



HAL
open science

A Practical Approach to Novel Class Discovery in Tabular Data

Colin Troisemaine, Alexandre Reiffers-Masson, Stéphane Gosselin, Vincent
Lemaire, Sandrine Vaton

► **To cite this version:**

Colin Troisemaine, Alexandre Reiffers-Masson, Stéphane Gosselin, Vincent Lemaire, Sandrine Vaton.
A Practical Approach to Novel Class Discovery in Tabular Data. 2023. hal-04283853v1

HAL Id: hal-04283853

<https://hal.science/hal-04283853v1>

Preprint submitted on 14 Nov 2023 (v1), last revised 7 Jun 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Practical Approach to Novel Class Discovery in Tabular Data

Troisemaine Colin^{1,2*}, Reiffers-Masson Alexandre¹,
Gosselin Stéphane², Lemaire Vincent², Vaton Sandrine¹

¹Department of Computer Science, IMT Atlantique, Brest, France.

²Orange Labs, Lannion, France.

*Corresponding author(s). E-mail(s): colin.troisemaine@gmail.com;

Abstract

The problem of Novel Class Discovery (NCD) consists in extracting knowledge from a labeled set of known classes to accurately partition an unlabeled set of novel classes. While NCD has recently received a lot of attention from the community, it is often solved on computer vision problems and under unrealistic conditions. In particular, the number of novel classes is usually assumed to be known in advance, and their labels are sometimes used to tune hyperparameters. Methods that rely on these assumptions are not applicable in real-world scenarios. In this work, we focus on solving NCD in tabular data when no prior knowledge of the novel classes is available. To this end, we propose to tune the hyperparameters of NCD methods by adapting the k -fold cross-validation process and hiding some of the known classes in each fold. Since we have found that methods with too many hyperparameters are likely to overfit these hidden classes, we define a simple deep NCD model. This method is composed of only the essential elements necessary for the NCD problem and performs impressively well under realistic conditions. Furthermore, we find that the latent space of this method can be used to reliably estimate the number of novel classes. Additionally, we adapt two unsupervised clustering algorithms (k -means and Spectral Clustering) to leverage the knowledge of the known classes. Extensive experiments are conducted on 7 tabular datasets and demonstrate the effectiveness of the proposed method and hyperparameter tuning process, and show that the NCD problem can be solved without relying on knowledge from the novel classes.

Keywords: novel class discovery, clustering, tabular data, open world learning, transfer learning

1 Introduction

Recently, remarkable progress has been achieved in supervised tasks, in part with the help of large and fully labeled sets such as ImageNet [1]. These advancements have predominantly focused on closed-world scenarios, where, during training, it is presumed that all classes are known in advance and have some labeled examples. However, in practical applications, obtaining labeled instances for all classes of interest can be a difficult task due to factors such as budget constraints or lack of comprehensive information. Furthermore, for models to be able to transfer learned concepts to new classes, they need to be designed with this in mind from the start, which is rarely the case. Yet this is an important skill that humans can use effortlessly. For example, having learnt to distinguish a few animals, a person will easily be able to recognise and “cluster” new species they have never seen before. The transposition of this human capacity to the field of machine learning could be a model capable of categorizing new products in novel categories.

This observation has led researchers to formulate a new problem called Novel Class Discovery (NCD) [2, 3]. Here, we are given a labeled set of known classes and an unlabeled set of different but related classes that must be discovered. Lately, this task has received a lot of attention from the community, with many new methods such as AutoNovel [4], OpenMix [5] or NCL [6] and theoretical studies [7, 8]. However, most of these works tackle the NCD problem under the unrealistic assumption that the number of novel classes is known in advance, or that the target labels of the novel classes are available for hyperparameter optimization [9]. These assumptions render these methods impractical for real-world NCD scenarios. To address these challenges, we propose a general framework for optimizing the hyperparameters of NCD methods where the ground-truth labels of novel classes are never used, as they are not available in real-world NCD scenarios. Furthermore, we show that the latent spaces obtained by such methods can be used to accurately estimate the number of novel classes.

We also introduce three new NCD methods. Two of them are unsupervised clustering algorithms modified to leverage the additional information available in the NCD setting. The first one improves the centroid initialization step of k -means, resulting in a fast and easy to use algorithm that can still give good results in many scenarios. The second method focuses on optimizing the parameters of the Spectral Clustering (SC) algorithm. This approach has a potentially higher learning capacity as the representation itself (i.e. the spectral embedding) is tuned to easily cluster the novel data. Finally, the last approach is a deep NCD method composed of only the essential components necessary for the NCD problem. Compared to SC, this method is more flexible in the definition of its latent space and effectively integrates the knowledge of the known classes.

While these contributions can be applied to any type of data, our work focuses on tabular data. The NCD community has focused almost exclusively on computer vision problems and, to the best of our knowledge, only one paper [9] has tackled the problem of NCD in the tabular context. However, this work required the meticulous tuning of a large number of hyperparameters to achieve optimal results. Methods designed for tabular data cannot take advantage of powerful techniques commonly employed in computer vision. Examples include convolutions, data augmentation or

Self-Supervised Learning methods such as DINO [10], which have been used with great success in NCD works [11–13], thanks to their strong ability to obtain representative latent spaces without any supervision. On the other hand, tabular data methods have to rely on finely tuned hyperparameters to achieve optimal results. For this reason, we believe that the field of tabular data will benefit the most from our contributions.

By making the following contributions, we demonstrate the feasibility of solving the NCD problem with tabular data and under realistic conditions:

- We develop a hyperparameter optimization procedure tailored to transfer the results from the known classes to the novel classes with good generalization.
- We show that it is possible to accurately estimate the number of novel classes in the context of NCD, by applying simple clustering quality metrics in the latent space of NCD methods.
- We modify two classical unsupervised clustering algorithms to effectively utilize the data available in the NCD setting.
- We propose a simple and robust method, called PBN (for *Projection-Based NCD*), that learns a latent representation that incorporates the important features of the known classes, without overfitting on them.

The code is available at <https://github.com/PracticalNCD/ECMLPKDD2024>.

2 Related work

The setup of NCD [2], which involves both labeled and unlabeled data, can make it difficult to distinguish from the many other domains that revolve around similar concepts. In this section, we review some of the most closely related domains and try to highlight their key differences in order to provide the reader with a clear and comprehensive understanding of the NCD domain.

