

Class Anchor Clustering: A Loss for Distance-based Open Set Recognition

Dimity Miller, Niko Sünderhauf, Michael Milford, Feras Dayoub
 Queensland University of Technology, Australian Centre for Robotic Vision

{d24.miller, niko.suenderhauf, michael.milford, feras.dayoub}@qut.edu.au

Abstract

In open set recognition, deep neural networks encounter object classes that were unknown during training. Existing open set classifiers distinguish between known and unknown classes by measuring distance in a network’s logit space, assuming that known classes cluster closer to the training data than unknown classes. However, this approach is applied post-hoc to networks trained with cross-entropy loss, which does not guarantee this clustering behaviour. To overcome this limitation, we introduce the Class Anchor Clustering (CAC) loss. CAC is a distance-based loss that explicitly trains known classes to form tight clusters around anchored class-dependent centres in the logit space. We show that training with CAC achieves state-of-the-art performance for distance-based open set classifiers on all six standard benchmark datasets, with a 15.2% AUROC increase on the challenging TinyImageNet, without sacrificing classification accuracy. We also show that our anchored class centres achieve higher open set performance than learnt class centres, particularly on object-based datasets and large numbers of training classes.

1. Introduction

Many practical applications require the deployment of trained visual perception models under *open set* conditions, such as autonomous systems, driverless cars, and robotics. In open set conditions, a model encounters object classes that were not present during training (referred to as ‘unknown’ classes) [20]. Deep convolutional neural networks (CNNs) degrade in performance in open set conditions, as they can confidently misclassify unknown classes as known training classes [4, 9, 15]. This behaviour raises serious concerns about the safety of using CNNs in open set environments [1] – particularly on autonomous systems where per-

ception failures may have severe consequences [4, 24].

Open set recognition extends object recognition to an open set environment [20]. During testing, an open set classifier must classify known object classes and reject unknown object classes [20]. In this paper, we propose a new distance-based loss that achieves state-of-the-art performance for distance-based open set recognition.

Many open set classifiers model the position of known training data in the final layer, or logit space, of a CNN [2, 28, 29]. Such approaches assume known classes cluster tightly in the logit space, and that unknown classes will maintain a distance from these clusters. Figure 1a shows this ideal performance. Currently, this concept is applied to networks trained with cross-entropy loss [2, 28, 29]. However, cross-entropy loss does not guarantee the clustering behaviour these methods seek to exploit. We exhibit this in Figure 1b, where we train a CNN with cross-entropy loss to classify trains, buses, and bicycles (CIFAR100 classes). The resulting logit space of this CNN appears crowded with inflated class clusters, and it is challenging to distinguish the unknown classes (bear and possum) from these clusters.

In this work, we introduce the Class Anchor Clustering (CAC) loss to address this limitation in prior work. CAC is a distance-based loss that explicitly encourages known training data to form tight clusters around anchored, class-specific centres in the logit space. CAC is compatible with existing classification networks, with only slight modifications to the network architecture. Compared to the cross-entropy trained CNN, a logit space trained with CAC exhibits tight, separate class clusters and an improved distinction of these clusters from unknown classes (see Figure 1c).

Our paper makes the following contributions:

1. We propose a new loss term for open set recognition that encourages known class training data to cluster tightly around class-specific centres in logit space.
2. We show that training with this novel *Class Anchor Clustering* (CAC) loss achieves new state-of-the-art open set performance for distance-based open set classifiers, without sacrificing classification accuracy.
3. We introduce the concept of *anchored* class centres as

The authors acknowledge continued support from the Queensland University of Technology (QUT) through the Centre for Robotics. This research was conducted by the Australian Research Council Centre of Excellence for Robotic Vision (project number CE140100016).

