

AutoNovel: Automatically Discovering and Learning Novel Visual Categories

Kai Han, Sylvestre-Alvise Rebuffi, Sébastien Ehrhardt, Andrea Vedaldi, Andrew Zisserman

Abstract—We tackle the problem of discovering novel classes in an image collection given labelled examples of other classes. We present a new approach called *AutoNovel* to address this problem by combining three ideas: (1) we suggest that the common approach of bootstrapping an image representation using the labelled data only introduces an unwanted bias, and that this can be avoided by using self-supervised learning to train the representation from scratch on the union of labelled and unlabelled data; (2) we use ranking statistics to transfer the model’s knowledge of the labelled classes to the problem of clustering the unlabelled images; and, (3) we train the data representation by optimizing a joint objective function on the labelled and unlabelled subsets of the data, improving both the supervised classification of the labelled data, and the clustering of the unlabelled data. Moreover, we propose a method to estimate the number of classes for the case where the number of new categories is not known a priori. We evaluate *AutoNovel* on standard classification benchmarks and substantially outperform current methods for novel category discovery. In addition, we also show that *AutoNovel* can be used for fully unsupervised image clustering, achieving promising results.

Index Terms—novel category discovery, deep transfer clustering, clustering, classification, incremental learning.

1 INTRODUCTION

MODERN machine learning systems can match or surpass human-level performance in tasks such as image classification [1], but at the cost of collecting large quantities of annotated training data. Semi-supervised learning (SSL) [2] can alleviate this issue by mixing labelled with unlabelled data, which is usually much cheaper to obtain. However, these methods still require some annotations for each of the classes that one wishes to learn. We argue that this is not always possible in real applications. For instance, consider the task of recognizing products in supermarkets; thousands of new products are introduced in stores every week, and it would be very expensive to annotate them all. However, new products do not differ drastically from the existing ones, so prior knowledge of older products should help to discover new products automatically as they arise in the data. Unfortunately, machines are still unable to effectively detect and learn new classes without manual annotations.

In this paper, we thus consider the problem of discovering new visual classes automatically, assuming that a certain number of classes are already known by the model [3], [4], [5] (see fig. 1). This knowledge comes in the form of a *labelled dataset* of images for a certain set of classes. Given that this data is labelled, off-the-shelf supervised learning techniques can be used to train a very effective classifier for the known classes, particularly if Convolutional Neural Networks (CNNs) are employed. However, this does not mean that the learned features are useful as a representation of the *new classes*. Furthermore, even if the representation transfers well, one still has the problem of identifying the

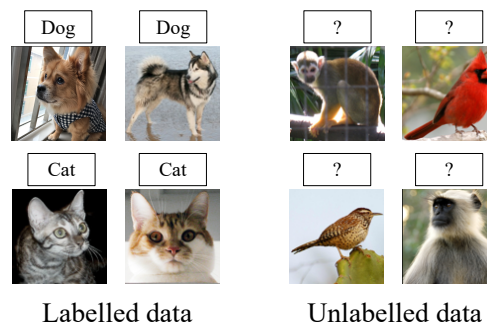


Fig. 1: Novel category discovery. Given labelled images from a few known categories (e.g., dog and cat), our objective is to automatically partition unlabelled images from new categories (e.g., monkey and bird) into proper clusters.

new classes in an unlabelled dataset, which is a clustering problem.

We tackle these problems by introducing a novel approach called *AutoNovel* that combines three key ideas (section 3 and fig. 2). The first idea is to pre-train the image representation (a CNN) using all available images, both labelled and unlabelled, using a self-supervised learning objective. Crucially, this objective *does not* leverage the known labels, resulting in features that are much less biased towards the labelled classes. Labels are used only after pre-training to learn a classifier specific to the labelled data as well as to fine-tune the last layers of the CNN.

The second idea is a new approach to transfer the information contained in the labelled images to the problem of clustering the unlabelled ones. Information is transferred by sharing the same representation between labelled and unlabelled images, in order to be able to reuse discriminative features learned on the labelled set. In more detail, pairs of

• The authors are with Visual Geometry Group, Department of Engineering Science, University of Oxford, Oxford, OX1 3PJ, United Kingdom. E-mail: {khan, srebuffi, hyenal, vedaldi, az}@robots.ox.ac.uk Corresponding author: Kai Han.

unlabelled images are compared via their representation vectors. The comparison is done using robust ranking statistics, by testing if two images share the same subset of k maximally activated representation components. This test is used to decide if two unlabelled images belong to the same (new) class or not, generating a set of noisy pairwise pseudo-labels. The pseudo-labels are then used to learn a similarity function for the unlabelled images.

The third idea is, after bootstrapping the representation, to optimise the model by minimizing a joint objective function, containing terms for both the labelled and unlabelled subsets. To do this, we use respectively the given labels and the generated pseudo-labels, thus avoiding the forgetting issue [6] that may arise with a sequential approach. A further boost is obtained by incorporating incremental learning of the discovered classes in the classification task, which allows information to flow between the labelled and unlabelled images.

However, this approach still requires knowing the number of new categories in the unlabelled data, which is not a realistic assumption in many applications. We propose a method to estimate the number of classes in the unlabelled data which also transfers knowledge from the set of known classes. The idea is to use part of the known classes as a probe set, adding them to the unlabelled set pretending that part of them are unlabelled, and then running the clustering algorithm described above on the extended unlabelled dataset. This allows us to cross-validate the number of new classes, according to the clustering accuracy on the probe set as well as a cluster quality index on the unlabelled set, resulting in a reliable estimate of the true number of unlabelled classes.

In addition, AutoNovel can be used for unsupervised image clustering by simply removing the requirement of labelled data, resulting in a simplified version of our method, achieving the state-of-the-art results on image clustering.

We evaluate AutoNovel on several public benchmarks, outperforming by a large margin all existing techniques that can be applied to novel category discovery, demonstrating the effectiveness of our approach. We also evaluate our category number estimation method, showing reliable estimation of the number of categories in the unlabelled data. We also demonstrate the effectiveness of our method by using it for unsupervised image clustering, achieving state-of-the-art clustering results.

We have presented preliminary results of this work in [7], and this paper extends them in several aspects. First, we include a solution to handle the case of an unknown number of categories in unlabelled data, which we initially introduced in [5]. Second, we study different alternatives to ranking statistics for generating pairwise pseudo labels and compare their effectiveness. Third, we expand the experiments and study transferring representations from ImageNet-pretrained models to new domains. Fourth, we test the effectiveness of different self-supervised learning methods when used as a component of our method. Fifth, we show that our method can also be used for unsupervised clustering, achieving the state-of-the-art results on public benchmarks.

The code reproducing our experiments can be found at http://www.robots.ox.ac.uk/~vgg/research/auto_novel.

2 RELATED WORK

Our work draws inspiration from semi-supervised learning, transfer learning, clustering, and zero-shot learning. We review below the most relevant contributions.

In semi-supervised learning (SSL) [8], a partially labelled training dataset is given and the objective is to learn a model that can propagate the labels from the labelled data to unlabelled data. Most SSL methods focus on the classification task where, usually, both labelled and unlabelled points belong to the same set of classes. On the contrary, our goal is to handle the case where the unlabelled data classes differ from the labelled data. [2] summarizes the state-of-the-art SSL methods. Among them, the consistency-based methods appear to be the most effective. [9] proposes a ladder network which is trained on both labelled and unlabelled data using a reconstruction loss. [10] simplifies this ladder network by enforcing prediction consistency between a data point and its augmented counterpart. As an alternative to data augmentation, they also consider a regularization method based on the exponential moving average (EMA) of the predictions. This idea is further improved by [11]: instead of using the EMA of predictions, they propose to maintain the EMA of model parameters. The consistency is then measured between the predictions of the current model (student) and the predictions of the EMA model (teacher). More recently (and closer to our work) practitioners have also combined SSL with self-supervision [12], [13] to leverage datasets with very few annotations. Finally, FixMatch [14] also uses pseudo-labels extracted from the most confident images reaching state-of-the-art results in SSL benchmarks. However, they use pseudo-labels as soft targets for the cross-entropy loss, while we use a binary score to evaluate similarity of sample pairs within a mini-batch.

Transfer learning [15], [16], [17] is an effective way to reduce the amount of data annotations required by pre-training the model on a different dataset. In image classification, for example, it is customary to start from a model pre-trained on the ImageNet [1] dataset. In most transfer learning settings, however, both the source data and the target data are fully annotated. In contrast, our goal is to transfer information from a labelled dataset to an unlabelled one.

Many classic (e.g., [18], [19], [20], [21]) and deep learning (e.g., [3], [4], [22], [23], [24], [25], [26]) clustering methods have been proposed to automatically partition an unlabelled data collection into different classes. However, this task is usually ill-posed as there are multiple, equally valid criteria to partition most datasets. We address this challenge by learning the appropriate criterion by using a labelled dataset, narrowing down what constitutes a proper class. We call this setting “transfer clustering”.