Semi-supervised Learning is another domain that is at the frontier between supervised and unsupervised learning. Specifically, a labeled set is given alongside an unlabeled set containing instances that are assumed to be from the same classes. Semi-supervised Learning can be particularly useful when labeled data is scarce or annotation is expensive. As unlabeled data is generally available in large quantities, the goal is to exploit it to obtain the best possible generalization performance given limited labeled data.

The main difference with NCD is that all the classes are known in advance. Some works have shown that the presence of novel classes in the unlabeled set negatively impacts the performance of Semi-Supervised Learning models [14, 15]. So as these works do not attempt to discover the novel classes, they are not applicable to NCD.

Transfer Learning aims at solving a problem faster or with better performance by leveraging knowledge from a different problem. It is commonly expressed in computer vision by pre-training models on ImageNet [1]. Transfer Learning can be either *cross-domain*, when a model trained on a given dataset is fine-tuned to perform the same task on a different (but related) dataset. Or it can be *cross-task*, where a model that can distinguish some classes is re-trained for other classes of the same domain.

NCD can be viewed as a cross-task Transfer Learning problem where the knowledge from a classification task on a source dataset is transferred to a clustering task on a target dataset. But unlike NCD, Transfer Learning typically requires the target spaces of both sets to be known in advance. Initially, NCD was characterized as a Transfer Learning problem (e.g. in DTC [16] and MCL [17]) and the training was done in two stages: first on the labeled set and then on the unlabeled set. This methodology seemed natural, as with Transfer Learning, both sets are not available at the same time.

Generalized Category Discovery (GCD) was first introduced by [11] and has also attracted attention from the community [12, 18, 19]. It can be seen as a less constrained alternative to NCD, since it does not rely on the assumption that samples belong exclusively to the novel classes during inference. However, this is a more difficult problem, as the models must not only cluster the novel classes, but also accurately differentiate between known and novel classes while correctly classifying samples from the known classes.

Some notable works in this area include ORCA [20] and OpenCon [21]. Namely, ORCA trains a discriminative representation by balancing a supervised loss on the known classes and unsupervised pairwise loss on the unlabeled data. And OpenCon proposes a contrastive learning framework which employs Out-Of-Distribution strategies to separate known vs. novel classes. Its clustering strategy is based on moving prototypes that enable the definition of positive and negative pairs of instances.

Novel Class Discovery has a rich body of papers in the domain of computer vision. Early works approached this problem in a *two-stage* manner. Some define a latent space using only the known classes, and project the unlabeled data into it (DTC[16] and MM [22]). Others train a pairwise labeling model on the known classes and use it to label and then cluster the novel classes (CCN [3] and MCL [17]). But both of these approaches suffered from overfitting on the known data when the high-level features were not fully shared by the known and novel classes.

Today, to alleviate this overfitting, the majority of approaches are *one-stage* and try to transfer knowledge from labeled to unlabeled data by learning a shared representation. In this category, AutoNovel [4] is one of the most highly influential works. After pre-training their latent representation with SSL [23], two classification networks are jointly trained. The first simply learns to distinguish the known classes with the ground-truth labels. And the other learns to separate unlabeled data from pseudo-labels defined for each epoch based on pairwise similarity. NCL [6] adopts the same architecture as AutoNovel, and extends the loss by adding a contrastive learning term to encourage the separation of novel classes. OpenMix [5] utilizes the MixUp strategy to generate more robust pseudo-labels.

As expressed before, although these methods have achieved some success, they are not applicable to tabular data. To date, and to the best of our knowledge, only TabularNCD [9] tackles this problem. Also inspired by AutoNovel, it pre-trains a dense-layer autoencoder with SSL and adopts the same loss terms and dual classifier architecture. Pseudo-labels are defined between pairs of unlabeled instances by checking if they are among the most similar pairs.

For a more complete overview of the state-of-the-art of NCD, we refer the reader to the survey [2].

3 Approaches

In this section, after introducing the notations, we define two simple but potentially strong models derived from classical clustering algorithms (Sections 3.2 and 3.3). The idea is to use the labeled data to improve the unsupervised clustering process, and make the comparison to NCD methods more challenging. Then, we present a new method, PBN (for Projection-Based NCD, Section 3.4), characterized by its low number of hyperparameters needed to be tuned.

3.1 Problem setting

We start by describing the Novel Class Discovery setup and the necessary notations. Here, data is provided in two distinct sets: a labeled set of known classes $D^l = \{(x_i^l, y_i^l)\}_{i=1}^N$ with $x_i^l \in \mathcal{X} = \mathbb{R}^d$ and $y_i^l \in \mathcal{Y}^l = \{1, \dots, C^l\}$ the ground-truth labels of x_i^l . And an unlabeled set $D^u = \{(x_i^u)\}_{i=1}^M$ where the data samples $x_i^u \in \mathcal{X}$ are only from novel classes $\mathcal{Y}^u = \{1, \dots, C^u\}$, which are different but related to the known classes. In other words, there is no overlap between the known and novel classes, so $\mathcal{Y}^l \cap \mathcal{Y}^u = \emptyset$. The objective is to exploit the knowledge from D^l to accurately partition D^u into the C^u clusters of the novel classes.

Following previous research, we first assume that the number of novel classes C^u is known in advance, and later propose an approach to estimate this number in a later section.

3.2 NCD k-means

This is a straightforward method that takes inspiration from k -means++ [24], which is an algorithm for choosing the initial positions of the centroids (or cluster centers). In k -means++, the first centroid is chosen at random from the data points. Then, each new centroid is chosen iteratively from one of the data points with a probability proportional to the squared distance from the point’s closest centroid. The resulting initial positions of the centroids are generally spread more evenly, which yields appreciable improvement in the final error of k -means and convergence time.

As shown in Figure 1a, we naively adapt k -means++ to the NCD setting by defining C^l initial centroids. They are set as the mean class points of the known classes using the ground-truth labels. Then, we follow k -means++ and randomly select C^u new centroids in the unlabeled set, with similarly decreasing probability when closer to existing centroids. We found experimentally (see Appendix D) that, after the initialization is complete, the best accuracy is achieved when only the centroids of the novel classes are updated, and using the unlabeled data only. In other words, the data of the known classes is only used during the initialization of the new centroids, but not during the convergence phase. Intuitively, if the centroids of the known and novel data are updated together, they have a higher risk of drifting and capturing data of the other set.