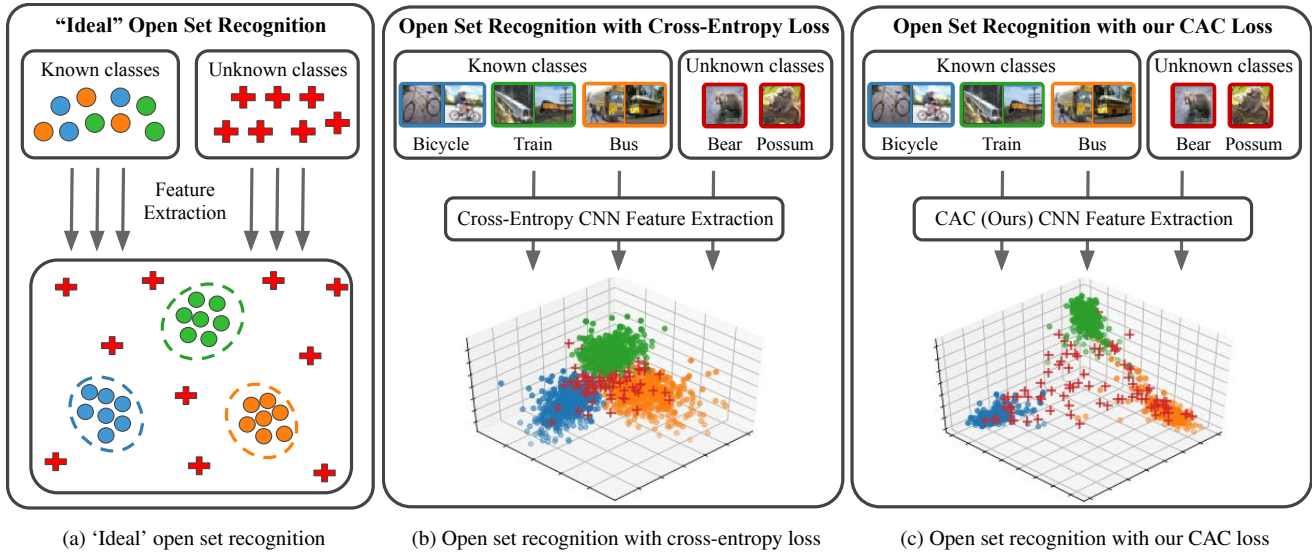


Figure 1: Left: An ‘ideal’ open set classifier will tightly cluster known classes in the feature space, and unknown classes will fall far away. Middle: A CNN trained on real image data from CIFAR100 with cross-entropy loss shows a final 3-D logit space with inflated known class clusters which cannot be easily distinguished from the unknown classes. Right: A CNN trained with our proposed CAC loss (on the same CIFAR100 data) shows a final 3-D logit space with tight, separated known class clusters and improved distinction of unknown classes.

an effective and scalable strategy for distance-based training. In contrast to learning class centres during training, distance losses using anchored centres perform well with object-based datasets with high intra-class variation and large numbers of classes.

2. Related Work

Open set recognition Open set recognition is multiclass classification with the additional requirement of rejecting inputs from unknown classes [20]. This is formalised as the task of minimising open space risk, the portion of classification space labelled as ‘known’ that is far from the known training data, while maintaining generalisation and classification accuracy on the known classes [20]. Related areas, such as out-of-distribution and novelty detection, exist as relaxed forms of open set classification where known and unknown classes are from different distributions [9] or multiclass classification is not required [5]. For this work, we focus specifically on open set recognition.

OpenMax was one of the first CNN open set classifiers, using the network’s final layer’s logits, or logit space, as the classification space with open space risk [2]. OpenMax models each known class as a single cluster, and uses a Weibull distribution to re-calibrate softmax scores based on an input’s distance to each cluster centre. OpenMax was the first ‘distance-based’ approach, using distance from the training data to minimise the open space risk of a CNN.

Several following works employed real or generated ‘known unknown’ data to augment the training dataset, either using the data to improve the feature representation for distance-based measures [29] or to bound the known classification space with an ‘other’ class [13, 21]. In [7], a network is trained to produce low feature magnitudes and uniform confidence scores for ‘known unknown’ data. While feature magnitudes are used for open set recognition, this method does not model or measure class-specific distances in the logit space [7] and therefore we do not define it as a distance-based method.

Other recent open set classifiers use a combined classifier and autoencoder network architecture [16, 28]. In [16], the reconstruction error from a class-conditioned autoencoder-classifier is used to distinguish between known and unknown inputs. Others [6, 22] observed that reconstruction error alone is not suitable as a measure of class novelty. In contrast, [28] jointly applies OpenMax to a classifier logit space and auto-encoder latent space, with the additional reconstruction-learned features improving the overall feature representation. Another approach [18] uses a self-supervision loss with random transformations to learn a more descriptive feature representation [18]. Additionally, the input to the network is augmented with their reconstruction from an autoencoder to further enable open set recognition [18].

In contrast to existing distance-based open set classifiers [2, 28, 29], which *assume* known classes will tightly

cluster but train with cross-entropy loss, our work is the first to train with a distance loss when using distance in the logit space for testing.

Distance Losses for Deep Neural Networks The field of metric learning uses distance loss functions to learn meaningful feature embeddings. Triplet loss is a popular distance loss that encourages inputs to minimise distance to a ‘positive’ example and maximise distance to a single ‘negative’ example. Tuplet loss was introduced as an extension of triplet loss that maximises an input’s distance to *multiple* ‘negative’ examples [23]. We adopt a modified version of Tuplet loss as one of two terms in our new CAC loss but show that Tuplet loss alone is not sufficient for best performance.

Center Loss [26] was proposed to improve discriminative learning for facial recognition by encouraging clustering in a feature space. It is used in conjunction with cross-entropy loss and encourages an input to minimise distance to its ground truth class centre. The class centres are learnt simultaneously with the feature embedding during training. In contrast, we propose to use *anchored*, i.e. fixed, class centres. This makes training more stable and, as we will show, more scalable to larger and more complex datasets.

Recently, [12] demonstrated the utility of metric learning for open set classification, however only for fine-grained image classification. Such metric learning approaches compute distances between individual instances of the training data, and the sampling technique used can have a significant effect on the convergence speed and stability of the training minimum [27]. As discussed in [19], this sampling typically makes metric learning computationally intractable on larger datasets, such as CIFAR10, CIFAR100, or ImageNet. Although recent work [19] adapted metric learning approaches for large-scale datasets, this technique degraded the classification accuracy of a standard cross-entropy network.