To the best of our knowledge, the work most related to ours are [3], [4], [5]. In [5], the authors also consider discovering new classes as a transfer clustering problem. They first learn a data embedding by using metric learning on the labelled data, and then fine-tune the embedding and learn the cluster assignments on the unlabelled data. In [3], [4], the authors introduce KCL and MCL clustering methods. In both, a similarity prediction network (SPN), also used in [27], is first trained on a labelled dataset. Afterwards, the pre-trained SPN is used to provide binary pseudo labels

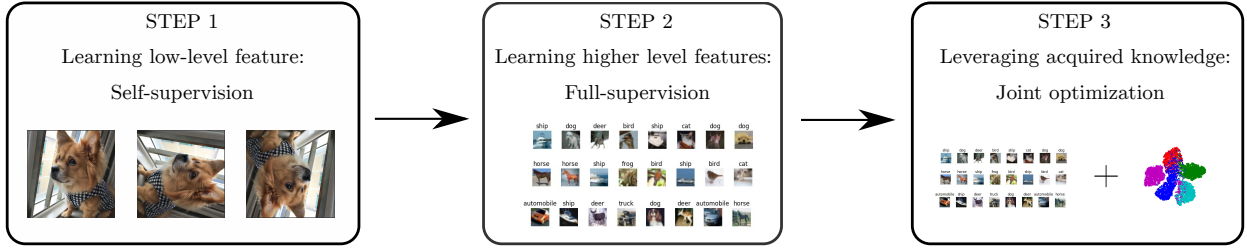


Fig. 2: Overview of the AutoNovel learning pipeline for novel category discovery. The first step is to learn an unbiased image representation via self-supervision using both labelled and unlabelled data. This learns well the early layers of the representation. The second step is to fine-tune only the last few layers of the representation using supervision on the labelled subset of the data. The final step is to use the fine-tuned representation, via ranking statistics, to induce clusters in the unlabelled data, while maintaining a good representation on the labelled set.

for training the main model on an unlabelled dataset. The overall pipelines of the two methods are similar, but the losses differ: KCL uses a contrastive loss based on Kullback-Leibler divergence, which is equivalent to the BCE used in this paper (eq. (3)), and MCL uses the Meta Classification Likelihood loss. Zero-shot learning (ZSL) [28], [29] can also be used to recognize new classes. However, differently from our work, ZSL also requires additional side information (e.g., class attributes) in addition to the raw images.

Finally, other works [30], [31] discuss the application of ranking statistics to measuring the similarity of vectors; however, to the best of our knowledge, we are the first to apply ranking statistics to the task of novel category discovery using deep neural networks.

3 METHOD

Given an *unlabelled* dataset $D^u = \{x_i^u; i = 1, \dots, M\}$ of images $x_i^u \in \mathbb{R}^{3 \times H \times W}$, our goal is to automatically cluster them into a number of classes C^u . We also assume to have a second *labelled* image dataset $D^l = \{(x_i^l, y_i^l); i = 1, \dots, N\}$ where $y_i^l \in \{1, \dots, C^l\}$ is the class label for image x_i^l , where the set of C^l labelled classes is disjoint from the set of C^u unlabelled ones. While the statistics of D^l and D^u differ, we hypothesize that a general notion of what constitutes a “good class” can be extracted from D^l and that the latter can be used to better cluster D^u .

We approach the problem by learning an image representation $\Phi: x \mapsto \Phi(x) \in \mathbb{R}^d$ in the form of a CNN. The goal of the representation is to recognize the known classes and to discover the new ones. In order to learn this representation, we propose AutoNovel, a method that combines three ideas detailed in the next three sections.

3.1 Self-supervised learning

Given that we have a certain number of labelled images D^l at our disposal, the obvious idea is to use these labels to bootstrap the representation Φ by minimizing a standard supervised objective such as the cross-entropy loss. However, experiments show that this causes the representation to overly-specialize for the classes in D^l , providing a poor representation of the new classes in D^u .

Thus we resist the temptation of using the labels right away and instead use a self-supervised learning method to bootstrap the representation Φ . Self-supervised learning has

been shown [32], [33] to produce robust low-level features, especially for the first few layers of typical CNNs. It has the benefit that no data annotations are needed, and thus it can be applied to both labelled and unlabelled images during training. In this way, we achieve the key benefit of ensuring that the representation is initialized without being biased towards the labelled data.

In detail, we first pre-train our model Φ with self-supervision on the union of D^l and D^u (ignoring all labels). We use the RotNet [33] approach¹ as our default choice due to its simplicity and efficacy, but any self-supervised method could be used instead. In our experiments, we also experimented with other self-supervised learning methods such as SimCLR [34] and MoCo [35]. Interestingly, we found these alternatives perform less effectively than RotNet in our setting, though they have shown better performance for fully supervised downstream tasks such as recognition and detection. We then extend the pre-trained network Φ with a classification head $\eta^l: \mathbb{R}^d \rightarrow \mathbb{R}^{C^l}$ implemented as a single linear layer followed by a softmax layer. The function $\eta^l \circ \Phi$ is fine-tuned on the labelled dataset D^l in order to learn a classifier for the C^l known classes, this time using the labels y_i and optimizing the standard cross-entropy (CE) loss:

$$L_{CE} = -\frac{1}{N} \sum_{i=1}^N \log[\eta^l \circ z_i^l]_{y_i} \quad (1)$$

where $z_i^l = \Phi(x_i^l) \in \mathbb{R}^d$ is the representation of image x_i^l . Only η^l and the last macro-block of Φ (see section 5 for details) are updated in order to avoid overfitting the representation to the labelled data.

3.2 Transfer learning via ranking statistics

Once the representation Φ and the classifier η^l have been trained, we are ready to look for the new classes in D^u . Since the classes in D^u are unknown, we represent them by defining a relation among pairs of unlabelled images (x_i^u, x_j^u) . The idea is that similar images should belong to the same

1. We present to the network Φ randomly-rotated versions $R(x)$ of each image and task it with predicting R . The problem is formulated as a 4-way classification of the rotation angle, with angle in $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$. The model $\eta \circ \Phi(R(x))$ is terminated by a single linear layer η with 4 outputs each scoring an hypothesis. The parameters of η and Φ are optimized by minimizing the cross-entropy loss on the rotation prediction.

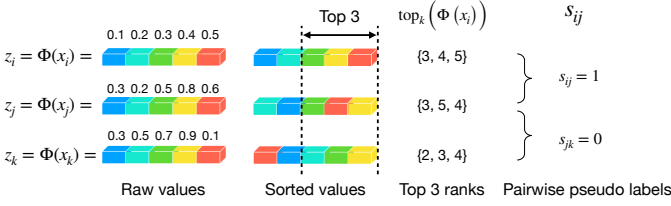


Fig. 3: Ranking statistics. In this example, we consider top-3 ranks. As the top-3 ranks of z_i and z_j are the same, $s_{ij} = 1$. While the top-3 ranks of z_j and z_k are the different, so $s_{jk} = 0$.

(new) class, which we denote by the symbol $s_{ij} = 1$, while dissimilar ones should not, which we denote by $s_{ij} = 0$. A similar idea has been applied in the literature (e.g., [3], [4], [36], [37]). The problem is then to obtain the labels s_{ij} .

Our assumption is that the new classes will have some degree of visual similarity with the known ones. Hence, the learned representation should be applicable to old and new classes equally well. As a consequence, we expect the descriptors $z_i^u = \Phi(x_i^u)$ and $z_j^u = \Phi(x_j^u)$ of two images x_i^u , x_j^u from the new classes to be close if they are from the same (new) class, and to be distinct otherwise. The way this is done is explained in the next section.

3.2.1 Ranking statistics

Rather than comparing vectors z_i^u and z_j^u directly (e.g., by a scalar product), we use a more robust ranking statistics. Specifically, we rank the values in vector z_i^u by magnitude. Then, if the rankings obtained for two unlabelled images x_i^u and x_j^u are the same, they are very likely to belong to the same (new) class, so we set $s_{ij} = 1$. Otherwise, we set $s_{ij} = 0$. In practice, it is too strict to require the two rankings to be identical if the dimension of z_i^u is high (otherwise we may end up with $s_{ij} = 0$ for all pairs $(i, j), i \neq j$). Therefore, we relax this requirement by only testing if the *sets* of the top- k ranked dimensions are the same (we use $k = 5$ in our experiments), i.e.:

$$s_{ij} = \mathbb{1} \{ \text{top}_k(\Phi(x_i^u)) = \text{top}_k(\Phi(x_j^u)) \}, \quad (2)$$

where $\text{top}_k : \mathbb{R}^d \rightarrow \mathcal{P}(\{1, \dots, d\})$ associates to a vector z the subset of indices $\{1, \dots, d\}$ of its top- k elements. Figure 3 shows an example of using ranking statistics to obtain pairwise pseudo labels.

Once the values s_{ij} have been obtained, we use them as pseudo-labels to train a comparison function for the unlabelled data. In order to do this, we apply a new head $\eta^u : \mathbb{R}^d \rightarrow \mathbb{R}^{C^u}$ to the image representation $z_i^u = \Phi(x_i^u)$ to extract a new descriptor vector $\eta^u(z_i^u)$ optimized for the unlabelled data. As in section 3.1, the head is composed of a linear layer followed by a softmax. Then, the inner product $\eta^u(z_i^u)^\top \eta^u(z_j^u)$ is used as a score for whether images x_i^u and x_j^u belong to the same class or not. Note that $\eta^u(z_i^u)$ is a normalized vector due to the softmax layer in η^u . This descriptor is trained by optimizing the *binary cross-entropy* (BCE) loss:

$$L_{\text{BCE}} = - \frac{1}{M^2} \sum_{i=1}^M \sum_{j=1}^M [s_{ij} \log \eta^u(z_i^u)^\top \eta^u(z_j^u) + (1 - s_{ij}) \log(1 - \eta^u(z_i^u)^\top \eta^u(z_j^u))]. \quad (3)$$

Furthermore, we structure η^u in a particular manner: We set its output dimension to be equal to the number of new classes C^u , which is a common practice for clustering in the literature (e.g., [3], [4], [36], [37]). In this manner, we can use the index of the maximum element of each vector $\hat{y}_i^u = \text{argmax}_y [\eta^u \circ \Phi(x_i^u)]_y$ as prediction \hat{y}_i^u for the class of image x_i^u (as opposed to assigning labels via a clustering method such as k -means).

Alternatives to ranking statistics. While we adopt ranking statistics to obtain the pairwise pseudo labels for unlabelled data, there exist many other options such as k -means, cosine similarity, and nearest neighbor. In the experiment, we evaluate applying k -means on the unlabelled data in each mini-batch and use the resulting cluster assignments to generate pairwise pseudo labels. We can also compute cosine similarity between vectors and generate binary pseudo-labels based on a predefined threshold τ . Another natural way to generate binary pseudo-labels is using the mutual nearest neighbor criteria, for which we follow [38] and define

$$s_{ij} = \begin{cases} 1 & \text{if } j = \kappa_i^1 \text{ or } \kappa_j^1 = i \text{ or } \kappa_i^1 = \kappa_j^1 \\ 0 & \text{otherwise} \end{cases}, \quad (4)$$

where κ_i^1 denotes the nearest neighbor of image i in the mini-batch. As will be shown in the experiments, these alternatives are less effective than using ranking statistics.