Similarly to k -means++, we repeat the initialization process a few times and keep the centroids that achieved the smallest inertia. Note that, to stay consistent with the

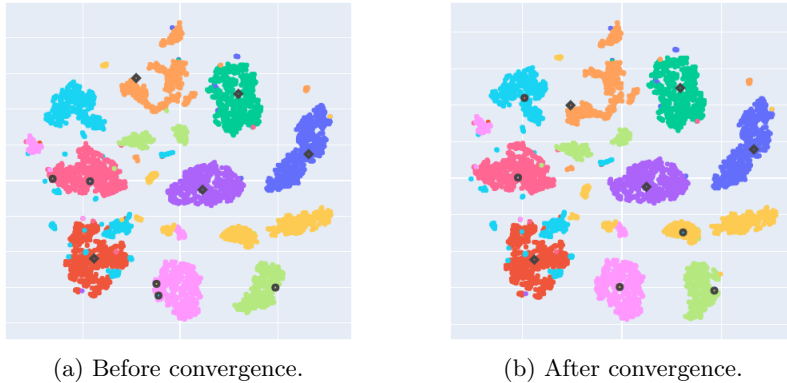


Fig. 1: t-SNE plots of the *Pendigits* dataset depicting the centroids before and after convergence. Note how the centroids of the known classes (the squares) don't move, as they stay the mean class point.

k -means algorithm, we use also the L_2 norm (i.e. the Euclidean distance) for NCD k -means.

3.3 NCD Spectral Clustering

Spectral Clustering (SC) is an alternative to distance-based clustering methods (such as k -means). It makes no assumptions about the structure of the data and considers the clustering problem as a graph partitioning problem and seeks to decompose the graph into connected components [25]. The input of the Spectral Clustering algorithm is an adjacency matrix (sometimes called similarity graph) which must accurately represent the neighborhood relationships between data points. There are multiple ways to construct such a graph, however there is no theoretical result on the relation between the similarity graph construction method and the Spectral Clustering results. Here, we employ a popular approach to construct the adjacency matrix, which is through a Gaussian kernel: $A_{i,j} = \exp(-\|x_i - x_j\|_2^2 / \sigma^2)$, $\forall x_i, x_j \in \mathcal{X}$ where the parameter σ controls the width of the neighborhood.

Following the Ng-Jordan-Weiss algorithm [26], we use the symmetric normalized Laplacian $L_{sym} = D^{-1/2}LD^{-1/2}$, where D is the degree matrix defined as: $D_{i,i} = \sum_{j=1}^n A_{i,j}$, in which n is the number of data samples. The next step consists in finding the first u eigenvectors of L_{sym} to form the spectral embedding $U \in \mathbb{R}^{n \times u}$, where row i corresponds to point x_i . Finally, the points in U are partitioned with k -means into clusters.

The optimal value of σ in the Gaussian kernel can vary widely depending on the distribution of inter-point distances. For this reason, we take inspiration from the rules of thumb given in [25] and employ a minimum spanning tree (MST) to choose σ . In the past few years, several graph-based clustering methods that use the MST have been proposed [27], as it reliably represents the layout of the data and is inexpensive to compute. In the approach proposed here, we denote d_{max} the length

of the longest edge in the MST of inter-point distances. The longest edge of the MST is a much studied object [28] and is representative of the scale of the dataset. The scaling factor σ is then calculated such that, after applying the Gaussian kernel, d_{max} is transformed into a chosen similarity s_{min} . This ensures that the resulting graph is safely connected. By optimizing this similarity s_{min} we can accurately represent the neighborhood relationships.

Therefore, for a given value of s_{min} , we derive σ from the length of the longest edge d_{max} in the MST:

$$\begin{aligned} \exp(-d_{max}^2/\sigma^2) &= s_{min} \\ \Leftrightarrow \sigma &= d_{max}/\sqrt{-\ln(s_{min})} \end{aligned} \quad (1)$$

Optimizing s_{min} instead of σ should be more robust to variations in the distribution of inter-point distances and give better results across different datasets or parts of a dataset.

To incorporate the knowledge from the known classes in the Spectral Clustering process, our initial approach was to utilize NCD k -means within the spectral embedding. In short, we would initially compute the full spectral embedding for all the data and determine the mean points of the known classes with the help of the ground truth labels. These mean points would then serve as the initial centroids. However, the observed performance improvement over the fully unsupervised SC was quite marginal.

Instead, the idea that we will use throughout this article (and that we refer to as ‘‘NCD Spectral Clustering’’) stems from the observation that SC can obtain very different results according to the parameters that are used. Among these parameters, the temperature parameter σ of the kernel holds particular importance, as it directly impacts the adjacency matrix’s accuracy in representing the neighborhood relationships of the data points. The rule of thumb of Equation 1 still requires to choose a value, but significantly reduces the space of possible values. Additionally, while the literature often sets the number of components u of the spectral embedding equal to the number of clusters, we have observed that optimizing it can also improve performance.

Therefore, rather than a specific method, we propose the parameter optimization scheme illustrated in Figure 2. For a given combination of parameters $\{s_{min}, u\}$, the




Fig. 2: NCD Spectral Clustering parameter optimization process.

corresponding spectral embedding of all the data is computed and then partitioned

with k -means. The quality of the parameters is evaluated from the clustering performance on the known classes, as the ground-truth labels are only available for the known classes. Indeed an important hypothesis behind the NCD setup is that the known and novel classes are related and share similarities, so they should have similar feature scales and distributions. Consequently, if the Spectral Clustering performs well on the known classes, the parameters are likely suitable to represent the novel classes.

Discussion. This idea can be applied to optimize the parameters of any unsupervised clustering algorithm in the NCD context. For example, the *Eps* and *MinPts* parameters of DBSCAN [29] can be selected in the same manner. It is also possible to use a different adjacency matrix in the SC algorithm. One option could be to substitute the Gaussian kernel with the k -nearest neighbor graph, and therefore optimise k instead of σ . However, for the sake of simplicity, we will only investigate SC using the Gaussian kernel.

