf{x}), s_2^{(m)}(\mathbf{x}), \dots, s_{C_{\text{main}}}^{(m)}(\mathbf{x})]. \quad (16)$$

Our goal is to aggregate  $\{S^{(j)}(\mathbf{x})\}_{j=1}^{C_{\text{aux}}}$  properly to make model more robust and reliable. We considered 3 different aggregation methods to extract predictive label  $i^*$ . The foremost intuitive way is averaging the main label scores across  $\{S^{(j)}(\mathbf{x})\}_{j=1}^{C_{\text{aux}}}$ .

$$i_{\text{avg}}^*(\mathbf{x}) = \arg\max_{i \in I} \frac{1}{|J|} \sum_{m \in J} S_i^{(m)}(\mathbf{x}). \quad (17)$$

Another variation is to select label index from the highest IND score:

$$i_{\text{max}}^*(\mathbf{x}) = \arg\max_{i \in I} \left\{ \max_{m \in J} S_i^{(m)}(\mathbf{x}) \right\}. \quad (18)$$

The last aggregation is the weighted-average, which gives adaptive weights to each score in  $S$ . Weights for each score are computed per  $j$  using the harmonic mean to penalize exceptionally low scores for better calibration:

$$W^{(m)}(\mathbf{x}) = \left( \sum_{n \in I} \frac{1}{S_n^{(m)}(\mathbf{x})} \right)^{-1}. \quad (19)$$

$$i_{\text{w-avg}}^*(\mathbf{x}) = \arg\max_{i \in I} \frac{\sum_{m \in J} W^{(m)}(\mathbf{x}) S_i^{(m)}(\mathbf{x})}{\sum_{m \in J} W^{(m)}(\mathbf{x})}. \quad (20)$$

Corresponding score to predictive label  $i^*$  is used to distinguish OOD data following Eq. 14. Depending on the aggregation method, the effect that the model can achieve varies. Further details are elaborated in Section 4.2 with experimental results.

## 4 Experiment

In this section, we demonstrate the effectiveness of MCL and SEI on several multi-class image classification datasets.

**Experiment configurations.** In the following experiments, we adopt ResNet-34 [He *et al.*, 2016] with a single projection head, following structure used to train CIFAR-10 in SimCLR. We also fixed hyper-parameters related to contrastive learning following SimCLR, which include transformation  $\mathcal{T} = \{\text{color jittering, horizontal flip, grayscale, inception crop}\}$ , the strength of color distortion to 0.5, batch size to 1024, and temperature  $\tau$  to 0.2, to keep our experiment tractable. For MCL hyper-parameters, we set  $\alpha$  to 0.05,  $\beta$  to 2.5, and  $\lambda$  to 1 which meets certain condition for MCL (See Appendix A for details). Unlike SimCLR, we used SGD optimizer with learning rate 1.2 ( $0.3 \times \text{batch size} / 256$ ), decay  $1e-6$ , and momentum 0.9. Furthermore, we use a cosine annealing scheduler without any warm-up.

Training Method	Test Acc.	AUROC			
		SVHN	LSUN(F)	ImageNet(F)	CIFAR-100
Baseline [Hendrycks and Gimpel, 2016]	93.6	89.9	84.3	88.0	86.4
ODIN* [Liang <i>et al.</i> , 2017]	93.6	85.8	83.2	87.7	85.8
Mahalanobis Distance* [Lee <i>et al.</i> , 2018]	94.1	99.1	87.9	90.9	88.2
Auxiliary Rotation [Hendrycks <i>et al.</i> , 2019]	94.3	97.3	94.5	94.7	90.7
Outlier Exposure <sup>†</sup> [Hendrycks <i>et al.</i> , 2018]	-	98.4	-	-	93.3
SupCLR <sup>‡</sup> [Khosla <i>et al.</i> , 2020]	93.8	97.3	91.6	90.5	88.6
CSI [Tack <i>et al.</i> , 2020] + SupCLR <sup>‡</sup>	94.8	96.5	92.1	92.4	90.5
CSI + SupCLR + Ensemble <sup>‡</sup>	96.1	97.9	93.5	94.0	92.2
Auxiliary Rotation + SEI	95.8	98.4	95.7	95.8	92.3
MCL + SEI (ours) <sup>‡</sup>	95.9	98.9	96.0	95.9	93.1
<b>MCL + SEI (ours)</b>	<b>96.4</b>	<b>99.3</b>	<b>96.3</b>	<b>96.5</b>	<b>94.0</b>

\* refers models with additional post-processing procedure. † denotes models with additional supervised (OOD) data. ‡ denotes models that use ResNet-18 as a backbone network, and models without the mark use ResNet-34.