Discussion. There are several reasons why we believe the ranking statistics do well. First, the statistics focus on the top- k most active feature components for each image. Intuitively, the magnitude of these components reflects the degree to which they are discriminative for the object in the image. Thus, our ranking statistics only considers the most salient feature components when comparing images, while ignoring noisy components with small values.

Second, other similarity measures like cosine similarity, which use the whole feature vectors for comparison in a high-dimensional vector space, can potentially suffer from the problem of distance concentration [39]. The distance concentration is the counter-intuitive phenomenon that, as the data dimensionality increases, all pairwise distances between points may converge to the same value, which is not desired in our case.

Third, we further relax the comparison by not requiring the order of the top- k ranks to be identical. Instead, we only check the *sets* of the top- k ranks. This further makes the pairwise comparisons robust to slight discrepancies among the most discriminative feature components.

3.3 Joint training on labelled and unlabelled data

We now have two losses that involve the representation Φ : the CE loss L_{CE} for the labelled data D^l and the pairwise BCE loss L_{BCE} for the unlabelled data D^u . They both share the same image embedding Φ . This embedding can be trained sequentially, first on the labelled data, and then on the unlabelled data using the pseudo-labels obtained above. However, in this way the model will very likely forget the knowledge learned from the labelled data, which is known as *catastrophic forgetting* in incremental learning [40], [41], [42], [43].

Instead, we jointly fine-tune our model using both losses at the same time. Note that most of the model Φ is frozen;

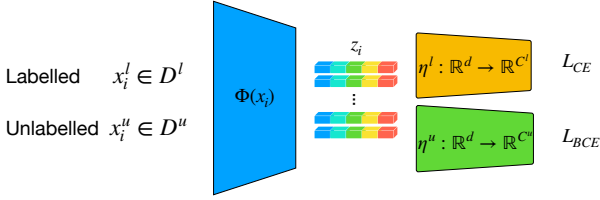


Fig. 4: Joint learning on labelled and unlabelled data.

we only fine-tune the last macro-block of Φ together with the two heads η^u and η^l . Figure 4 demonstrates the overall architecture for joint learning. Importantly, as we fine-tune the model, the labels s_{ij} are changing at every epoch as the embedding η^l is updated. This in turn affects the ranking statistics used to determine the labels s_{ij} as explained in section 3.2. This leads to a “moving target” phenomenon that can introduce some instability in learning the model. This potential issue is addressed in the next section.

3.4 Enforcing predictions to be consistent

In addition to the CE and BCE losses, we also introduce a consistency regularization term, which is used for both labelled and unlabelled data. In semi-supervised learning [2], [10], [11], the idea of consistency is that the class predictions on an image x and on a randomly-transformed counterpart $t(x)$ (for example an image rotation) should be the same. In our case, as will be shown in the experiments, consistency is important to obtain good performance. One reason is that, as noted above, the pairwise pseudo-labels for the unlabelled data are subject to change on the fly during training. Indeed, for an image x_i^u and a randomly-transformed counterpart $t(x_i^u)$, if we do not enforce consistency, we can have $\text{top}_k(\Phi(x_i^u)) \neq \text{top}_k(\Phi(t(x_i^u)))$. According to eq. (2) defining s_{ij} , this could result in different s_{ij} for (x_i^u, x_j^u) depending on the data augmentation applied to the images. This variability of the ranking labels for a given pair could then confuse the training of the embedding.

Following the common practice in semi-supervised learning, we use the *Mean Squared Error* (MSE) as the consistency cost. This is given by:

$$L_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N (\eta^l(z_i^l) - \eta^l(\hat{z}_i^l))^2 + \frac{1}{M} \sum_{i=1}^M (\eta^u(z_i^u) - \eta^u(\hat{z}_i^u))^2, \quad (5)$$

where \hat{z} is the representation of $t(x)$.

The overall loss of our model can then be written as

$$L = L_{\text{CE}} + L_{\text{BCE}} + \omega(r)L_{\text{MSE}}, \quad (6)$$

where the coefficient $\omega(r)$ is a ramp-up function. This is widely used in semi-supervised learning [10], [11]. Following [10], [11], we use the sigmoid-shaped function $\omega(r) = \lambda e^{-5(1-\frac{r}{T})^2}$, where r is current time step and T is the ramp-up length and $\lambda \in \mathbb{R}_+$. As opposed to contrastive learning [34] this loss is only minimising the distance between positive pairs of samples within a minibatch.

3.5 Incremental learning scheme

We also explore a setting analogous to incremental learning. In this approach, after tuning on the labelled set (end of section 3.1), we extend the head η^l to C^u new classes, so that $\eta^l: \mathbb{R}^d \rightarrow \mathbb{R}^{C^l+C^u}$. The head parameters for the new classes are initialized randomly. The model is then trained using the same loss eq. (6), but the cross-entropy part of the loss is evaluated on both labelled and unlabelled data D^l and D^u . Since the cross-entropy requires labels, for the unlabelled data we use the *pseudo-labels* \hat{y}_i^u , which are generated on-the-fly from the head η^u at each forward pass.

The advantage is that this approach *increments* η^l to discriminate both old and new classes, which is often desirable in real applications. It also creates a feedback loop that causes the features z_i^u to be refined, which in turn generates better pseudo-labels \hat{y}_i^u for D^u from the head η^u . In this manner, further improvements can be obtained by this cycle of positive interactions between the two heads during training.

3.6 Unsupervised clustering

Rather than working on the task of novel category discovery, AutoNovel can also be used for standard (unsupervised) clustering by simply removing the use of labelled data (thus dropping step two of the method), obtaining a two-step approach. In the first step, we pre-train our model with self-supervised learning as before. In the second step, we finetune the last macro block and the linear layer for clustering, using the ranking statistics to provide pseudo pairwise labels. This way, our method can simultaneously learn the feature embedding for clustering as well as the cluster assignments. The training loss in eq. (6) becomes:

$$L = L_{\text{BCE}} + \omega(r)L_{\text{MSE}}. \quad (7)$$

To our knowledge, Deep Adaptive Clustering (DAC) [36] is the deep clustering method most related to this approach, in the sense that DAC also reduces clustering as a binary classification problem. However, our approach differs from DAC in several aspects. First, our method uses ranking statistics to generate pairwise pseudo labels instead of the cosine similarity (though any pairwise labeling method could also be used in our approach). Second, our method optimizes the standard binary cross-entropy loss with a consistency constraint, while DAC optimizes a Bhattacharyya distance with an ad-hoc sample number penalty term. Third, our method incorporates self-supervised learning for pretraining lower level features and only needs to train the last macro-block and the linear layers of the model, which is less likely to suffer from overfitting. By comparison, DAC updates all parameters of the model. As will be seen in the experiment, our method significantly outperforms DAC as well as the recent state-of-the-art method IIC [44] in several benchmarks. Note that the main objective of this work is novel category discovery rather than clustering, and here we only show that our method can be easily adopted for the problem of clustering achieving superior performance than exiting alternatives.

Method summary 1 Estimating the number of classes

- 1: **Preparation:**
- 2: Split the probe set D_r^l into D_{ra}^l and D_{rv}^l .
- 3: Extract features of D_r^l and D^u using Φ .
- 4: **Main loop:**
- 5: **for** $0 \leq C_i^u \leq C_{\max}^u$ **do**
- 6: Run k -means on $D_r^l \cup D^u$ assuming $C_r^{lu} = C_r^l + C_i^u$ classes in semi-supervised mode (i.e. forcing data in D_{ra}^l to map to the ground-truth class labels).
- 7: Compute ACC for D_{rv}^l and CVI for D^u .
- 8: **end for**
- 9: **Obtain optimal:**
- 10: Let C_a^{u*} be the value of C_i^u that maximise ACC for D_{rv}^l and C_v^{u*} be the value that maximise CVI for D^u and let $\hat{C}^u = (C_a^{u*} + C_v^{u*})/2$. Run semi-supervised k -means on $D_r^l \cup D^u$ again assuming $C_r^l + \hat{C}^u$ classes.
- 11: **Remove outliers:**
- 12: Look at the resulting clusters in D^u and drop any that has a mass less than τ of the largest cluster. Output the number of remaining clusters.

4 ESTIMATING THE NUMBER OF CLASSES

So far, we have assumed that the number of classes C^u in the unlabelled data is known, but this is usually not the case in real applications. Here we propose a new approach to estimate the number of classes in the unlabelled data by making use of labelled probe classes. The probe classes are combined with the unlabelled data and the resulting set is clustered using k -means multiple times, varying the number of classes. The resulting clusters are then examined by computing two quality indices, one of which checks how well the probe classes, for which ground truth data is available, have been identified. The number of categories is then estimated to be the one that maximizes these quality indices.

In more details, we first split the C^l known classes into a probe subset D_r^l of C_r^l classes and a training subset $D_r^l \setminus D_{ra}^l$ containing the remaining $C^l - C_r^l$ classes. These $C^l - C_r^l$ classes are used for supervised feature representation learning, while the C_r^l probe classes are combined with the unlabelled data for class number estimation. We then further split the C_r^l probe classes into a subset D_{ra}^l of C_{ra}^l classes and a subset D_{rv}^l of C_{rv}^l classes (e.g., $C_{ra}^l : C_{rv}^l = 4 : 1$), which we call anchor probe set and validation probe set respectively (see fig. 5). We then run a constrained (semi-supervised) k -means on $D_r^l \cup D^u$ to estimate the number of classes in D^u . Namely, during k -means, we force images in the anchor probe set D_{ra}^l to map to clusters following their ground-truth labels, while images in the validation probe set D_{rv}^l are considered as additional “unlabelled” data. We launch this constrained k -means multiple times by sweeping the number of total categories C_r^{lu} in $D_r^l \cup D^u$, and measure the constrained clustering quality on $D_r^l \cup D^u$. We consider two quality indices, given below, for each value of C_r^{lu} . The first measures the cluster quality in the labelled validation probe set D_{rv}^l , whereas the second measures the quality in the unlabelled data D^u . Each index is used to determine an optimal number of classes and the results are averaged. Finally, k -means is run one last time with this value as the

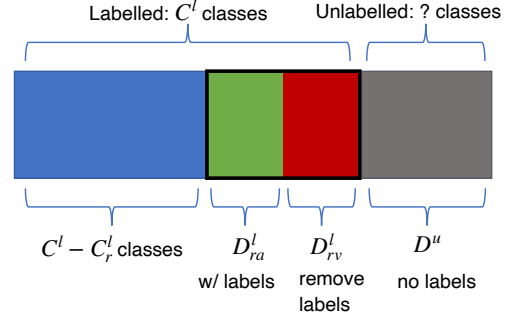


Fig. 5: Data split for category number estimation.

number of classes and any outlier clusters in D^u , defined as containing less than τ (e.g., $\tau = 1\%$) the mass of the largest clusters, are dropped. The details are given in method summary 1.