3. Class Anchor Clustering (CAC) for Open Set Classification

We now introduce the two core ideas of our paper that enable distance-based training for large-scale image open set classification: (1) the Class Anchor Clustering (CAC) loss that encourages training data to form tight, class-specific clusters. Tight clusters make it easier to distinguish between known and unknown class inputs during deployment. (2) the concept of using *anchored* class centres in the logit space to fix cluster centre positions during training.

Before introducing our approach, we briefly explain why cross-entropy loss is not sufficient for distance-based open set recognition. Cross-entropy loss minimises the negative log-probability of an input’s ground truth class, which is

obtained by normalising the logits with a softmax function. The softmax function is *not injective* and multiple input logit vectors map to the same output softmax probability vector [8]. As a result, cross-entropy loss cannot guarantee clustering in the logit space.

General Architecture CAC is compatible with existing classification networks, with slight modifications to the architecture. Our proposed CAC-trained open set classifier has three main components:

1. A base network, f , that projects an input image \mathbf{x} to a vector of class logits $\mathbf{z} = f(\mathbf{x})$. This network can be any existing classifier with an N -dimensional logit space, where N is the number of known classes.
2. A non-trainable parameter, \mathbf{C} , representing a set of class centre points $(\mathbf{c}_1, \dots, \mathbf{c}_N)$, one for each of the N known classes.
3. A new layer, $e(\mathbf{z}, \mathbf{C})$, that calculates \mathbf{d} , a vector of Euclidean distances between a logit vector \mathbf{z} and the set of class centres \mathbf{C} .

In summary, the output of our distance-based classifier is

$$\mathbf{d} = e(\mathbf{z}, \mathbf{C}) = (\|\mathbf{z} - \mathbf{c}_1\|_2, \dots, \|\mathbf{z} - \mathbf{c}_N\|_2)^\top \quad (1)$$

where $\|\cdot\|_2$ denotes the Euclidean norm.

3.1. Training with a Distance-based Loss Function

During training, we wish to learn a logit space embedding $f(\mathbf{x})$ where known inputs form tight, class-specific clusters. This clustering enables us to use a distance-to-class-centre metric during testing to reject unknown class inputs and classify known class inputs.

3.1.1 Class Anchor Clustering Loss

We require a distance-based loss that a) encourages training inputs to minimise the distance to their ground-truth class centre, while b) maximising the distance to all other class centres to encourage discriminative learning.

To do this, we use a modified Tuplet loss term \mathcal{L}_T [23] that forces an input \mathbf{x} to maximise the difference in distance to the correct class centre \mathbf{c}_y and all other class centres. Remembering that $\mathbf{d} = (d_1, \dots, d_N)^\top$ is defined as in (1), we define this loss component as

$$\mathcal{L}_T(\mathbf{x}, y) = \log \left(1 + \sum_{j \neq y}^N e^{d_y - d_j} \right). \quad (2)$$

\mathcal{L}_T differs from Tuplet loss [23] because it is based on class centres \mathbf{C} rather than sampled class instances. Our modified Tuplet loss term is equivalent to cross-entropy loss

applied to the distance vector \mathbf{d} , but used with a softmin function rather than softmax (see supplementary material for proof). The softmin function is the opposite of softmax: it assigns a large value (≈ 1) to the *smallest* value of the input vector and is defined as:

$$\text{softmin}(\mathbf{d})_i = \frac{e^{-\mathbf{d}_i}}{\sum_{k=1}^N e^{-\mathbf{d}_k}}. \quad (3)$$

While effective for discriminative learning, \mathcal{L}_T aims to maximise the *margin* between distance to the correct class centre and distance to the incorrect class centres. To ensure an input is explicitly forced to lower its *absolute* distance to the correct class centre, we also penalise the Euclidean distance between the training logit and the ground truth class centre. We refer to this as the Anchor loss term:

$$\mathcal{L}_A(\mathbf{x}, y) = d_y = \|f(\mathbf{x}) - \mathbf{c}_y\|_2. \quad (4)$$

We combine the Anchor and Tuplet loss terms to form our final distance-based loss, which we refer to as the Class Anchor Clustering (CAC) loss:

$$\mathcal{L}_{CAC}(\mathbf{x}, y) = \mathcal{L}_T(\mathbf{x}, y) + \lambda \mathcal{L}_A(\mathbf{x}, y). \quad (5)$$

A hyperparameter of our method is λ , which balances these two individual loss terms (explored in section 5.2.3). By combining the Anchor and Tuplet loss terms, our loss minimises training inputs distance to their ground-truth anchored class centre, while maximising the distance to other anchored class centres.

3.1.2 Anchored Class Centres

We introduce anchored class centre points as a method of *anchoring*, i.e. fixing, cluster centres for each class in the logit space during training. By anchoring our class centres during training, we eliminate the need to learn another parameter (as done in previous approaches to distance losses, e.g. [26]).