Table 1: Test accuracy of in-domain classification and AUROC of OOD data for each model trained on CIFAR-10 dataset.

Training Method	Test Acc.	AUROC			
		SVHN	LSUN(F)	ImageNet(F)	CIFAR-100
SupCLR [Khosla <i>et al.</i> , 2020]	93.8	97.3	91.6	90.5	88.6
MCL (Ours)	93.1	97.9	93.8	93.6	90.8
MCL + SEI (Ours)	95.9	98.9	96.0	95.9	93.1

Table 2: Performance comparison between MCL and SupCLR. Models are trained on CIFAR-10 dataset with ResNet-18.

**Evaluation metrics.** To evaluate IND detection performance, we measured the label classification accuracy. For OOD detection performance, we used the *area under the receiver operating characteristic curve* (AUROC), which is a threshold ( $\delta$  in Section 3.1) free metric and the most common metric in anomaly detection literature.

**Selecting self-supervision tasks.** Since the complex auxiliary task is not our primary concern, we considered rotation, horizontal flip, and translation as candidates for auxiliary tasks, which are simple and commonly used in the area of anomaly detection [Golan and El-Yaniv, 2018]. However, MCL contains inception crop [Szegedy *et al.*, 2015] and horizontal flip in contrastive transformation  $\mathcal{T}$ , so using translations or horizontal flip as an auxiliary task only confuses the model. To this end, we employed predicting 4-directional rotations ( $0^\circ$ ,  $90^\circ$ ,  $180^\circ$ , and  $270^\circ$ ) as our auxiliary task.

### 4.1 Multi-Class Anomaly Detection

We trained our model on CIFAR-10 [Krizhevsky *et al.*, 2009] as IND, and used CIFAR-100, SVHN [Netzer *et al.*, 2011], ImageNet [Deng *et al.*, 2009], and LSUN [Yu *et al.*, 2015] datasets for OOD. Note that all the classes in OOD datasets are disjoint with CIFAR-10. In particular, for ImageNet and LSUN dataset, we use ImageNet-Fix and LSUN-Fix datasets [Tack *et al.*, 2020] which are fixed versions of the previously released dataset [Liang *et al.*, 2017]. The previous version contains unintended artifacts caused by resizing large images. (See Appendix I in [Tack *et al.*, 2020] for more details.)

### Main Results

We compared MCL’s performance with several other methods in OOD detection. Table 1 summarizes our experimental results. MCL with SEI shows significant performance gain over previous methods in all setting; moreover, it outperforms supervised method [Hendrycks *et al.*, 2018], which utilizes additional explicit OOD data.

### Comparison with SupCLR

We compared MCL’s performance with SupCLR [Khosla *et al.*, 2020], another *task-specific* variant of CL. Table 2 summarizes performance comparison with SupCLR. In terms of test accuracy, SupCLR performs slightly better than MCL, while in terms of AUROC, MCL shows better performance. The superiority of MCL in detecting anomalies comes from several differences between the two frameworks. Specifically, SupCLR additionally attracts all same label views from  $\tilde{\mathcal{B}}_i^+$  in the augmented batch  $\tilde{\mathcal{B}}$ . Since SupCLR forces all positive views to have a high similarity to the query view, the ability to distinguish data with the same label diminishes. Unlike SupCLR, MCL penalizes a small amount  $\alpha$  to positive views, which endows an ability to discern each data with same label.

### 4.2 Ablation Study

In this section, we perform an ablation study on our proposed methods, along with baselines. In all experiments, we treated CIFAR-10 as IND and CIFAR-100 as OOD.

### Masked Contrastive Learning

We conduct ablation experiments to explore the effectiveness of the main components (CCM, SPA, and auxiliary 4-way

Model	Acc	AUROC
Baseline (w/o NT-Xent)	93.61	86.40
SimCLR (CE fine-tuned)	93.88	87.56
SimCLR (joint fine-tuned)	93.91	88.51
MCL (w/o SPA, w/o aux)	91.41	85.03
MCL (w/o aux)	<b>94.35</b>	89.49
MCL (CE fine-tuned)	94.22	88.89
<b>MCL</b>	94.03	<b>91.12</b>

Table 3: Ablation studies for each component in MCL.