3.4 Projection-Based NCD

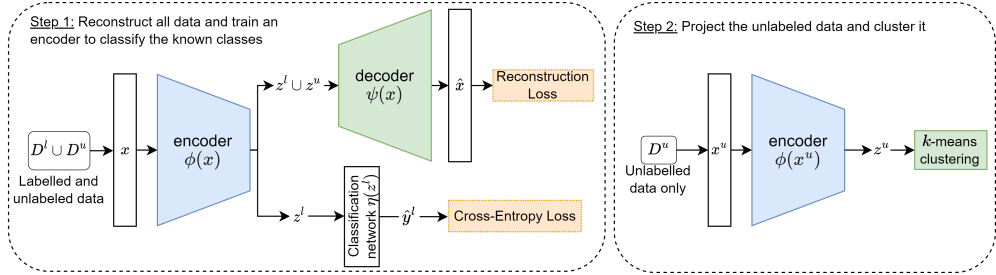


Fig. 3: Architecture of the PBN model.

Projection-based NCD (PBN) can be seen as an extension of the baseline method used in TabularNCD [9]. PBN is illustrated in Figure 3 and consists of 3 key components: (1) an encoder that learns a shared representation between the known and novel classes; (2) a classification network trained to distinguish the known classes of D^l in order to incorporate their relevant features into the representation; and (3) a decoder that reconstructs the data for both known and novel classes of $D^l \cup D^u$, ensuring that the latent space contains the information necessary to represent all classes. The decoder serves a dual purpose: it provides regularization and mitigates overfitting on the known classes, thus improving generalization, as shown in [30].

The training loss is defined as:

$$\mathcal{L}_{PBN} = w \times \mathcal{L}_{CE} + (1 - w) \times \mathcal{L}_{MSE} \quad (2)$$

where $w \in (0, 1)$ is a trade-off parameter that allows to balance the strength of the cross-entropy loss and the reconstruction loss.

The cross-entropy loss on the known classes is defined as:

$$\mathcal{L}_{CE} = - \sum_{c=1}^{C^l} y_c \log(\eta_c(z)) \quad (3)$$

where $\eta(z) = (\eta_c(z))_{c=1}^{C^l}$ is the output of a classification network composed of a single dense layer of neurons, $z = \phi(x)$ is the projection of instance x through the encoder ϕ and $(y_c)_{c=1}^{C^l}$ is the one-hot encoded ground-truth label of instance x .

The reconstruction loss of the instances from all classes is written as:

$$\mathcal{L}_{MSE} = \frac{1}{d} \sum_{j=1}^d (x_j - \hat{x}_j)^2 \quad (4)$$

where $\hat{x} = \psi(z)$ is the reconstruction of instance $x \in \mathbb{R}^d$.

Once the encoder, decoder and classification network have been trained (step 1), unlabeled data D^u is projected by the trained encoder into the latent space and then clustered with k -means to discover novel classes (step 2).

Projection-based NCD requires tuning of four hyperparameters. The trade-off parameter w is inherent to the method and the other three come from the choice of architecture: the learning rate, the dropout rate and the size of the latent space.

Note that this method doesn't employ complex schemes to define pseudo-labels unlike many NCD works. They have been proven to be accurate with image data (notably thanks to data augmentation techniques) [4, 31], but we found in preliminary results not detailed here, that for tabular data, they introduce variability in the results and new hyperparameters that need to be tuned.

Discussion. Similarly to PBN, the baseline method of TabularNCD [9] relies on the assumption that known and novel classes share similar high-level features, and defines a latent space that highlights these features. This baseline first trains a deep classifier to distinguish only the known classes of D^l . After training, the output and softmax layers are discarded, and the last hidden layer is now considered as the output of an encoder. It then projects the novel data of D^u into this latent space and partitions it using k -means. This is the basic workflow of two-stage latent space-based NCD methods identified in [2]. It is also similar to DTC [16], which uses the more refined DEC [32] clustering model instead of k -means in the baseline. The problem with such two-stage methods is that the resulting representations are at risk of being heavily biased towards the known classes. Thus, if some concepts or high-level features are not shared between the known and novel classes, the novel classes will not be well represented and these approaches will fail.

3.5 Summary of proposed approaches

In this section, we have proposed 3 distinct methods for solving the NCD problem, all of which leverage knowledge from the known classes in different ways. Firstly, NCD k -means uses the labeled data to improve the initialization of its centroids. Secondly,

instead of using the labels of the known classes during the clustering process itself, NCD Spectral Clustering uses them to find parameters that are likely to be suitable for the whole domain. More precisely, by clustering $D^l \cup D^u$ together, the adequacy of the parameters s_{min} and u can be evaluated on the known classes. Finally, PBN is a straightforward method that includes only the essential components to define a latent representation suitable for clustering the novel classes. In this case, an encoder is trained with a classification loss on the known classes and a reconstruction loss on all the data to ensure that the novel classes are not misrepresented. The novel data is then projected into this representation and clustered with k -means.

In the next section, we present an approach to finding hyperparameters without using the labels of the novel classes, which are not available in realistic scenarios. Indeed in the experiments (see Section 7), it should become clear why the simplicity of the proposed approach is a desirable feature for hyperparameter optimization in the NCD context.

4 Hyperparameter optimization

The success of machine learning algorithms (including NCD) can be attributed in part to the high flexibility induced by their hyperparameters. In most cases, a target is available and approaches such as the k -fold Cross-Validation (CV) can be employed to tune the hyperparameters and achieve optimal results. However, in a realistic scenario of Novel Class Discovery, the labels of the novel classes are never available. We must therefore find a way to optimize hyperparameters without ever relying on the labels of the novel classes. In this section, we present a method that leverages the known classes to find hyperparameters applicable to the novel classes. This tuning method is designed specifically for NCD algorithms that require both labeled data (known classes) and unlabeled data (novel classes) during training¹. This is the case for Projection-based NCD, as described in Section 3.4.

The process that we devised is represented in Figure 4. For each of the splits, the instances of around half of the known classes are selected to form the set D^{hid} and their labels are hidden. The labeled set now becomes the instances of $D^l \setminus D^{hid}$ and the unlabeled set becomes the instances of $D^u \cup D^{hid}$. After training the model with this new data split, it is evaluated for its performance for partitioning the instances of D^{hid} only since their labels are available.