For each known class i , our network has an anchored class centre \mathbf{c}_i in the logit space. Given an N -dimensional logit space for N known classes, we place the anchored centre for each known class at a point on its class coordinate axis. This anchored centre point is therefore equivalent to a scaled standard basis vector \mathbf{e}_i , or scaled one-hot vector, for each class. The magnitude of the anchored centre, α , is a hyperparameter of our method (explored in section 5.2.3). We summarise this below:

$$\mathbf{C} = (\mathbf{c}_1, \dots, \mathbf{c}_N) = (\alpha \cdot \mathbf{e}_1, \dots, \alpha \cdot \mathbf{e}_N) \quad (6)$$

$$\mathbf{e}_1 = (1, 0, \dots, 0)^\top, \quad \mathbf{e}_N = (0, \dots, 0, 1)^\top. \quad (7)$$

After completing training, the anchored class centre positions \mathbf{C} are adjusted to the mean position of the correctly

classified training data. This allows us to model the class cluster centres for complex datasets more accurately, where visual and semantic similarities between classes can cause slight divergence from the original anchored class centre positions. Note that while anchoring our class centres equal distances apart in the logit space may limit the learning of semantically meaningful features, for this work we aim to learn only a discriminative feature representation that exhibits tight clustering behaviour.

3.2. Using Distance-based Measures during Testing

During testing, the network has to reject unknown class inputs and correctly classify known class inputs. Our CAC loss trains known inputs to have two distance-based properties: (1) a high softmin score for the distance to correct known class centre (as per the modified Tuplet loss term \mathcal{L}_T) and (2) a low absolute distance to the correct known class centre (as per the Anchor loss term \mathcal{L}_A). Based on this, we calculate rejection scores $\gamma = (\gamma_1, \dots, \gamma_N)^\top$ that express the classifier’s *disbelief* that the input \mathbf{x} belongs to each of the N known classes. We calculate the rejection scores γ as the element-wise product (\circ) of the distance vector \mathbf{d} and its inverted softmin:

$$\gamma = \mathbf{d} \circ (1 - \text{softmin}(\mathbf{d})) \quad (8)$$

By weighting the absolute distance with the inverted softmin score, inputs must have both a low absolute distance and high softmin score to be assigned a low rejection score for a known class. If all values in γ are above a threshold θ , the input does not belong to any known class and is rejected as unknown. Otherwise, the class label corresponding to the smallest value in γ is assigned:

$$\text{decision} = \begin{cases} \text{rejected as unknown} & \text{if } \min(\gamma) > \theta \\ \text{class } i = \text{argmin } \gamma & \text{if } \min(\gamma) \leq \theta \end{cases} \quad (9)$$

Using this distance-based decision procedure minimises open space risk [20]: the further away an input \mathbf{x} projects from the class-specific centres, the more likely it is to be rejected as unknown.

4. Experimental Setup

To simulate open set conditions, a set of ‘known’ and ‘unknown’ classes must be established. The *openness* \mathbf{O} of the classification task [20] can then be defined as

$$\mathbf{O} = 1 - \sqrt{\frac{2 \cdot N_{\text{train}}}{N_{\text{test}} + N_{\text{target}}}} \quad (10)$$

where N_{train} is the number of ‘known’ classes during training, N_{target} is the number of ‘known’ classes for classification during testing and N_{test} is the total number of testing classes (‘known’ and ‘unknown’). A higher openness

indicates a more difficult problem setup, but other factors such as the visual similarity between known and unknown classes also influence the difficulty. We follow the established benchmark evaluation protocol [13], where standard classification datasets are adapted to the open set task by randomly splitting into ‘known’ and ‘unknown’ classes.

4.1. Datasets

The details of each dataset in its open set configuration are summarised below. For each dataset, performance is evaluated over 5 trials with random known and unknown class splits.

MNIST [11]: grayscale 32×32 images of handwritten digits, 6 known and 4 unknown classes, $\mathbf{O} = 13.39\%$.

SVHN [14]: RGB 32×32 images of street view house digits, 6 known and 4 unknown classes, $\mathbf{O} = 13.39\%$.

CIFAR10 [10]: RGB 32×32 images of animals and objects, 6 known and 4 unknown classes, $\mathbf{O} = 13.39\%$.

CIFAR+10/+50: considers the 4 non-animal classes of CIFAR10 as known, and 10 or 50 randomly sampled animal classes from CIFAR100 [10] as unknown ($\mathbf{O} = 33.33\%$ and 62.86%).

TinyImageNet [25]: RGB 64×64 images of animals and objects, 20 known and 180 unknown classes, $\mathbf{O} = 57.35\%$. TinyImageNet images can contain significant background information unrelated to the object class, a number of classes are very visually and semantically related (e.g. different breeds of dogs), and there is high visual variation within individual classes. Examples are provided in the supplementary material.

4.2. Metrics

We use the following metrics to assess the performance of an open set classifier.

Area Under the ROC Curve (AUROC) is a calibration-free measure of the open set performance of a classifier. The Receiver Operating Characteristic (ROC) curve represents the trade-off between true positive rate (known inputs correctly retained as ‘known’) and the false positive rate (unknown inputs incorrectly retained as ‘known’) when applying varying thresholds to a given score. We modify the threshold θ that is compared to our network’s rejection scores γ as discussed in (9).