Model	Agg	Acc	AUROC
MCL + w/o SEI	-	94.03	91.12
MCL + 4-way SEI	avg	94.68	93.20
	max	95.97	92.30
	w-avg	96.12	93.29
<b>MCL + 8-way SEI</b>	avg	94.74	93.37
	max	96.40	92.00
	<b>w-avg</b>	<b>96.43</b>	<b>94.06</b>

Table 4: Ablation studies for SEI. MCL is trained with additional 4-way rotations auxiliary task.

rotation task) in MCL. Tab. 3 reports our ablation experiments along with baseline models. For SimCLR based models, we fine-tuned the pre-trained model in two ways. One way is to use *cross-entropy loss* which is a traditional fine-tuning method in the classification task, and the other way is to use *cross-entropy loss* along with SimCLR loss (Eq.2) jointly [Winkens *et al.*, 2020]. The two methodologies show almost the same accuracy, while there is a slight difference in AUROC. Our justification for this phenomenon is due to the nature of the *cross-entropy loss*, which solely reflects the class label leading diminish in the distributional discrepancy between IND and OOD. We conjectured that the additional SimCLR loss in joint fine-tuning mitigates this phenomenon and shows better AUROC performance. This phenomenon is more evident when applied to MCL, showing substantial performance degradation in AUROC. Therefore, we used MCL with neither a fine-tuning procedure nor any *task-specific* layer on the top level of the network.

### Self-Ensemble Inference 1

In this part, we share our findings on SEI with its variations (average, maximum, and weighted-average). Since we employed a 4-directional rotation prediction as our auxiliary task, SEI is done in a 4-way correspondingly. It is also possible to add additional 4-way SEI with a horizontally flipped image, which we dubbed 8-way SEI. (Fig. 3 provides a visual explanation.) 8-way SEI follows the same strategy introduced earlier. The only difference is the number of augmented images to ensemble. As claimed in [Hendrycks *et al.*, 2019], learning the auxiliary task alone can not improve accuracy. But with the help of SEI, we were able to achieve performance gains in both accuracy and AUROC regardless of its variations as can be seen in Tab. 4. Interestingly, the

Model	Agg	Acc	AUROC
MCL + w/o Aux	-	94.35	90.49
	w-avg	83.96	71.17
MCL + DA	-	92.55	90.07
	w-avg	94.70	92.08
<b>MCL + Aux</b>	-	94.03	91.12
	<b>w-avg</b>	<b>96.43</b>	<b>94.06</b>

Table 5: Ablation studies for SEI. data augmentation is abbreviated to DA. Symbol - indicates a model without SEI.

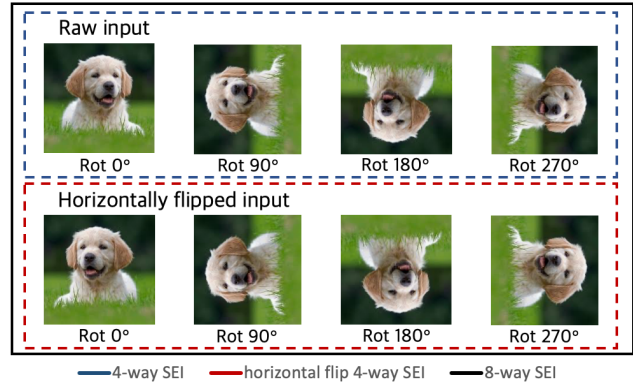


Figure 3: 4-way SEI and 8-way SEI.

effect of each ensemble differs depending on the aggregation method. For example, average SEI makes the model robust to input variation which is a commonly known benefit of the ensemble, bringing noticeable gain in AUROC contrary to accuracy. On the other hand, maximum SEI yields a significant gain in accuracy, which indicates MCL’s prediction with high confidence score is quite precise. The weighted-average SEI absorbs the advantages of both ensembles by assigning adaptive weights to its score. As a side note, the SEI can be used for the general classifier trained with an auxiliary task, which also yields significant performance gains as can be seen in Tab. 1.