To illustrate, in the split 1 of Figure 4, the model will be trained with the subsets of classes $\{C_2, C_3, C_4\}$ as known classes and $\{C_0, C_1, C_5, \dots, C_9\}$ as novel classes. It will be evaluated for its performance on the hidden classes $\{C_0, C_1\}$ only.

To evaluate a given combination of hyperparameters, this approach is applied to all the splits, and the performance on the hidden classes is averaged. After repeating this process for many combinations, the combination that achieved the best performance is selected. For the final evaluation on the novel classes, in a realistic scenario of NCD their labels are never available. However, in the datasets employed in this article, the novel classes are comprised of pre-defined classes. Therefore, even though these labels

¹To optimize purely unsupervised clustering methods for NCD, we refer the reader to the optimization process of Section 3.3.




Fig. 4: The k -fold cross-validation approach for hyperparameter optimisation of NCD methods.

are not employed during training, they can still be used to assess the final performance on the novel classes of different models and compare them against each other.

This tuning method stems from the same idea behind the NCD Spectral Clustering parameterization process. Namely, if the clustering in the learned representation successfully partitions the hidden classes in D^{hid} , it is also likely suitable for the novel classes in D^u . Furthermore, keeping the unlabeled data during training even though the model is not evaluated on the novel classes is important, as it increases the chances of the representation being adapted for all the classes. For the same reason, the k -means of PBN (see Section 3.4) is fitted on $D^u \cup D^{hid}$ together (instead of just D^{hid}) and the performance is then computed on D^{hid} only. So cases where the classes in D^u and D^{hid} are tangled will be penalized.

In Table 1, we report for all datasets used in our experiments the number of known classes that are hidden in each split, as well as the number of splits. Note that when the number of known classes is small (e.g. 3 for *Human*), this approach may be difficult to apply.

Discussion. Similarly to NCD, there are no labels available in unsupervised clustering problems, which makes the task of hyperparameter selection very difficult. To address this issue, clustering algorithms are sometimes tuned using *internal* metrics that do not rely on labeled data for computation. These metrics offer a means of comparing the results obtained from different clustering approaches. Examples of such metrics include the Silhouette coefficient, Davies-Bouldin index, or Calinski-Harabasz index [33]. However, it is important to note that these metrics make assumptions about the structure of the data and can be biased towards algorithms which make a similar assumption. But unlike unsupervised clustering, the NCD setting provides known classes that are related to the novel classes we are trying to cluster.

Table 1: Classes splits of the k -fold cross-validation.

Dataset	Known classes	Novel classes	Hidden classes	Splits
Human	3	3	2	3
Letter	19	7	7	5
Pendigits	5	5	2	5
Census	12	6	6	5
m feat	5	5	2	5
Optdigits	5	5	2	5
CNAE-9	4	5	2	5

5 Estimating the number of novel classes

Cluster Validity Indices (CVIs) are commonly used in unsupervised data analysis to estimate the number of clusters and are also applicable to the NCD problem. CVIs are scores that compare the compactness and separation of clusters without the help of external information such as ground truth labels. However, the knowledge from the known classes isn't used if the CVIs are directly applied to estimate the number of novel classes. Therefore, we propose to apply the CVIs in the latent representation learned by PBN. Projection-based NCD methods such as PBN are designed to create a latent space that emphasizes the relevant features of the known classes. Since these features are shared to some extent with the novel classes, this representation should be better at revealing the clusters we are trying to discover than the original feature space. Consequently, it makes sense that applying the different estimation techniques in the learned latent space should yield better results.

Note that this is only applicable to NCD methods such as PBN that don't require the number of novel classes C^u to train their latent space (unlike TabularNCD). For the others, the estimation can be done once in the original feature space, but should have higher error.

Some NCD works have also previously attempted to estimate the number of novel classes. For instance, [3] defines a large number of output neurons in their clustering network (e.g. 100). In this case, the clustering network is expected to use only the necessary number of clusters while leaving the remaining output neurons unused. Clusters were counted if they contained more instances than a certain threshold. However, since, with the exception of TabularNCD, the models studied in this paper do not use a clustering network, we will not evaluate this method.

Another technique, proposed by [11], consists in training a k -means on the combined dataset $D^l \cup D^u$ and selecting the k that yielded the highest accuracy on D^l . While this approach worked well for balanced datasets [12], it has been shown to underperform in the case of unbalanced class distributions [34]. For the sake of simplicity, we will call this method KM-ACC (for k -means ACC) in the remainder of this paper.

To select the CVI that we will use for our application, we rely on the results of [33]. Here, the authors conducted an extensive performance evaluation of 30 CVIs. They concluded that the Silhouette, Davies–Bouldin, Calinski–Harabasz and Dunn indices behaved better than other indices in almost all cases. In the experiments, the performance of these 4 indices will be compared, with the addition of the elbow method and the NCD-specific method KM-ACC.

6 Full training procedure

In the previous sections, we presented the models, the hyperparameter optimization and the estimation procedure of the number of novel classes independently. In this section, these components are brought together to form a complete training procedure. To ensure that no prior knowledge about the novel classes is ever used in this process, the number of novel classes is naturally estimated during the k -fold CV introduced in Section 4. As the whole process is quite complex, we try to summarize it in clear terms in this section and in Algorithm 1.

To gauge a given set of hyperparameters, we evaluate the performance of the model over n_{folds} , where in each fold, a random combination of known classes is “hidden” and merged with the unlabeled data of D^u . In a fold, the encoder of the NCD model is first trained on this new data split. The number of novel classes is then estimated with a CVI in the projection of the unlabeled data. At this point, the novel and hidden classes are partitioned by the model in the latent space using the previous estimate of the number of clusters. And the accuracy for this fold is calculated on the hidden classes. This process is repeated for all folds and for many combinations of hyperparameters, and the combination that achieved the best performance on average is selected for the final evaluation of the model.