Classification Accuracy measures the classifier’s accuracy when applied to only the known classes in the dataset, equivalent to closed set classification. An open set classifier should maintain the classification accuracy of a standard closed set classifier.

Correct Classification Rate (CCR) measures the fraction of known inputs that are correctly retained as ‘known’ and correctly classified as their ground truth class [7]. It can be computed for various false positive rates when varying a threshold (we threshold our network’s rejection scores γ).

4.3. State-of-the-art Methods for Comparison

We compare to seven existing state-of-the-art open set classifiers [2, 9, 13, 16, 17, 28, 29], following the established protocol that uses AUROC as the primary evaluation metric. The results of these experiments are shown in Table 1, where we also highlight the three methods [2, 28, 29] that use distance in the logit space during *testing* to distinguish between known and unknown inputs. In contrast to CAC, none of those methods uses distance during *training*.

We additionally compare to another state-of-the-art open set classifier [7] that uses a different evaluation protocol (as described in Section 5.1). Instead of assessing the overall performance with AUROC, [7] evaluate with CCR, which quantifies performance at specific false positive operating points. This is an important metric for many safety-critical domains where performance at reasonable low false positive rates is relevant. We hope that including this comparison encourages future open set works to do so also.

4.4. Implementation Details

We use the network architecture specified by the benchmark evaluation protocol [13]. We use a Stochastic Gradient Descent (SGD) optimizer with a learning rate of 0.01 and train until convergence. We then complete another training cycle with a lower learning rate of 0.001 and train again until convergence. More details about the training procedure are in the supplementary material. For all datasets, we use an Anchor loss weight λ of 0.1 and a logit anchor magnitude α of 10.

5. Results and Discussion

Our evaluation revealed four main results that we discuss in the following: (1) CAC outperforms the existing distance-based open set classifiers [2, 28, 29] on every tested dataset, without sacrificing classification accuracy (Section 5.1). (2) Compared to other distance losses, CAC achieves better open set performance (Section 5.2.1). (3) Training with *anchored* class centres achieves better open set performance than learnt class centres on nearly all tested datasets, particularly on object-based datasets with high intra-class visual variations. Anchored centres also maintain open set performance better with increasing numbers of known classes (Section 5.2.2). (4) Training with CAC is insensitive to the choice of its two hyperparameters over a wide range of values (Section 5.2.3).

5.1. Comparison with State-of-the-Art Open Set Classifiers

The open set performance of our proposed approach is compared to the state-of-the-art methods in Table 1.

Method	Use Distance in		MNIST	SVHN	CIFAR10	CIFAR+10	CIFAR+50	TinyImageNet
	Training	Testing						
Softmax [9]	✗	✗	97.8 ± 0.6	88.6 ± 1.4	67.7 ± 3.8	$81.6 \pm \text{N.R.}$	$80.5 \pm \text{N.R.}$	$57.7 \pm \text{N.R.}$
OSRCI [13]	✗	✗	98.8 ± 0.4	91.0 ± 1.0	69.9 ± 3.8	$83.8 \pm \text{N.R.}$	$82.7 \pm \text{N.R.}$	$58.6 \pm \text{N.R.}$
C2AE [16]	✗	✗	-	89.2 ± 1.3	71.1 ± 0.8	81.0 ± 0.5	80.3 ± 0.0	58.1 ± 1.9
GFROR [17]	✗	✗	-	93.5 ± 1.8	80.7 ± 3.9	92.8 ± 0.2	92.6 ± 0.0	60.8 ± 1.7
OpenMax [2]	✗	✓	98.1 ± 0.5	89.4 ± 1.3	69.5 ± 4.4	$81.7 \pm \text{N.R.}$	$79.6 \pm \text{N.R.}$	$57.6 \pm \text{N.R.}$
G-OpenMax [29]	✗	✓	98.4 ± 0.5	89.6 ± 1.7	67.5 ± 4.4	$82.7 \pm \text{N.R.}$	$81.9 \pm \text{N.R.}$	$58.0 \pm \text{N.R.}$
CROSR [28]	✗	✓	99.1 ± 0.4	89.9 ± 1.8	-	-	-	$58.9 \pm \text{N.R.}$
CAC (Ours)	✓	✓	99.1 ± 0.5	94.1 ± 0.7	80.1 ± 3.0	87.7 ± 1.2	87.0 ± 0.0	76.0 ± 1.5

Table 1: Open set AUROC for state-of-the-art methods and our proposed approach. Best and second best performance are bolded and italicised respectively.

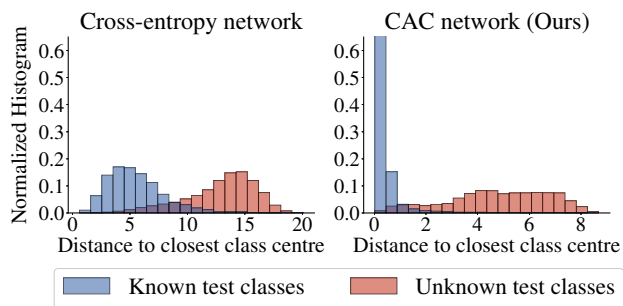


Figure 2: Training with CAC (right) causes known test data to cluster more tightly to the class centres than training with cross-entropy. This allows for distance to better separate known and unknown test data. This plot shows the distributions for MNIST, see Table 2 for other datasets.