Cluster quality indices. We measure our clustering for class number estimation with two indices. The first index is the *average clustering accuracy* (ACC), which is applicable to the C_{rv}^l labelled classes in the validation probe set D_{rv}^l and is given by

$$\max_{g \in \text{Sym}(C_{rv}^l)} \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{\bar{y}_i = g(y_i)\}, \quad (8)$$

where \bar{y}_i and y_i denote the ground-truth label and clustering assignment for each data point $x_i \in D_{rv}^l$ and $\text{Sym}(C_{rv}^l)$ is the group of permutations of C_{rv}^l elements (this discounts the fact that the cluster indices may not be in the same order as the ground-truth labels). Permutations are optimized using the Hungarian algorithm [45].

The other index is a *cluster validity index* (CVI) [46] which, by capturing notions such as intra-cluster cohesion vs inter-cluster separation, is applicable to the unlabelled data D^u . There are several CVI metrics, such as Silhouette [47], Dunn [48], Davies–Bouldin [49], and Calinski–Harabasz [50]; while no metric is uniformly the best, the Silhouette index generally works well [46], [51], and we found it to be a good choice for our case too. This index is given by

$$\sum_{x \in \mathcal{D}^u} \frac{b(x) - a(x)}{\max\{a(x), b(x)\}}, \quad (9)$$

where x is a data sample, $a(x)$ is the average distance between x and all other data samples within the same cluster, and $b(x)$ is the smallest average distance of x to all points in any other cluster (of which x is not a member).

5 EXPERIMENTS

5.1 Data and implementation details

We evaluate AutoNovel on a variety of standard benchmark datasets: CIFAR10 [52], CIFAR100 [52], SVHN [53], OmniGlott [54], and ImageNet [1]. Following [5], we split these to have 5/20/5/654/30 classes respectively in the unlabelled set. The splits are summarized in table 1. In addition, for OmniGlott and ImageNet we use 20 and 3 different splits respectively, as in [5], and report average clustering accuracy (as defined in eq. (8)) on the unlabelled data. While we follow standard practice to split the datasets we note here

that most of the time, the number of unlabelled classes is under a hundred. This is a potential limitation of clustering which still proves to be very difficult for classification over thousands of categories [55].

TABLE 1: Data splits in the experiments.

	labelled classes	unlabelled classes
CIFAR10	5	5
CIFAR100	80	20
SVHN	5	5
OmniGlot	964	659
ImageNet	882	118

We use the ResNet-18 [56] architecture, except for OmniGlot for which we use a VGG-like network [57] with six layers to make our setting directly comparable to prior work. We use SGD with momentum [58] as the optimizer for all but the OmniGlot dataset, for which we use Adam [59]. For all experiments we use a batch size of 128 and $k = 5$ which we found work consistently well across datasets.

In the first self-supervised training step, unless otherwise mentioned, we train our model with the pretext task of rotation predictions (i.e., a four-class classification: 0° , 90° , 180° , and 270°) for 200 epochs and a step-wise decaying learning rate starting from 0.1 and divided by 5 at epochs 60, 120, and 160.

In the second step of our framework (i.e., supervised training using labelled data), we fine-tune our model on the labelled set for 100 epochs and a step-wise decaying learning rate starting from 0.1 and halved every 10 epochs. From this step onward we fix the first three convolutional blocks of the model, and fine-tune the last convolutional block together with the linear classifier.

Finally, in the last joint training step, we fine-tune our model for 200/100/90 epochs for {CIFAR10, CIFAR100, SVHN}/OmniGlot/ImageNet, which is randomly sampled from the merged set of both labelled and unlabelled data. The initial learning rate is set to 0.1 for all datasets, and is decayed with a factor of 10 at the 170th/{30th, 60th} epoch for {CIFAR10, CIFAR100, SVHN}/ImageNet. The learning rate of 0.01 is kept fixed for OmniGlot. For the consistency regularization term, we use the ramp-up function as described in section 3.4 with $\lambda = \{5.0, 50.0, 50.0, 100.0, 10.0\}$, and $T = \{50, 150, 80, 1, 50\}$ for CIFAR10, CIFAR100, SVHN, OmniGlot, and ImageNet respectively.

In the incremental learning setting, all previous hyper parameters remain the same for our method. We only add a ramp-up on the cross entropy loss on unlabelled data. The ramp-up length is the same as the one used for eq. (4) and we use for all experiments a coefficient of 0.05. For all other methods we train the classifier for 150 epochs with SGD with momentum and learning rate of 0.1 divided by 10 at epoch 50.

For hyper-parameter tuning, we create a probe validation set from the labelled data by dropping the labels of a few classes. We then tune the hyper-parameters based on the ACC on this probe validation set. We construct the probe validation set to have the same number of classes as the actual unlabelled set. For example, for CIFAR100, we split the 80 labelled classes into a 60-class labelled subset and a 20-class probe validation set. We then tune the hyper-parameters

TABLE 2: Ablation study of AutoNovel. “MSE” means MSE consistency constraint; “CE” means cross entropy loss for training on labeled data; “BCE” means binary cross entropy loss for training on unlabeled data; “S.S.” means self-supervision; “I.L.” means incremental learning. The evaluation metric is the ACC.

	CIFAR-10	CIFAR-100	SVHN
Ours w/o MSE	82.6±12.0%	61.8±3.6%	61.3±1.9%
Ours w/o CE	84.7±4.4%	58.4±2.7%	59.7±6.6%
Ours w/o BCE	26.2±2.0%	6.6±0.7%	24.5±0.5%
Ours w/o S.S.	89.4±1.4%	67.4±2.0%	72.9±5.0%
Ours full	90.4±0.5%	73.2±2.1%	95.0±0.2%
Ours w/ I.L.	91.7±0.9%	75.2±4.2%	95.2±0.3%

based on the novel category discovery performance on the probe validation set. For CIFAR10 and SVHN, due to the small number of labelled classes, we only take 2 classes from the labelled data to construct the probe validation set.

We implement our method using PyTorch 1.1.0 and run experiments on NVIDIA Tesla M40 GPUs. Following [5], our results are averaged over 10 runs for all datasets, except ImageNet, for which the results are averaged over the three 30-class subsets. In general, we found the results are stable. Our code is publicly available at http://www.robots.ox.ac.uk/~vgg/research/auto_novel.

5.2 Ablation study

We validate the effectiveness of the components of AutoNovel by ablating them and measuring the resulting ACC on the unlabelled data. Note that, since the evaluation is restricted to the unlabelled data, we are solving a clustering problem. The same unlabelled data points are used for both training and testing, except that data augmentation (i.e. image transformations) is not applied when computing the cluster assignments. As can be seen in table 2, all components have a significant effect as removing any of them causes the performance to drop substantially. Among them, the BCE loss is by far the most important one, since removing it results in a dramatic drop of 40–60% absolute ACC points. For example, the full method has ACC 90.4% on CIFAR10, while removing BCE causes the ACC to drop to 26.2%. This shows that that our rank-based embedding comparison can indeed generate reliable pairwise pseudo labels for the BCE loss. Without consistency, cross entropy, or self-supervision, the performance drops by a more modest but still significant 7.8%, 5.7% and 1.0% absolute ACC points, respectively, for CIFAR10. It means that the consistency term plays a role as important as the cross-entropy term by preventing the “moving target” phenomenon described in section 3.4. Finally, by incorporating the discovered classes in the classification task, we get a further boost of 1.3%, 2.0% and 0.2% points on CIFAR10, CIFAR100 and SVHN respectively.

We also evaluate the evolution of performances of our method with respect to k for ranking statistics. The results on SVHN/CIFAR10/CIFAR100 are shown in fig. 6. We found that $k = \{5, 7\}$ give the best results overall. We also found that for all values of k except 1 results are in general stable.

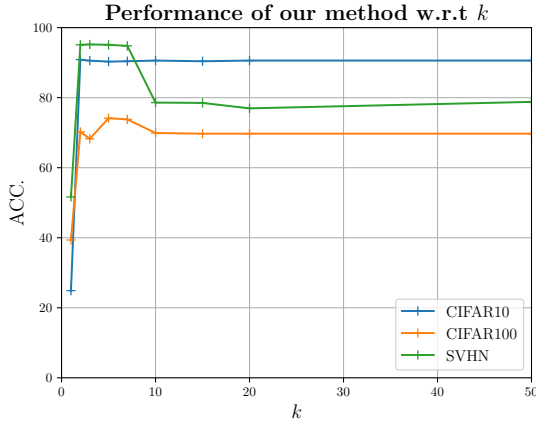


Fig. 6: Performance evolution w.r.t. k for ranking statistics. We report results for $k = \{1, 2, 3, 5, 7, 10, 15, 20, 50\}$.

TABLE 3: Novel category discovery results on CIFAR10, CIFAR100, and SVHN. ACC on the unlabelled set. ‘‘S.S.’’ means self-supervision; ‘‘I.L.’’ means incremental learning.