### Self-Ensemble Inference 2

We also conduct ablation experiments on SEI and the relationship with the auxiliary task. We applied SEI to 3 differently trained model as follows:

- **MCL without auxiliary task:** MCL is trained neither with auxiliary task nor data augmentation
- **MCL with data augmentation:** MCL trained with additional 4-way rotated images without rotation labels.
- **MCL with auxiliary task:** MCL is trained with additional 4-way rotated images with rotation labels.

Tab. 5 summarizes our extended ablation study for SEI. Applying SEI to MCL without auxiliary task only confuses the model and undermines both IND and OOD performance substantially, as rotated images are unseen during training phase. If the rotated image is augmented without a label, the model performance is degraded but the SEI shows a slight

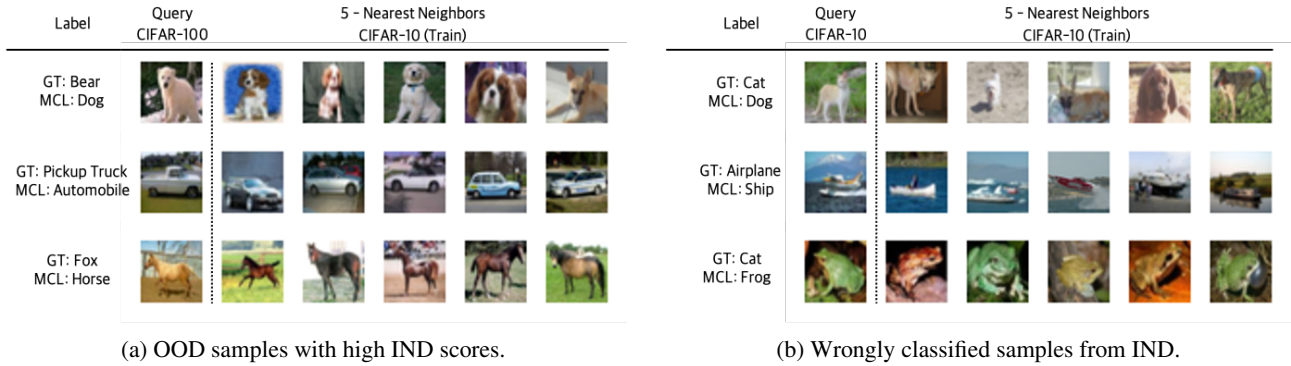


Figure 4: Case studies on OOD samples with high confidence and wrongly classified IND samples.

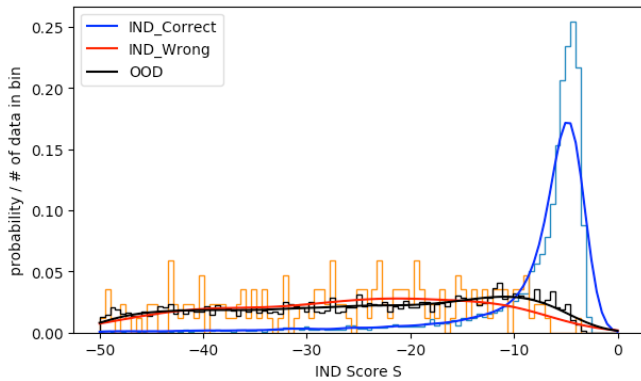


Figure 5: Histogram and PDF for correctly classified IND, wrongly classified IND and OOD distribution.

improvement. Finally, when the auxiliary task is explicitly trained with the main downstream task, SEI shows significant gain without any performance degradation. Such experimental results reveal that it is necessary to additionally train the auxiliary task to use SEI properly.

### 4.3 Qualitative Results

In this section, we analyze the data distribution along with several failure cases of our model. As can be seen in Fig. 5, both OOD data distribution and wrongly classified data distribution have lower scores compared to correctly classified samples, which indicates that our model can measure predictive uncertainty quite precisely. Furthermore, to analyze our failure cases in detail, we conducted a case study on OOD samples with high confidence and wrongly classified IND samples. In MCL, representation is learned based on semantic similarity, so it is feasible to conjecture model’s decision-making process via observing nearest neighbors of the input. For most failure cases, their predicted label and ground-truth label (GT) belongs to the same super-class regardless of the aforementioned failure types. For example, MCL predicts a bear (OOD) as a dog (IND); likewise, a cat as a dog (both IND), which belongs to the same super-class, mammal. As a side note, few data were mislabeled, as can be seen in Fig. 4.(a).(GT: Fox)

## 5 Related Work

**Anomaly detection.** Recent approaches in OOD can be categorized as follows: reconstruction based [Oza and Patel, 2019; Li *et al.*, 2018], density estimation based [Malinin and Gales, 2018], post-processing based [Lee *et al.*, 2018; Liang *et al.*, 2017], and self-supervised learning based. Self-supervised learning based methods can be split again into auxiliary self-supervision based [Golan and El-Yaniv, 2018; Hendrycks *et al.*, 2019] and contrastive learning based [Tack *et al.*, 2020; Winkens *et al.*, 2020]. Our method belongs to self-supervised learning, which exploits both auxiliary self-supervision task and contrastive learning.