Comparison with other distance-based approaches:

Compared to other state-of-the-art methods that use distance in the logit space during testing [2, 28, 29], we achieve the best open set performance on all six of the benchmark datasets. Our performance increase is most substantial on TinyImageNet and CIFAR10, where there is an increase of 17.1% and 10.6%.

Our proposed approach is the first method that trains with a distance-based loss when using distance in the logit space during testing. To analyse the impact of distance-based training, we examine the distributions of known class and unknown class distances to a class centre for a network trained with cross-entropy loss (as used by [2, 28, 29]) and a network trained with our proposed CAC loss. As shown in Figure 2, the CAC-trained network has a known distribution that clusters more tightly to the class centres (behaviour it was trained for), and as a result, there is a lower overlap with the unknown distribution. By reducing the overlap with the unknown distribution, the open set classifier can more accurately identify and reject unknown inputs, thus improving open set performance.

Dataset	Cross-Entropy	CAC (Ours)
MNIST	0.414	0.324
SVHN	0.700	0.573
CIFAR10	0.946	0.868
CIFAR+10	0.899	0.766
CIFAR+50	0.889	0.751
TinyImageNet	0.984	0.913

Table 2: Compared to cross-entropy loss, training with CAC gives a lower Bhattacharyya coefficient between distributions of known and unknown class distances to the closest class centre. This represents less overlap between distributions, enabling better distance-based open set recognition.

In Table 2, we quantitatively show that training with CAC decreases the overlap between the known and unknown class distance distributions in comparison to cross-entropy loss. The table shows the Bhattacharyya coefficient [3], an established measure of the overlap between two distributions. For each of the datasets, CAC loss results in a lower Bhattacharyya coefficient, on average by 14.3%.

Comparison to non-distance-based approaches: Compared to non-distance-based open set classifiers, we achieve state-of-the-art performance on TinyImageNet, MNIST and SVHN, and come second to GFROR [17] on CIFAR10 and CIFAR+10/+50.

While CAC achieves a 15.2% performance increase on TinyImageNet with 20 known classes, it performs less well to [17] on CIFAR+10/+50 variations with only 4 known classes. When presented with only 4 known classes, CAC has less data to learn a rich feature representation that ensures known and unknown class inputs do not project to the same region in the logit space. In contrast, [17] specifically uses reconstruction and self-supervision techniques during training to improve the feature representation.




Figure 3: Our CAC classifier maintains the classification accuracy of a standard classifier trained with cross-entropy loss.

Unknown Data	Method	CCR at FPR of		
		1%	5%	10%
SVHN	Softmax [9]	46.0	N.R.	64.7
	Objectosphere [7] [†]	54.5	N.R.	70.1
	CAC (Ours)	68.3	78.5	82.6
CIFAR100 Subset	Softmax [9]	23.4	N.R.	51.4
	Objectosphere [7] [†]	43.3	N.R.	66.5
	CAC (Ours)	32.5	63.8	73.8

Table 3: At False Positive Rates (FPR) 1%, 5% and 10%, CAC achieves state-of-the-art Correct Classification Rates (CCR) for the unknown datasets (with CIFAR10 as the known dataset). [†] [7] uses a separate subset of CIFAR100 as known unknowns during training.

Maintaining classification accuracy: In Figure 3, we show that training with CAC loss maintains the closed set classification accuracy of a standard network. The standard network uses the same architecture but is trained with cross-entropy loss and uses the softmax score for classification. This result demonstrates that our method improves open set performance without compromising classification accuracy.

Open set performance at low false positive operating points: For many safety-critical domains (e.g. robotics or medical applications), the performance at specific low false positive rates is practically relevant. Although the main benchmark used for open set classification (in Table 1) does not provide such an evaluation, we can compare CAC’s performance at different false positive rates against [7] and [9]. Table 3 shows the results after training a ResNet-18 on CIFAR10, and using a CIFAR100 subset and SVHN as the unknown datasets. We note that [7] uses a ‘known unknown’ subset of CIFAR100 during training, which potentially contributes to their better performance at the very low 1% FPR on the CIFAR100 experiment. We refer the reader to the Supplementary for details on the experimental setup.

Dataset	Center [26]	Tuplet [23]	\mathcal{L}_A only (Ours)	CAC (Ours)
MNIST	0.988	0.957	0.979	0.987
SVHN	0.941	0.833	0.888	0.942
CIFAR10	0.786	0.739	0.751	0.803
CIFAR+10	0.854	0.844	0.804	0.863
CIFAR+50	0.863	0.837	0.816	0.872
TinyImageNet	0.765	0.717	0.749	0.772

Table 4: CAC provides better open set AUROC performance than the compared distance losses on nearly all the benchmark datasets.