Algorithm 1 Agnostic NCD model evaluation

Require: Training data $\{D^l, D^u\}$, hyperparameters θ , number of classes to hide n_{hid} , number of folds n_{folds}

- 1: **Initialize:** $folds \leftarrow$ set of n_{folds} random combinations of n_{hid} known classes
- 2: **for each fold in folds do**
- 3: $D^{hid} \leftarrow$ the data from D^l of the classes in $fold$
- 4: $D^{l'} \leftarrow D^l \setminus D^{hid}$
- 5: $D^{u'} \leftarrow D^u \cup D^{hid}$
- 6: Train model on $\{D^{l'}, D^{u'}\}$ with hyperparameters θ
- 7: $Z^u \leftarrow \phi_\theta(X^u)$ the projection of the novel data
- 8: $k' \leftarrow$ the estimation of C^u in Z^u with a CVI
- 9: Get the clustering prediction of the model for $D^{u'}$ using $k = n_{hid} + k'$
- 10: $ACC_{hid} \leftarrow$ clustering performance on D^{hid}
- 11: **end for**
- 12: **Return:** Average of all the ACC_{hid}

7 Experiments

7.1 Experimental setup

Datasets. To evaluate the performance of the methods compared in this paper, 7 tabular classification datasets were selected: Human Activity Recognition [35], Letter Recognition [36], Pen-Based Handwritten Digits [37], 1990 US Census Data [37], Multiple Features [37], Handwritten Digits [37] and CNAE-9 [37].

Following the previous NCD works [4, 6, 16], the instances of about 50% of the classes are hidden a priori to form the unlabeled set of novel classes D^u , while the rest form the labeled set D^l . We use a 70/30% train/test split if it was not already provided. Statistical information on the datasets is shown in Table 2, and the number of known/novel classes (along with the number of classes hidden during the k -fold CV) can be found in Table 1. The numerical features of all the datasets are pre-processed to have zero mean and unit variance, while the categorical features are one-hot encoded.

Table 2: Details of the datasets.

Dataset	Human	Letter	Pendigits	Census	m feat	Optdigits	CNAE-9
Features	562	16	16	67	515	62	856
Known classes	3	19	5	12	5	5	4
Known data	3733	10229	3777	12000	802	1918	377
Novel classes	3	7	5	6	5	5	5
Novel data	3619	3770	3717	6000	798	1905	487
Test data	1453	1704	1734	6000	202	905	113

Metrics. We report the *clustering accuracy* (ACC) on the unlabeled data. It is defined as:

$$ACC = \max_{perm \in P} \frac{1}{M} \sum_{i=1}^M \mathbb{1}\{y_i = perm(\hat{y}_i)\} \quad (5)$$

where y_i and \hat{y}_i are the ground truth labels and predicted labels for instance x_i respectively. Here, $M = |D^u|$. P is the set of all possible permutations between ground truth and predicted labels. It can be easily computed using the Hungarian algorithm [38]. The Normalized Mutual Information (NMI) and Adjusted Rand Index (ARI) are also reported in Annex A.

Competitors. We report the performance of k -means and Spectral Clustering, along with their NCD-adapted versions introduced in Sections 3.2 and 3.3. We also include PBN and the pioneering NCD work for tabular data, TabularNCD [9]. Finally, we implement the same baseline used in [9]. It is a simple deep classifier that is first trained to distinguish the known classes. Then, the penultimate layer is used to project the data of the novel classes before clustering it with k -means. See the discussion of Section 3.4 for more details. We will call this approach the “baseline” for the remainder of this article.

Implementation details. The neural-network based methods (i.e. PBN, TabularNCD and the baseline) are trained with the same architecture: an encoder of 2

hidden layers of decreasing size, and a final layer for the latent space whose dimension is optimized as an hyperparameter. The dropout probability and learning rate are also hyperparameters to be optimized. The classification networks are all a single linear layer followed by a softmax layer. All methods are trained for 200 epochs and with a fixed batch size of 512. For a fair comparison, the hyperparameters of these neural-network based methods are optimized following the process described in Section 4 (whereas the parameters of NCD SC are simply optimized following Section 3.3). Thus, the labels of the novel classes are never used except for the computation of the evaluation metrics reported in the result tables. The hyperparameters are optimized by maximizing the ARI, and their values can be found in Annex B.

7.2 Results analysis

7.2.1 Results when the number of novel classes is known in advance

Clustering. In Table 3, we first examine the performance of the unsupervised clustering methods when the number of novel classes C^u is known in advance. The aim is to determine which of the clustering algorithms performs best and should be compared with the NCD methods. We observe that NCD k -means is never worse than k -means, and NCD SC is only once worse than SC. This result confirms the efficacy of both NCD approaches and demonstrates that even simple clustering techniques can benefit from the known classes, although the improvements are sometimes only marginal.

The comparison between NCD k -means and NCD SC confirms the idea that no single clustering algorithm is universally better than the others in all scenarios, as noted by [39]. However, NCD SC outperforms its competitors on 4 occasions and has a the highest average accuracy. Therefore, this algorithm is selected for the next step of comparisons.

Table 3: Test ACC of the clustering algorithms averaged over 10 runs.

Dataset	k -means	NCD k -means	SC	NCD SC
Human	75.7±0.2	75.9±0.0	76.3±0.3	93.1±9.7
Letter	50.7±0.2	51.9±2.3	55.9±0.0	57.4±5.8
Pendigits	81.7±0.0	81.7±0.0	83.0±0.0	81.7±2.7
Census	49.9±4.0	50.4±1.1	48.5±0.3	48.0±1.8
m feat	89.1±0.3	89.7±0.4	89.6±0.3	89.2±2.3
Optdigits	79.1±4.5	94.2±0.0	89.7±0.0	95.4±5.3
CNAE-9	60.6±5.9	61.2±4.5	53.8±4.8	69.0±6.7
Average	69.5	72.1	71.0	76.3

NCD. As shown in Table 4, PBN outperforms both the baseline and TabularNCD by an average of 21.6% and 12.9%, respectively. It is only outperformed by the baseline on the *Letter Recognition* [36] dataset. This dataset consists of primitive numeric attributes describing the 26 capital letters in the English alphabet, which suggests a

high feature correlation between the features used to distinguish the known and novel classes. Since the baseline learns a latent space that is strongly discriminative for the known classes, this gives the baseline model a distinct advantage in this specific context. On the other hand, we observe that it is at a disadvantage when the datasets do not share as many high-level features between the known and novel classes.