No		CIFAR10	CIFAR100	SVHN
(1)	k -means [19]	65.5±0.0 %	56.6±1.6%	42.6±0.0
(2)	KCL [3]	66.5±3.9%	14.3±1.3%	21.4±0.6
(3)	MCL [4]	64.2±0.1%	21.3±3.4%	38.6±10.8
(4)	DTC [5]	87.5±0.3%	56.7±1.2%	60.9±1.6
(5)	k -means [19] w/ S.S.	72.5±0.0%	56.3±1.7%	46.7±0.0%
(6)	KCL [3] w/ S.S.	72.3±0.2%	42.1±1.8%	65.6±4.9%
(7)	MCL [4] w/ S.S.	70.9±0.1%	21.5±2.3%	53.1±0.3%
(8)	DTC [5] w/ S.S.	88.7±0.3%	67.3±1.2%	75.7±0.4%
(9)	Ours	90.4±0.5%	73.2±2.1%	95.0±0.2%
(10)	Ours w/ I.L.	91.7±0.9%	75.2±4.2%	95.2±0.2%

5.3 Novel category discovery

We compare AutoNovel to baselines and state-of-the-art methods for new class discovery, starting from CIFAR10, CIFAR100, and SVHN in table 3. The first baseline (row 5 in table 3) amounts to applying k -means [19] to the features extracted by the fine-tuned model (the second step in section 3.1), for which we use the k -means++ [60] initialization. The second baseline (row 1 in table 3) is similar, but uses as feature extractor a model trained from scratch using only the labelled images, which corresponds to a standard transfer learning setting. By comparing rows 1, 5 and 9 in table 3, we can see that our method substantially outperforms k -means. Next, we compare with the KCL [3], MCL [4] and DTC [5] methods. By comparing rows 2–4 to 9, we see that our method outperforms these by a large margin. We also try to improve KCL, MCL and DTC by using the same self-supervised initialization we adopt (section 3.1), which indeed results in an improvement (rows 2–4 vs 6–8). However, their overall performance still lags behind ours by a large margin. For example, our method of section 3.4 achieves 95.0% ACC on SVHN, while ‘‘KCL w/ S.S.’’, ‘‘MCL w/ S.S.’’ and ‘‘DTC w/ S.S.’’ achieve only 65.6%, 53.1% and 75.7% ACC, respectively. Similar trends hold for CIFAR10 and CIFAR100. Finally, the incremental learning scheme of section 3.5 results in further improvements, as can be seen by comparing rows 9 and 10 of table 3.

In fig. 7, we show the evolution of the learned representation on the unlabelled data from CIFAR10 using t-

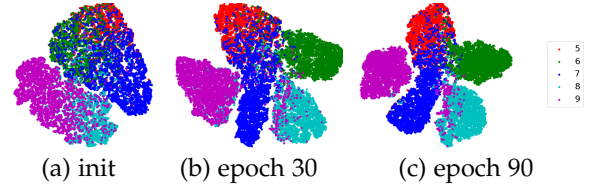


Fig. 7: Evolution of the t-SNE during the training of CIFAR-10. Performed on unlabelled data (i.e., instances of dog, frog, horse, ship, truck). Colors of data points denote their ground-truth labels.

TABLE 4: Novel category discovery results on OmniGlot and ImageNet. ACC on the unlabelled set.

No		OmniGlot	ImageNet
(1)	k -means [19]	77.2%	71.9%
(2)	KCL [3]	82.4%	73.8%
(3)	MCL [4]	83.3%	74.4%
(4)	DTC [5]	89.0%	78.3%
(5)	Ours	89.1%	82.5%

SNE [61]. As can be seen, while the clusters overlap in the beginning, they become more and more separated as the training progresses, showing that our model can effectively discover novel visual categories without labels and learn meaningful embeddings for them.

We further compare AutoNovel to others on two more challenging datasets, OmniGlot and ImageNet, in table 4. For OmniGlot, results are averaged over the 20 alphabets in the *evaluation* set; for ImageNet, results are averaged over the three 30-class unlabelled sets used in [3], [4]. Since we have a relatively larger number of labelled classes in these two datasets, we follow [5] and use metric learning on the labelled classes to pre-train the feature extractor, instead of the self-supervised learning. We empirically found that self-supervision does not provide obvious gains for these two datasets. This is reasonable since the data in the labelled sets of these two datasets are rather diverse and abundant, so metric learning can provide good feature initialization as there is less class-specific bias due to the large number of pre-training classes. However, by comparing rows 1 and 5 in table 4, it is clear that metric learning alone is not sufficient for the task of novel category discovery. Our method substantially outperforms the k -means results obtained using the features from metric learning — by 11.9% and 10.6% on OmniGlot and ImageNet respectively. Our method also substantially outperforms the current state-of-the-art, achieving 89.1% and 82.5% ACC on OmniGlot and ImageNet respectively, compared with 89.0% and 78.3% of [5], thus setting the new state-of-the-art. By comparing table 3 and table 4, we observe that KCL and MCL perform better on the more challenging ImageNet than the smaller datasets CIAR10, CIFAR100 and SVHN. This can be explained by the fact that the pairwise pseudo labels are provided by a similarity prediction network (SPN) which is pretrained on the labelled data. As there are much less labelled data in CIFAR10, CIFAR100 and SVHN than ImageNet, the learned SPN is less reliable, thus resulting in relatively poor performance for novel category discovery on unlabelled data from new classes.

TABLE 5: Incremental Learning with the novel categories. “old” refers to the ACC on the labelled classes while “new” refers to the unlabelled classes in the *testing set*. “all” indicates the whole testing set. It should be noted that the predictions are not restricted to their respective subset. “S.S.” means self-supervision; “I.L.” means incremental learning.

Classes	CIFAR10			CIFAR100			SVHN		
	old	new	all	old	new	all	old	new	all
KCL w/ S.S.	79.4±0.6%	60.1±0.6%	69.8±0.1%	23.4±0.3%	29.4±0.3%	24.6±0.2%	90.3±0.3%	65.0±0.5%	81.0±0.1%
MCL w/ S.S.	81.4±0.4%	64.8±0.4%	73.1±0.1%	18.2±0.3%	18.0±0.1%	18.2±0.2%	94.0±0.2%	48.6±0.3%	77.2±0.1%
DTC w/ S.S.	58.7±0.6%	78.6±0.2%	68.7±0.3%	47.6±0.2%	49.1±0.2%	47.9±0.2%	90.5±0.3%	72.8±0.2%	84.0±0.1%
Ours w/ I.L.	90.6±0.2%	88.8±0.2%	89.7±0.1%	71.2±0.1%	56.8±0.3%	68.3±0.1%	96.3±0.1%	96.1±0.0%	96.2±0.1%

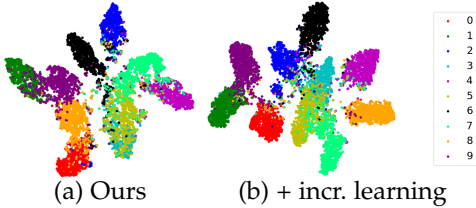


Fig. 8: t-SNE on CIFAR10: impact of incremental Learning. Colors of data points denote their ground-truth labels (“old” classes 0-4; “new” classes 5-9). We observe a bigger overlap in (a) between the “old” class 3 and the “new” class 5 when not incorporating Incremental Learning.

5.4 Incremental learning scheme

Here, we further evaluate our incremental scheme for novel category discovery as described in section 3.5. Methods for novel category discovery such as [3], [4], [5] focus on obtaining the highest clustering accuracy for the new unlabelled classes, but may forget the existing labelled classes in the process. In practice, forgetting is not desirable as the model should be able to recognize both old and new classes. Thus, we argue that the classification accuracy on the labelled classes should be assessed as well, as for any incremental learning setting. Note however that our setup differs substantially from standard incremental learning [40], [41], [42], [43] where every class is labelled and the focus is on using limited memory. In our case, we can store and access the original data without memory constraints, but the new classes are unlabelled, which is often encountered in practical applications.

By construction (section 3.5), our method learns the new classes on top of the old ones incrementally, out of the box. In order to compare AutoNovel to methods such as KCL, MCL and DTC that do not have this property, we proceed as follows. First, the method runs as usual to cluster the unlabelled portion of the data, thus obtaining pseudo-labels for it, and learning a feature extractor as a byproduct. Then, the feature extractor is used to compute features for both the labelled and unlabelled training data, and a linear classifier is trained using labels and pseudo-labels, jointly on all the classes, old and new.

We report in table 5 the performance of the resulting joint classifier networks on the *testing set* of each dataset (this is now entirely disjoint from the training set). Our method has similar performances on the old and new classes for CIFAR10 and SVHN, as might be expected as the split between old and new classes is balanced. In comparison, the feature extractor

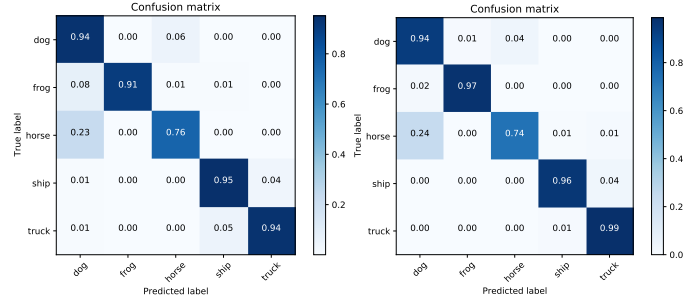


Fig. 9: Confusion matrix on unlabelled classes of CIFAR10. Left: our method; Right: our method w/ I.L.

learned by KCL and MCL works much better for the old classes (e.g., the accuracy discrepancy between old and new classes is 25.3% for KCL on SVHN). Conversely, DTC learns features that work better for the new classes, as shown by the poor performance for the old classes on CIFAR10. Thus, KCL, MCL and DTC learn representations that are biased to either the old or new classes, resulting overall in suboptimal performance. In contrast, our method works well on both old and new classes; furthermore, it drastically outperforms existing methods on both. In fig. 8, we show the t-SNE projection of the learned feature representation on both old and new classes. It can be seen, with incremental learning, the embedding becomes more discriminative between old and new classes. Similarly, in fig. 9 we compare the confusion matrices w/ and w/o the incremental learning scheme. It can be seen that, with the incremental learning scheme, the clusters for new classes turn out to be more accurate. We notice that the errors are mainly due to the confusion between dog and horse. By looking into the images, we found that images of dogs and horses are confused because of having similar colors or poses.