5.2. Ablation Studies

5.2.1 Comparison with Existing Distance Losses

We proposed CAC loss specifically for the task of *training* a distance-based open set classifier. However, other distance losses have been proposed for other computer vision tasks, e.g. metric learning [23] and facial recognition [26]. In this experiment, we compare the open set performance achieved when training with Center loss [26], Tuplet loss [23], the Anchor loss component \mathcal{L}_A of CAC, and our proposed CAC loss. We train the same network architecture with each loss function and use our anchored class centres. Table 4 summarises the open set AUROC results for each of the losses.

CAC outperforms all other distance losses [23, 26] on SVHN, CIFAR10, CIFAR+10, CIFAR+50 and TinyImageNet, with Center loss [26] achieving second best performance. Center loss uses cross-entropy loss on the logits to *implicitly* encourage inputs to maximise distance to other class centres. In contrast, CAC *explicitly* forces this behaviour by applying Tuplet loss directly to the output distance vector. Interestingly, when used alone, our Anchor loss term and Tuplet loss cannot achieve the same performance as when they are combined to create CAC loss. This validates that both loss terms are important for distance-based open set classification, as together they simultaneously ensure minimised distance to the correct class centre as well as maximised distance to all other class centres.

5.2.2 Anchored versus Learnt Class Centres

In this section we investigate the benefits of using *anchored* class centres in the context of open set classification. While our work is the first to anchor class centres during the training process, previous distance losses such as Center loss [26] encourage clustering around class centres that are simultaneously *learnt* during training.

We compare the open set performance when training with learnt and anchored class centres, and repeat this experiment with CAC loss and Center loss [26]. To learn class centres, we use the approach described in [26]. Learn-

Dataset	Center [26]		CAC (Ours)	
	Learnt	Anchored	Learnt	Anchored
MNIST	0.985	0.988	0.987	0.987
SVHN	0.937	0.941	0.946	0.942
CIFAR10	0.763	0.786	0.791	0.803
CIFAR+10	0.831	0.854	0.856	0.863
CIFAR+50	0.848	0.863	0.865	0.872
TinyImNet	0.738	0.765	0.764	0.772

Table 5: Anchored class centres yields better open set AUROC than learnt class centres, particularly on the object-based CIFAR10 and TinyImageNet datasets.




Figure 4: Anchored class centres perform better with increasing numbers of classes than learnt class centres. Results are averaged over 5 trials of random class splits.

ing centres with CAC required the addition of cross-entropy loss for stability (see supplementary material for details).

In Table 5, we show that anchored class centres yield better open set performance than learnt class centres, for both Center loss [26] and our proposed CAC loss. The performance difference between anchored and learnt centres is greatest for the object-based datasets (CIFAR10 variants and TinyImageNet), with an average 2.2% improvement for Center loss [26] and 0.85% for CAC. Learning class centres during training relies on a stable learning signal from the images in each batch. However, CIFAR10 and TinyImageNet can exhibit considerable visual variations within each class, thus providing a potentially noisy learning signal for the class centre positions. By anchoring our class centres in the logit space, we eliminate this difficulty and allow for high performance on object-based datasets.

Learning class centres is even more difficult for tasks with large numbers of classes, as each batch will provide less data per class. To investigate this effect, we train with learnt and anchored class centres for Center loss [26] and CAC loss on *increasing* numbers of known TinyImageNet classes, while keeping the openness of the open set task fixed at 18.35%. As we can see in Figure 4, open set performance of a network trained with learnt class centres degrades at a faster rate than a network trained with anchored class centres for both Center loss [26] and CAC loss.

In summary, we show that training with anchored class centres yields better performance on object-based datasets

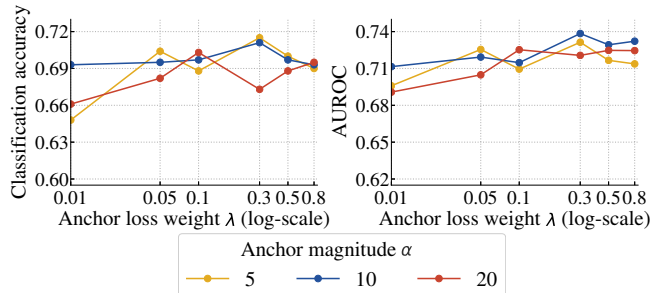


Figure 5: CAC Performance is not sensitive to hyperparameters anchor loss weight λ and anchor magnitude α across a range of values (results from 1 trial on TinyImageNet).

and scales better to larger numbers of training classes. We found this to be consistent for both tested loss functions. In addition, we observed training with anchored centres requires approximately half the epochs to converge, speeding up the training process (see supplementary material).

5.2.3 Analysis of Hyperparameters of CAC loss

Our proposed CAC loss has two hyperparameters, the Anchor loss term weight α and the anchored centre magnitude λ , and their sensitivity is shown in Figure 5. With an Anchor loss weight $0.05 \leq \lambda \leq 0.8$ and an anchor magnitude $5 \leq \alpha \leq 20$, both the classification accuracy and open set AUROC vary by less than 4%.