Table 4 also demonstrates the remarkable competitiveness of the NCD Spectral Clustering method, despite its low complexity. On average, it trails behind PBN by only 1.0% in ACC and manages to outperform PBN twice over 7 datasets.

Table 4: Test ACC of the NCD methods averaged over 10 runs.

Dataset	Baseline	NCD SC	TabularNCD	PBN
Human	71.6±1.7	93.1±9.7	72.2±2.6	76.7±1.8
Letter	64.9±2.6	57.4±5.8	62.1±3.0	62.4±2.0
Pendigits	53.4±6.6	81.7±2.3	57.0±6.0	82.8±0.6
Census	59.1±0.8	48.0±1.8	45.2±4.8	62.4±0.9
m feat	66.7±4.1	89.2±2.3	90.2±2.7	91.7±0.8
Optdigits	40.7±5.1	95.4±5.3	73.0±8.4	92.6±2.3
CNAE-9	40.2±3.2	69.0±6.7	51.3±5.2	72.6±4.6
Average	55.7	76.3	64.4	77.3

To investigate the reasons behind the subpar performance of TabularNCD, we look at the correlation between the ARI on the hidden classes and the final ARI of the model on the novel classes. A strong correlation would imply that if a combination of hyperparameters performed well on the hidden classes, it would also perform well on the novel classes. To examine this, we plot the average ARI on the hidden classes against the ARI on the novel classes. Figure 5 is an example of such a plot. It shows that, in the case of the *Letter Recognition* dataset, PBN has a much stronger correlation than TabularNCD. We attribute this difference to the large number of hyperparameters of TabularNCD (7, against 4 for PBN), which causes the method to overfit on the hidden classes, resulting in a lack of effective transfer of hyperparameters to the novel classes.




Fig. 5: Comparison between the ARI on the hidden and novel classes. Each point is a different hyperparameter combination.

In conclusion, this section has shown that when the number C^u of novel classes is known, NCD SC performs almost as well as PBN. Therefore, in this specific scenario, NCD SC is a viable candidate for addressing the NCD problem due to its lower complexity and shorter training time. Conversely, despite its strong learning capacity, TabularNCD is penalized by its high number of hyperparameters.

7.2.2 Results when the number of novel classes is estimated

As expressed in Section 5, we leverage the representation learned by PBN during the k -fold CV to estimate the number of clusters. The result is a method that never relies on any kind of knowledge from the novel classes. This approach is also applicable to the baseline and the spectral embedding of SC, but not to TabularNCD as it requires a number of clusters to be defined during the training of its representation. For a fair comparison, TabularNCD is trained here with a number of clusters that was estimated beforehand with a CVI.

To determine which CVI will perform the best in this application, we estimate the number of classes in the latent spaces learned by PBN when C^u was known in advance. Figure 6 displays the average ranks of the CVIs in the latent spaces, and the details of the results can be found in Annex C. The Nemenyi post-hoc test was used to compare the methods against all others. However, given the relatively small number of datasets, the Critical Difference (CD) is large and the CVIs are not statistically different from each another according to the Nemenyi test.




Fig. 6: Comparison of the CVIs in the latent space of PBN using the Nemenyi test with a 95% confidence interval.

Nevertheless, we find that the Silhouette coefficient performed the best in the latent space of PBN, closely followed by the Calinski-Harabasz index. The elbow method is ranked last, which could be explained by the difficulty in defining an elbow ².

To summarize, in the following results, the Silhouette coefficient will be used for all estimations of C^u . The full training procedure described in Section 6 will be used to train the baseline and PBN, with C^u estimated in their latent spaces as detailed in Algorithm 1. For TabularNCD, C^u is estimated once in the original feature space (see Annex C Table C4 for the values used). And NCD SC is trained as it was described in Section 3.3, but with C^u estimated in its spectral embedding.

²There are no widely accepted approaches, as the concept of an “elbow” is subjective. In this study, we employed the *kneedle* algorithm [40] through the *kneed* Python Library [41].

Table 5: Test ACC averaged over 10 runs. With C^u estimated with the Silhouette coefficient.

Dataset	Baseline	NCD SC	TabularNCD	PBN
Human	70.8±2.9	30.2±4.2	71.1±0.0	71.1±0.0
Letter	64.0±6.1	34.8±2.3	41.8±4.9	61.3±4.7
Pendigits	46.7±3.6	74.1±2.1	57.0±6.0	83.0±0.3
Census	56.6±3.6	29.0±3.5	35.7±1.6	49.8±0.1
m feat	59.5±7.7	73.2±3.9	41.1±0.2	90.6±2.2
Optdigits	42.1±4.6	79.5±4.0	96.9±1.3	90.5±4.8
CNAE-9	33.8±3.8	44.6±3.9	39.3±0.2	50.8±1.5
Average	53.4	52.2	54.7	71.0

As emphasised earlier, Table 5 reports the results of the different NCD algorithms in the most realistic scenario possible, where both the labels and the number of novel classes are not known in advance. This is the setting where PBN exhibits the greatest improvement in performance compared to the other competitors, achieving an ACC that is 17.6%, 18.8% and 16.3% higher than the baseline, NCD SC, and TabularNCD, respectively.

Remarkably, TabularNCD outperforms PBN on the Optdigits datasets where the number of clusters was overestimated by the Silhouette coefficient in the original feature space. This suggests that TabularNCD probably only utilized the output neurons necessary for clustering, leaving the others unused, which was the method proposed in [3] for estimating C^u . This, however, is not true for the Letter dataset where C^u was significantly overestimated, indicating that accurate estimations will likely result in improved performance.

Compared to the case where C^u is known in advance, the ACC of the baseline falls from 55.7% to 53.4% and NCD SC falls from 74.3% to 52.2%. This shows that they are both unable to find a latent space suitable for the estimation of C^u .

However, the ACC of PBN remains an impressive 71.0%, demonstrating that this simple method, comprising of only two loss terms, is the most appropriate for tackling the NCD problem in a realistic scenario. In contrast to the baseline and NCD SC methods, its reconstruction term enables it find a latent space where the unlabeled data are correctly represented. And unlike TabularNCD, it has a low number of hyperparameters which decreases the probability of overfitting on the hidden classes during the k -fold CV procedure.