5.5 Finding the number of novel categories

We now experiment under the more challenging (and realistic) scenario where the number of categories in the unlabelled data is unknown. KCL and MCL assume the number of categories to be a large value (i.e., 100) instead of estimating the number of categories explicitly. By contrast, we choose to estimate the number of categories as described in method summary 1 (with $C_{\max}^u = 100$ for all our experiments), before running the transfer clustering algorithm, and only then apply our ranking based method to learn the representation and find the cluster assignment. Results for novel category number estimation are reported in table 6 and table 7 on

OmniGlot and ImageNet respectively. The average errors are 4.6 on OmniGlot and 2.33 on ImageNet, which validates the effectiveness of our approach. In table 8, we show the clustering results on OmniGlot and ImageNet, by substituting these estimates to our ranking based method for novel category discovery, and also compare with other methods. The results of traditional methods are those reported in [4] using raw images for OmniGlot and pretrained features for ImageNet. AutoNovel outperforms the previous state-of-the-art MCL by 5.2% and 9.0% ACC on OmniGlot and ImageNet respectively. We also experiment on KCL, MCL and DTC by using our estimated number of clusters. With this augmentation, both KCL and MCL improve significantly, indicating that our category number estimation method can also be beneficial for other methods. DTC slightly outperforms our ranking based method by 1.6% on OmniGlot with our estimated category number, while our method outperforms DTC by 2.9% on ImageNet. In addition, we also validate the sensitivity of different methods to the choice of cluster number on the 30-class ImageNet_A. We vary the cluster number from 20 to 100. The results are shown in fig. 10. It can be seen that the sensitivity to the cluster number is similar for all methods. All methods achieve the best performance when the cluster number equals the ground truth, while the performance drops when the cluster number is off the ground truth. Our method consistently outperforms all others for cluster numbers 25 to 40 (note that our estimated cluster number is 34). For the extreme case with the cluster number of 100, MCL performs the best.

TABLE 6: Category number estimation on OmniGlot.

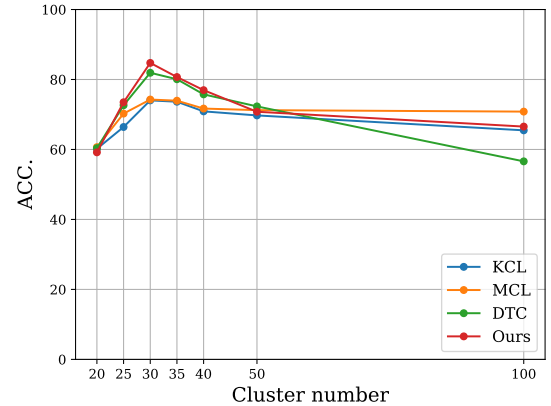
Alphabet	GT	SKMS [62]	KCL [3]	MCL [4]	Ours
Angelic	20	16	26	22	23
Atemayar Q.	26	17	34	26	25
Atlantean	26	21	41	25	34
Aurek_Besh	26	14	28	22	34
Avesta	26	8	32	23	31
Ge_ez	26	18	32	25	31
Glagolitic	45	18	45	36	46
Gurmukhi	45	12	43	31	34
Kannada	41	19	44	30	40
Keble	26	16	28	23	25
Malayalam	47	12	47	35	42
Manipuri	40	17	41	33	39
Mongolian	30	28	36	29	33
Old Church S.	45	23	45	38	51
Oriya	46	22	49	32	33
Sylheti	28	11	50	30	22
Syriac_Serto	23	19	38	24	26
Tengwar	25	12	41	26	28
Tibetan	42	15	42	34	43
ULOG	26	15	40	27	33
<i>Avg_{error}</i>	-	16.3	6.35	5.10	4.60

TABLE 7: Category number estimation results.

Data	GT	Ours	Error
ImageNet _A	30	34	4
ImageNet _B	30	31	1
ImageNet _C	30	32	2
<i>Avg_{error}</i>	-	-	2.33

TABLE 8: Novel category discovery with an unknown class number C^u .

	OmniGlot	ImageNet
Method	ACC	ACC
<i>k</i> -means [19]	18.9%	34.5%
LPNMF [63]	16.3%	21.8%
LSC [64]	18.0%	33.5%
KCL [3]	78.1%	65.2%
MCL [4]	80.2%	71.5%
KCL [3] w/our C^u	80.3%	71.4%
MCL [4] w/our C^u	80.5%	72.9%
DTC [5] w/our C^u	87.0%	77.6%
Ours	85.4%	80.5%

Fig. 10: Performance of different methods with different cluster number on ImageNet_A. The ground truth is 30. We vary the cluster number from 20 to 100.

5.6 Transferring from ImageNet pretrained model

Rather than pretraining the model with self-supervised learning, one may also think to transfer representation learned from other datasets. The most common way of transfer learning with modern deep convolutional neural networks is to use ImageNet pretrained models. Here, we explore the potential of leveraging the ImageNet pretrained model to transfer features for novel category discovery. In particular, we take the ImageNet pretrained model as our feature extractor, and finetune the last macro-block and the linear heads of the model using our ranking based method. We experiment on CIFAR10, CIFAR100, and SVHN. The results are shown in table 9. As can be seen, with the ImageNet pretrained representation, the performance of our method on CIFAR10 and CIFAR100 are further improved w.r.t to the results in table 3. The incremental learning scheme successfully boosts the performance by 0.7% and 3.8% on CIFAR10 and CIFAR100 respectively. For the performance on SVHN, we notice a significant drop between table 3 and table 9. This is likely due to the small correlation between ImageNet and SVHN, as also noted in other works that try to transfer features from ImageNet to other datasets [2], [5].

TABLE 9: Transferring from ImageNet to CIFAR10/CIFAR100/SVHN.

	CIFAR10	CIFAR100	SVHN
<i>k</i> -means [19]	92.4%	78.8%	23.4%
Ours	95.4%	87.1%	40.2%
Ours w/I.L.	96.1%	90.9%	38.8%

5.7 Alternatives to ranking statistics

The ranking statistics is the key to transfer knowledge from old classes to new classes in our model. As discussed in section 3.2.1, other methods like *k*-means, cosine similarity, and nearest neighbor can potentially be used as alternatives to ranking statistics to generate pairwise pseudo labels in AutoNovel. We experiment with such alternatives and the results are shown in table 10. We experiment on two cases for *k*-means, one by running *k*-means on the mini-batch (denoted as *k*-means (batch)) and the other by running *k*-means (denoted as *k*-means (all)) on the whole unlabelled set. As it can be seen, ranking statistics, nearest neighbor and cosine similarity work significantly better than *k*-means on CIFAR10 and SVHN, while ranking statistics and cosine similarity work notably better than nearest neighbor on CIFAR100 and SVHN. Note that the performance of cosine similarity depends on the choice of a proper threshold τ . Here, we report the results using the best thresholds on each dataset (0.85/0.8/0.9 for CIFAR10/CIFAR100/SVHN). The effect of different thresholds is shown in fig. 11. It can be seen that, with a carefully chosen threshold, cosine similarity can also be a good measure to generate pairwise pseudo labels in our method, though the results turn to be relatively sensitive to τ . When $\tau < 0.6$, the cosine similarity fails to provide reliable pairwise pseudo labels. Meanwhile, as τ lies in the continues space while *k* in our ranking based method lies in the discrete integer space, it is easier to set a proper *k* than τ . Overall, while it can be seen that ranking statistics and cosine similarity exhibit a similar behaviour when grid-searching with a relatively low sensitivity to the best value, we still find that ranking statistics is an interesting alternative to cosine similarity and is relatively unexplored in the context of deep learning. Throughout all of our experiments we demonstrate that ranking statistics performs consistently well and could open the way for more applications. Therefore, unless stated otherwise, ranking statistics is our default choice for all experiments.

TABLE 10: Different methods for pairwise pseudo labels.

	CIFAR10	CIFAR100	SVHN
<i>k</i> -means (batch)	42.9%	74.3%	45.3%
<i>k</i> -means (all)	62.2%	55.5%	61.5%
cosine	90.1%	73.3%	95.0%
nearest neighbor	90.2%	69.7%	78.2%
ranking statistics	90.4%	73.2%	95.0%
soft ranking statistics (<i>k</i> = 5)	62.2%	65.2%	72.5%
soft ranking statistics (<i>k</i> = 15)	89.7%	71.1%	95.2%

Through ranking statistics, instead of generating hard (binary) pseudo targets, we can also encode soft rank similarities. To do so, we calculate the shared elements in the top-*k* rank between two images. Let *c* be the number

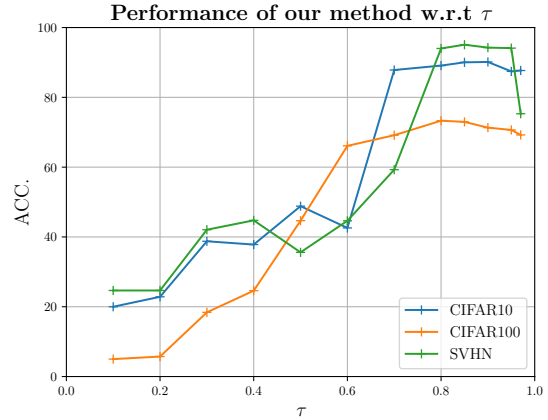


Fig. 11: Performance evolution w.r.t. the threshold τ for cosine similarity. We report results for $\tau = \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.85, 0.9, 0.95, 0.97\}$.

of shared elements. The soft similarity is then defined as $\frac{c}{k}$, which can be used to replace the s_{ij} in eq. (3). The results are shown in table 10. We find that *k* = 5 is not optimal for the soft rank similarity, as the results largely lag behind the hard counterpart. We adopt the validation method introduced in section 5.1 to get *k* = 15 as a better choice for the soft rank similarity. The results are in general on par with the hard (binary) rank similarity.