6. Conclusions

The deployment of deep neural networks under open set conditions remains an important and difficult challenge for computer vision. Reliability and robustness in the presence of unknown class inputs is crucial for many safety-critical applications such as driverless cars or robotics.

We introduced and demonstrated the benefits of anchored class centres and the novel Class Anchor Clustering loss for open set recognition. Future work could focus on learning more meaningful and diverse feature representations, potentially by allowing for more complex arrangements of the anchored class centres. A more semantically meaningful feature representation may allow for better distinction of unknown classes and increased interpretability of open set errors, both of which have the potential to improve distance-based open set recognition even further. We additionally look forward to engaging with the community to develop new evaluation protocols and datasets beyond the simple ones commonly used in open set recognition, such as MNIST, SVHN, or even CIFAR10. Many practical applications rely on open set robustness, and we believe benchmark datasets should better reflect the complexity and richness of those real world applications.

References

- [1] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [2] Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1563–1572, 2016.
- [3] Anil Bhattacharyya. On a measure of divergence between two statistical populations defined by their probability distributions. *Bull. Calcutta Math. Soc.*, 35:99–109, 1943.
- [4] Hermann Blum, Paul-Edouard Sarlin, Juan Nieto, Roland Siegwart, and Cesar Cadena. The fishyscapes benchmark: measuring blind spots in semantic segmentation. *arXiv preprint arXiv:1904.03215*, 2019.
- [5] Terrance E Boult, Steve Cruz, Akshay Raj Dhamija, Manuel Günther, James Henrydoss, and Walter J Scheirer. Learning and the unknown: Surveying steps toward open world recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9801–9807, 2019.
- [6] Taylor Denouden, Rick Salay, Krzysztof Czarnecki, Vahdat Abdelzad, Buu Phan, and Sachin Vernekar. Improving reconstruction autoencoder out-of-distribution detection with mahalnobis distance. *arXiv preprint arXiv:1812.02765*, 2018.
- [7] Akshay Raj Dhamija, Manuel Günther, and Terrance E Boult. Reducing network agnostophobia. In *Advances in Neural Information Processing Systems*, 2018.
- [8] Bolin Gao and Lacta Pavel. On the properties of the softmax function with application in game theory and reinforcement learning. *arXiv preprint arXiv:1704.00805*, 2017.
- [9] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *International Conference on Learning Representations*, 2017.
- [10] Alex Krizhevsky. Learning multiple layers of features from tiny images. *Technical Report*, 2009.
- [11] Yann LeCun, Corinna Cortes, and Christopher CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2010.
- [12] Benjamin J Meyer and Tom Drummond. The importance of metric learning for robotic vision: Open set recognition and active learning. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2924–2931, 2019.
- [13] Lawrence Neal, Matthew Olson, Xiaoli Fern, Weng-Keen Wong, and Fuxin Li. Open set learning with counterfactual images. In *Proceedings of the European Conference on Computer Vision*, pages 613–628, 2018.
- [14] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *Neural Information Processing Systems Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [15] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 427–436, 2015.
- [16] 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.
- [17] Pramuditha Perera, Vlad I Morariu, Rajiv Jain, Varun Manjunatha, Curtis Wigington, Vicente Ordonez, and Vishal M Patel. Generative-discriminative feature representations for open-set recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11814–11823, 2020.
- [18] Pramuditha Perera and Vishal M Patel. Deep transfer learning for multiple class novelty detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11544–11552, 2019.
- [19] Qi Qian, Jiasheng Tang, Hao Li, Shenghuo Zhu, and Rong Jin. Large-scale distance metric learning with uncertainty. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8542–8550, 2018.
- [20] Walter J Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E Boult. Toward open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(7):1757–1772, 2013.
- [21] Patrick Schlachter, Yiwen Liao, and Bin Yang. Open-set recognition using intra-class splitting. In *Proceedings of the European Signal Processing Conference*. IEEE, 2019.
- [22] Alireza Shafaei, Mark Schmidt, and James J Little. A less biased evaluation of out-of-distribution sample detectors. *arXiv preprint arXiv:1809.04729*, 2018.
- [23] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in Neural Information Processing Systems*, pages 1857–1865, 2016.
- [24] Niko Sünderhauf, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Ucroft, Pieter Abbeel, Wolfram Burgard, Michael Milford, and Peter Corke. The limits and potentials of deep learning for robotics. *The International Journal of Robotics Research*, 37(4-5):405–420, 2018.
- [25] TinyImageNet. Tiny ImageNet Visual Recognition Challenge.
- [26] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision*, pages 499–515. Springer, 2016.
- [27] Chao-Yuan Wu, R Manmatha, Alexander J Smola, and Philipp Krahenbuhl. Sampling matters in deep embedding learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2840–2848, 2017.
- [28] Ryota Yoshihashi, Wen Shao, Rei Kawakami, Shaodi You, Makoto Iida, and Takeshi Naemura. Classification-reconstruction learning for open-set recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4016–4025, 2019.
- [29] Sergey Demyanov Zongyuan Ge and Rahil Garnavi. Generative openmax for multi-class open set classification. In *Proceedings of the British Machine Vision Conference*, pages 42.1–42.12, 2017.