**Contrastive learning.** Contrastive learning is a specific framework of self-supervised learning, which has shown impressive results in visual representation learning tasks [Chen *et al.*, 2020a; Chen *et al.*, 2020b]. Most recent work in OOD [Tack *et al.*, 2020; Winkens *et al.*, 2020] report that employing CL improves OOD performance. Our work goes further from the previous papers and proposes a *task-specific* variant of CL. [Khosla *et al.*, 2020] proposed SupCLR, another *task-specific* variant of SimCLR, which is a noteworthy work. Similar to MCL, SupCLR also leverages label information in the batch while training and shows superior accuracy over SimCLR. Despite its performance in IND accuracy, representation from SupCLR which is not entirely appropriate for discerning anomalous data, as it attracts all same label views which discrepancy in each class disappears.

## 6 Conclusion

In this paper, we propose a novel training method called *masked contrastive learning* (MCL) and an inference method called *self-ensemble inference* (SEI). MCL can shape class-conditional clusters by inheriting advantages of CL and SEI fully leverages trained features from auxiliary self-supervised tasks in the inference phase. By combining our methods, our model reaches the new state-of-the-art performance.

## Acknowledgements

We thank Taeuk Kim, Ye-seul Song and the anonymous reviewers for their thoughtful feedback and comments.

## References

- [Chen *et al.*, 2020a] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [Chen *et al.*, 2020b] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020.
- [Deng *et al.*, 2009] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [Devlin *et al.*, 2018] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [Golan and El-Yaniv, 2018] Izhak Golan and Ran El-Yaniv. Deep anomaly detection using geometric transformations. In *Advances in Neural Information Processing Systems*, pages 9758–9769, 2018.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [Hein *et al.*, 2019] Matthias Hein, Maksym Andriushchenko, and Julian Bitterwolf. Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 41–50, 2019.
- [Hendrycks and Gimpel, 2016] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- [Hendrycks *et al.*, 2018] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606*, 2018.
- [Hendrycks *et al.*, 2019] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. In *Advances in Neural Information Processing Systems*, pages 15663–15674, 2019.
- [Khosla *et al.*, 2020] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- [Lee *et al.*, 2018] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Processing Systems*, pages 7167–7177, 2018.
- [Li *et al.*, 2018] Dan Li, Dacheng Chen, Jonathan Goh, and See-kiang Ng. Anomaly detection with generative adversarial networks for multivariate time series. *arXiv preprint arXiv:1809.04758*, 2018.
- [Liang *et al.*, 2017] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. *arXiv preprint arXiv:1706.02690*, 2017.
- [Malinin and Gales, 2018] Andrey Malinin and Mark Gales. Predictive uncertainty estimation via prior networks. In *Advances in Neural Information Processing Systems*, pages 7047–7058, 2018.
- [Netzer *et al.*, 2011] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [Oza and Patel, 2019] Poojan Oza and Vishal M Patel. C2ae: Class conditioned auto-encoder for open-set recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2307–2316, 2019.
- [Szegedy *et al.*, 2015] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [Tack *et al.*, 2020] Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. Csi: Novelty detection via contrastive learning on distributionally shifted instances. *arXiv preprint arXiv:2007.08176*, 2020.
- [Tan and Le, 2019] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019.
- [Tan *et al.*, 2020] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10781–10790, 2020.
- [Winkens *et al.*, 2020] Jim Winkens, Rudy Bunel, Abhijit Guha Roy, Robert Stanforth, Vivek Natarajan, Joseph R Ledsam, Patricia MacWilliams, Pushmeet Kohli, Alan Karthikesalingam, Simon Kohl, et al. Contrastive training for improved out-of-distribution detection. *arXiv preprint arXiv:2007.05566*, 2020.
- [Yu *et al.*, 2015] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.