5.8 Other self-supervised learning methods

We adopt the RotNet [33] for the first stage of AutoNovel. However, any other self-supervised representation methods can be applied. Here, we further experiment with the latest self-supervised representation learning methods including SimCLR [34], MoCo [35] and MoCo v2 [65], which are the state-of-the-art representation learning methods for the tasks of object recognition and detection. In this experiment, we replace RotNet by the latest self-supervised learning methods in our pipeline, and the other two steps remain the same as before. In table 11, we first directly compare the learned feature representations of different methods by running *k*-means on the output of the global average pooling layer on the unlabelled data. All the three alternatives work better than RotNet when comparing the raw features learned with self-supervised learning on CIFAR10 and CIFAR100, while RotNet performs slightly better on SVHN. This is likely due to nature of the pretext tasks used in these self-supervised learning methods. RotNet uses the rotation prediction task, which is less relevant to the down stream task of partitioning the unlabelled data based on their semantic meaning. Differently, the other three methods are contrastive learning based methods, which encourage the images of the same instance to be close in the feature space while the images of different instances to be further away. Interestingly, we find that RotNet consistently outperforms the other three methods for AutoNovel. This reveals that better feature initialization does not necessarily mean better representation fine-tuned on downstream tasks like novel category discovery. Overall, by taking any of these self-supervised learning methods to pre-train our model, the performance can be significantly boosted on novel category discovery by our method.

TABLE 11: Different self-supervised learning methods.

		CIFAR10	CIFAR100	SVHN
<i>k</i> -means [19]	SimCLR [34]	84.7%	41.2%	30.6%
	MoCo [35]	58.7%	34.5%	21.3%
	MoCo v2 [65]	61.8%	39.3%	28.5%
	RotNet [33]	25.5%	10.3%	31.7%
Ours	SimCLR [34]	89.0%	54.6%	67.8%
	MoCo [35]	87.6%	61.1%	74.6%
	MoCo v2 [65]	89.0%	62.5%	76.6%
	RotNet [33]	90.4%	73.2%	95.0%

To investigate why RotNet appears to be more effective in our experiments, we carry out experiments by freezing different layers of the network and finetuning the rest of the layers employing different self-supervised learning approaches. ResNet18 is composed of four macro-blocks, denoted as $layer_{\{1,2,3,4\}}$. In table 12, for each column, we freeze all the parameters before $layer_i$ and finetune $layer_i$ together with the subsequent layers. *head* denotes the case where we only finetune the two linear heads while $layer_0$ denotes the case where we finetune all the parameters. We measure the novel category performance on CIFAR10 for all methods and report the ACC on the unlabelled data. It can be seen that, if we only finetune the linear heads, SimCLR, MoCo and MoCoV2 significantly outperform RotNet, which is consistent with the conclusion in the literature that the contrastive learning based methods can learn more meaningful higher level feature representation. The higher level features for RotNet is focusing on the task of rotation prediction, which is loosely related to the target task of novel category discovery, thus the performance is poor. However, if we finetune more layers, we can see that the performance are similar, while RotNet appears to be more effective for $layer_4$ and $layer_3$. The strong augmentation is essential for the performance of contrastive learning based self-supervision. However, for SVHN, where multiple digits appear in the same image and only the center digit is to be recognized, strong augmentations like cropping is not suitable for training, because random cropping will change the location of the center digits, which is harmful for training. Therefore, the performance of contrastive learning based methods lags behind RotNet on SVHN as in table 11.

TABLE 12: Performance on fine-tuning different layers on CIFAR10. ACC on the unlabelled set.

Method	<i>head</i>	$layer_4$	$layer_3$	$layer_2$	$layer_1$	$layer_0$
RotNet	39.9%	90.4%	90.8%	88.4%	89.3%	88.9%
SimCLR	73.1%	89.0%	89.1%	89.5%	90.4%	88.6%
MoCo	80.8%	87.6%	88.8%	89.3%	90.4%	89.5%
MoCoV2	84.6%	89.0%	89.5%	89.0%	90.3%	89.2%

5.9 Unsupervised image clustering

As discussed in section 3.6, by removing the requirement of labelled data, AutoNovel turns to an unsupervised clustering method that can learn both feature representation and clustering assignment. Here, we compare our method on the clustering problem with the state-of-the-art methods on three popular benchmarks CIFAR10, CIFAR100 and

STL10 [66]. We follow the common practice to use all 10 classes in CIFAR10 and STL10, and the 20 meta classes in CIFAR100 (denoted as CIFAR100-20) in our experiment for fair comparison. The results are presented in table 13. Our method performs on par with the state-of-the-art method IIC [44] on CIFAR10 and STL10, while significantly outperforms IIC on CIFAR100-20 by 9.3%. Compared with IIC, which requires extra Sobel filtering on the input data and large batch sizes (660/1000/700 on CIFAR10/CIFAR100-20/STL10), our method only needs the conventional data augmentation (random cropping and horizontal flipping) and a small batch size of 128 for all three datasets. Therefore, our method is a good alternative to state-of-the-art methods for the task of unsupervised image clustering, though this is not the main objective of this work. Moreover, we report the *k*-means results on the feature representation of the base self-supervised model (i.e., RotNet) on each dataset. Unsurprisingly, the results are not satisfactory, because the high level features of RotNet are learned for the task of rotation prediction, making it less effective in capturing useful semantic information for downstream tasks like clustering. Meanwhile, we also validate the effectiveness of self-supervised pretraining and consistency regularization in eq. (7). We can see that by dropping each of them in our method, the performance drops. Without the self-supervised pretraining, the performance drops significantly. This suggests that self-supervised learning captures discriminative low level features for the task of image clustering. Similar to the task of novel category discovery, when applying our approach to unsupervised clustering, the MSE consistency loss is also effective in preventing the ‘‘moving target’’ phenomenon described in section 3.4 during training. We show the confusion matrix on CIFAR10 by our full method in fig. 12. As can be seen from the diagonal of the matrix, our method can properly cluster objects into proper clusters. We found airplane and bird are confused with ship because of the shared blue background; cat and dog are confused because of similar poses and colors.

TABLE 13: Unsupervised image clustering. ‘‘*k*-means on S.S.’’ refers to the *k*-means results on the representation of the self-supervised model.

	CIFAR10	CIFAR100-20	STL10
<i>k</i> -means [19]	22.9%	13.0%	19.2%
JULE [26]	27.2%	13.7%	27.7%
DEC [22]	30.1%	18.5%	35.9%
DAC [36]	52.2%	23.8%	47.0%
IIC [44]	61.7%	25.7%	59.6%
<i>k</i> -means on S.S.	14.3%	8.8%	15.7%
Ours w/o S.S.	18.8%	13.0%	22.7%
Ours w/o MSE	57.7%	31.6%	48.6%
Ours	61.7%	35.0%	56.4%

6 CONCLUSIONS

In this paper, we have looked at the problem of discovering new classes in an image collection, leveraging labels available for other, known classes. We have proposed AutoNovel to successfully address this task by combining a few new ideas. First, the use of self-supervised learning for bootstrapping

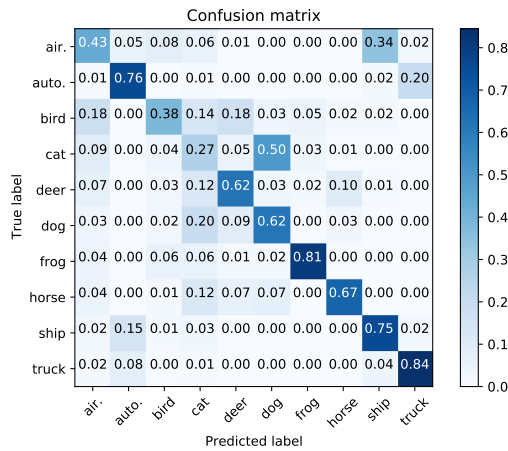


Fig. 12: Confusion matrix of clustering on CIFAR10.

the image representation trades off the representation quality with its generality, and for our problem this leads to a better solution overall. Second, we have shown that ranking statistics are an effective method to compare noisy image descriptors, resulting in robust data clustering. Third, we have shown that jointly optimizing both labelled recognition and unlabelled clustering in an incremental learning setup can reinforce the two tasks while avoiding forgetting. On standard benchmarks, the combination of these ideas results in much better performance than existing methods that solve the same task. For larger datasets with more classes and diverse data (e.g., ImageNet) we note that self-supervision can be bypassed as the pretraining on labelled data already provides a powerful enough representation. In such cases, we still show that the ranking statistics for clustering gives drastic improvement over existing methods. Besides, we have also proposed a method to estimate the number of categories in the unlabelled data, by transferring knowledge from the labelled data to the unlabelled data, allowing our method to handle the more challenging case when number of categories is unknown. Finally, we have shown that AutoNovel can also serve as a simple and effective method for unsupervised image clustering by simply removing the requirement of labelled data, performing on par with the state-of-the-art methods.

One key assumption in our work is that the classes in labelled and unlabelled data follow a similar category definition, in the sense that all classes belong to the same vision dataset. We assume this is the case for categories collected in the same dataset, because we normally follow a consistent procedure to define classes during data curation. Ideally, it would be great to have a precise measure about the relevance between labelled and unlabelled tasks so that we can have a clear sense on whether certain algorithms are applicable or not for novel category discovery. We consider this as a potential future research direction.

ACKNOWLEDGMENTS

This work is supported by the EPSRC Programme Grant Seebibyte EP/M013774/1, Mathworks/DTA DFR02620, and ERC IDIU-638009. We also gratefully acknowledge the support of Nielsen.

REFERENCES

- [1] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *CVPR*, 2009.
- [2] A. Oliver, A. Odena, C. Raffel, E. D. Cubuk, and I. J. Goodfellow, "Realistic evaluation of deep semi-supervised learning algorithms," in *NeurIPS*, 2018.
- [3] Y.-C. Hsu, Z. Lv, and Z. Kira, "Learning to cluster in order to transfer across domains and tasks," in *ICLR*, 2018.
- [4] Y.-C. Hsu, Z. Lv, J. Schlosser, P. Odom, and Z. Kira, "Multi-class classification without multi-class labels," in *ICLR*, 2019.
- [5] K. Han, A. Vedaldi, and A. Zisserman, "Learning to discover novel visual categories via deep transfer clustering," in *ICCV*, 2019.
- [6] M. McCloskey and N. J. Cohen, "Catastrophic interference in connectionist networks: The sequential learning problem," *Psychology of Learning and Motivation*, 1989.
- [7] K. Han, S.-A. Rebuffi, S. Ehrhardt, A. Vedaldi, and A. Zisserman, "Automatically discovering and learning new visual categories with ranking statistics," in *ICLR*, 2020.
- [8] O. Chapelle, B. Scholkopf, and A. Zien, *Semi-Supervised Learning*. MIT Press, 2006.
- [9] A. Rasmus, H. Valpola, M. Honkala, M. Berglund, and T. Raiko, "Semi-supervised learning with ladder networks," in *NeurIPS*, 2015.
- [10] S. Laine and T. Aila, "Temporal ensembling for semi-supervised learning," in *ICLR*, 2017.
- [11] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *NeurIPS*, 2017.
- [12] S.-A. Rebuffi, S. Ehrhardt, K. Han, A. Vedaldi, and A. Zisserman, "Semi-supervised learning with scarce annotations," *arxiv*, 2019.
- [13] X. Zhai, A. Oliver, A. Kolesnikov, and L. Beyer, "S4L: Self-supervised semi-supervised learning," in *ICCV*, 2019.
- [14] K. Sohn, D. Berthelot, C.-L. Li, Z. Zhang, N. Carlini, E. D. Cubuk, A. Kurakin, H. Zhang, and C. Raffel, "Fixmatch: Simplifying semi-supervised learning with consistency and confidence," in *NeurIPS*, 2020.
- [15] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, 2010.
- [16] K. Weiss, T. M. Khoshgoftaar, and D. Wang, "A survey of transfer learning," *Journal of Big Data*, 2016.
- [17] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, "A survey on deep transfer learning," in *International Conference on Artificial Neural Networks*, 2018.
- [18] C. C. Aggarwal and C. K. Reddy, *Data Clustering: Algorithms and Applications*. CRC Press, 2013.
- [19] J. MacQueen, "Some methods for classification and analysis of multivariate observations," in *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 1967.
- [20] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *IEEE TPAMI*, 1979.
- [21] A. Y. Ng, M. I. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *NeurIPS*, 2001.
- [22] J. Xie, R. Girshick, and A. Farhadi, "Unsupervised deep embedding for clustering analysis," in *ICML*, 2016.
- [23] J. Chang, L. Wang, G. Meng, S. Xiang, and C. Pan, "Deep adaptive image clustering," in *ICCV*, 2017.
- [24] K. G. Dizaji, A. Herandi, C. Deng, W. Cai, and H. Huang, "Deep clustering via joint convolutional autoencoder embedding and relative entropy minimization," in *ICCV*, 2017.
- [25] B. Yang, X. Fu, N. D. Sidiropoulos, and M. Hong, "Towards k-means-friendly spaces: Simultaneous deep learning and clustering," in *ICML*, 2017.
- [26] J. Yang, D. Parikh, and D. Batra, "Joint unsupervised learning of deep representations and image clusters," in *CVPR*, 2016.
- [27] Y.-C. Hsu, Z. Lv, and Z. Kira, "Deep image category discovery using a transferred similarity function," *arxiv*, 2016.
- [28] Y. Xian, C. H. Lampert, B. Schiele, and Z. Akata, "Zero-shot learning - a comprehensive evaluation of the good, the bad and the ugly," *IEEE TPAMI*, 2018.
- [29] Y. Fu, T. Xiang, Y.-G. Jiang, X. Xue, L. Sigal, and S. Gong, "Recent advances in zero-shot recognition: Toward data-efficient understanding of visual content," *IEEE Signal Processing Magazine*, 2018.
- [30] T. Dean, M. A. Ruzon, M. Segal, J. Shlens, S. Vijayanarasimhan, and J. Yagnik, "Fast, accurate detection of 100,000 object classes on a single machine," in *CVPR*, 2013.
- [31] J. Yagnik, D. Strelow, D. A. Ross, and R. sung Lin, "The power of comparative reasoning," in *ICCV*, 2011.

- [32] A. Kolesnikov, X. Zhai, and L. Beyer, "Revisiting self-supervised visual representation learning," in *CVPR*, 2019.
- [33] S. Gidaris, P. Singh, and N. Komodakis, "Unsupervised representation learning by predicting image rotations," *ICLR*, 2018.
- [34] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *ICML*, 2020.
- [35] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *CVPR*, 2020.
- [36] J. Chang, G. Meng, L. Wang, S. Xiang, and C. Pan, "Deep self-evolution clustering," in *IEEE TPAMI*, 2018.
- [37] S.-A. Rebuffi, S. Ehrhardt, K. Han, A. Vedaldi, and A. Zisserman, "Lsd-c: Linearly separable deep clusters," *arXiv preprint arXiv:2006.10039*, 2020.
- [38] M. S. Sarfraz, V. Sharma, and R. Stiefelwagen, "Efficient parameter-free clustering using first neighbor relations," in *CVPR*, 2019.
- [39] A. Kabán, "On the distance concentration awareness of certain data reduction techniques," *Pattern Recognition*, 2011.
- [40] S.-A. Rebuffi, A. Kolesnikov, G. Sperl, and C. H. Lampert, "icarl: Incremental classifier and representation learning," in *CVPR*, 2017.
- [41] D. Lopez-Paz and M. Ranzato, "Gradient episodic memory for continual learning," in *NeurIPS*, 2017.
- [42] K. Shmelkov, C. Schmid, and K. Alahari, "Incremental learning of object detectors without catastrophic forgetting," in *ICCV*, 2017.
- [43] R. Aljundi, F. Babiloni, M. Elhoseiny, M. Rohrbach, and T. Tuytelaars, "Memory aware synapses: Learning what (not) to forget," in *ECCV*, 2018.
- [44] X. Ji, J. F. Henriques, and A. Vedaldi, "Invariant information clustering for unsupervised image classification and segmentation," in *ICCV*, 2019.
- [45] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, 1955.
- [46] O. Arbelaitz, I. Gurrutxaga, J. Muguerza, J. M. Pérez, and I. Perona, "An extensive comparative study of cluster validity indices," *Pattern Recognition*, 2012.
- [47] P. J. Rousseeuw, "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics*, 1987.
- [48] J. C. Dunn, "Well-separated clusters and optimal fuzzy partitions," *Journal of Cybernetics*, 1974.
- [49] D. L. Davies and D. W. Bouldin, "A cluster separation measure," *IEEE TPAMI*, 1979.
- [50] T. Caliński and J. Harabasz, "A dendrite method for cluster analysis," *Communications in Statistics-theory and Methods*, 1974.
- [51] J. C. Bezdek and N. R. Pal, "Some new indexes of cluster validity," *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 1998.
- [52] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," *Technical report*, 2009.
- [53] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, "Reading digits in natural images with unsupervised feature learning," in *NeurIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [54] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum, "Human-level concept learning through probabilistic program induction," *Science*, 2015.
- [55] W. Van Gansbeke, S. Vandenhende, S. Georgoulis, M. Proesmans, and L. Van Gool, "Scan: Learning to classify images without labels," in *ECCV*, 2020.
- [56] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.
- [57] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015.
- [58] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *ICML*, 2013.
- [59] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2014.
- [60] D. Arthur and S. Vassilvitskii, "k-means++: the advantages of careful seeding," in *ACM-SIAM symposium on Discrete algorithms*, 2007.
- [61] L. van der Maaten and G. Hinton, "Visualizing data using t-sne," *JMLR*, 2008.
- [62] S. Anand, S. Mittal, O. Tuzel, and P. Meer, "Semi-supervised kernel mean shift clustering," *IEEE TPAMI*, 2014.
- [63] D. Cai, X. He, X. Wang, H. Bao, and J. Han, "Locality preserving nonnegative matrix factorization," in *IJCAI*, 2009.
- [64] X. Chen and D. Cai, "Large scale spectral clustering with landmark-based representation," in *AAAI*, 2011.
- [65] X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," *arXiv preprint arXiv:2003.04297*, 2020.
- [66] A. Coates, H. Lee, and A. Y. Ng, "An analysis of single layer networks in unsupervised feature learning," in *AISTATS*, 2011.



Kai Han received the Ph.D. degree from the Department of Computer Science, The University of Hong Kong, in 2018. He was a Postdoctoral Researcher with the Visual Geometry Group (VGG), Department of Engineering Science, University of Oxford, from 2018 to 2020. He is currently a Lecturer (Assistant Professor) in Computer Vision with the Department of Computer Science at University of Bristol. His research interests include computer vision and deep learning.



Sylvestre-Alvise Rebuffi studied at Ecole Centrale Paris, ENS Cachan and then obtained in 2020 his DPhil from the University of Oxford where he was a member of the Visual Geometry Group (VGG) from 2016 to 2020. He is currently a Research Scientist at Google DeepMind. His research spans diverse machine learning and computer vision topics ranging from continual learning, multi-task learning, interpretability and robustness.



Sébastien Ehrhardt received his DPhil from the University of Oxford in 2021. He is a former member of the Visual Geometry Group (VGG), Department of Engineering Science, University of Oxford, where he worked under Andrea Vedaldi's supervision from 2016 to 2020. He is currently an Applied Scientist at Onfido. His research interests revolves the application of deep learning to vision datasets with limited number of annotated data.



Andrea Vedaldi is currently a professor of computer vision and machine learning with the University of Oxford, where he has been co-leading Visual Geometry Group since 2012. He is also a research scientist with Facebook AI Research, London. He has authored or coauthored more than 130 peer-reviewed publications in the top machine vision and artificial intelligence conferences and journals. His research interests include unsupervised learning of representations and geometry in computer vision. He was the

recipient of the Mark Everingham Prize for selfless contributions to the computer vision community, the Open Source Software Award by the ACM, and the Best Paper Award from the Conference on Computer Vision and Pattern Recognition.

Andrew Zisserman is the Professor of Computer Vision Engineering in the Department of Engineering Science at the University of Oxford.