8 Conclusion

In this article, we have shown that in the NCD setting, unsupervised clustering algorithms can benefit from knowledge of the known classes and reliably improve their performance by implementing simple modifications. We have also introduced a novel NCD algorithm called PBN, which is characterized by its simplicity and low number of hyperparameters, which proved to be a decisive advantage under realistic conditions. In addition, we have proposed an adaptation of the k -fold cross-validation process to tune the hyperparameters of NCD methods without depending on the labels of the novel classes. Finally, we have demonstrated that the number of novel classes can be

accurately estimated within the latent space of PBN. These two previous contributions have shown that the NCD problem can be solved in realistic situations where no prior knowledge of the novel classes is available during training.

Declarations

Funding

Colin Troisemaine, Alexandre Reiffers-Masson, Stéphane Gosselin, Vincent Lemaire and Sandrine Vaton received funding from Orange SA.

Competing Interests

Colin Troisemaine, Stéphane Gosselin and Vincent Lemaire received research support from Orange SA. Alexandre Reiffers-Masson and Sandrine Vaton received research support from IMT Atlantique.

Ethics approval

Not applicable.

Consent to participate

All authors have read and approved the final manuscript.

Consent for publication

Not applicable.

Availability of data and materials

All data used in this study are available publicly online. The datasets were extracted directly in the repositories available with the links in the corresponding section.

Code availability

The code for experiments is available at the following url: <https://github.com/PracticalNCD/ECMLPKDD2024>.

Authors' contributions

Colin Troisemaine, Alexandre Reiffers-Masson, Stéphane Gosselin, Vincent Lemaire and Sandrine Vaton contributed to the manuscript equally.

References

- [1] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR, pp. 248–255 (2009)
- [2] Troisemaine, C., Lemaire, V., Gosselin, S., Reiffers-Masson, A., Flocon-Cholet, J., Vaton, S.: Novel class discovery: an introduction and key concepts. ArXiv (2023)
- [3] Hsu, Y.-C., Lv, Z., Kira, Z.: Learning to cluster in order to transfer across domains and tasks. In: ICLR (2018)
- [4] Han, K., Rebuffi, S.-A., Ehrhardt, S., Vedaldi, A., Zisserman, A.: Autonovel: Automatically discovering and learning novel visual categories. PAMI (2021)
- [5] Zhong, Z., Zhu, L., Luo, Z., Li, S., Yang, Y., Sebe, N.: Openmix: Reviving known knowledge for discovering novel visual categories in an open world. In: CVPR, pp. 9462–9470 (2021)
- [6] Zhong, Z., Fini, E., Roy, S., Luo, Z., Ricci, E., Sebe, N.: Neighborhood contrastive learning for novel class discovery. In: CVPR (2021)
- [7] Sun, Y., Shi, Z., Liang, Y., Li, Y.: When and how does known class help discover unknown ones? provable understanding through spectral analysis. In: ICML, vol. 202, pp. 33014–33043 (2023)
- [8] Li, Z., Otholt, J., Dai, B., Hu, D., Meinel, C., Yang, H.: A closer look at novel class discovery from the labeled set. In: NeurIPS 2022 Workshop on Distribution Shifts: Connecting Methods and Applications (2022)
- [9] Troisemaine, C., Flocon-Cholet, J., Gosselin, S., Vaton, S., Reiffers-Masson, A., Lemaire, V.: A method for discovering novel classes in tabular data. In: ICKG, pp. 265–274 (2022)
- [10] Caron, M., Touvron, H., Misra, I., Jegou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: ICCV, pp. 1–21 (2021)
- [11] Vaze, S., Han, K., Vedaldi, A., Zisserman, A.: Generalized category discovery. In: CVPR, pp. 7492–7501 (2022)
- [12] Fei, Y., Zhao, Z., Yang, S., Zhao, B.: Xcon: Learning with experts for fine-grained category discovery. In: British Machine Vision Conference (BMVC) (2022)
- [13] Zhang, L., Qi, L., Yang, X., Qiao, H., Yang, M.-H., Liu, Z.: Automatically Discovering Novel Visual Categories with Self-supervised Prototype Learning (2022)

- [14] Chen, Y., Zhu, X., Li, W., Gong, S.: Semi-supervised learning under class distribution mismatch. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, pp. 3569–3576 (2020)
- [15] Guo, L.-Z., Zhang, Z.-Y., Jiang, Y., Li, Y.-F., Zhou, Z.-H.: Safe deep semi-supervised learning for unseen-class unlabeled data. In: ICML, pp. 3897–3906 (2020)
- [16] Han, K., Vedaldi, A., Zisserman, A.: Learning to discover novel visual categories via deep transfer clustering. In: ICCV (2019)
- [17] Hsu, Y.-C., Lv, Z., Schlosser, J., Odom, P., Kira, Z.: Multi-class classification without multi-class labels. In: ICLR (2019)
- [18] Zheng, J., Li, W., Hong, J., Petersson, L., Barnes, N.: Towards open-set object detection and discovery. In: CVPR, pp. 3961–3970 (2022)
- [19] Yang, M., Zhu, Y., Yu, J., Wu, A., Deng, C.: Divide and conquer: Compositional experts for generalized novel class discovery. In: CVPR, pp. 14268–14277 (2022)
- [20] Cao, K., Brbic, M., Leskovec, J.: Open-world semi-supervised learning. In: ICLR (2022)
- [21] Sun, Y., Li, Y.: Opencon: Open-world contrastive learning. In: TMLR (2023)
- [22] Chi, H., Liu, F., Yang, W., Lan, L., Liu, T., Han, B., Niu, G., Zhou, M., Sugiyama, M.: Meta discovery: Learning to discover novel classes given very limited data. In: ICLR (2022)
- [23] Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: ICLR (2018)
- [24] Arthur, D., Vassilvitskii, S.: K-means++ the advantages of careful seeding. In: ACM-SIAM SODA, pp. 1027–1035 (2007)
- [25] Luxburg, U.: A tutorial on spectral clustering. *Stat. Comput.* **17**(4), 395–416 (2007)
- [26] Ng, A., Jordan, M., Weiss, Y.: On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems* **14** (2001)
- [27] Khan, A.A., Mohanty, S.K.: A fast spectral clustering technique using mst based proximity graph for diversified datasets. *Information Sciences* **609**, 1113–1131 (2022)
- [28] Stuetzle, W.: Estimating the cluster tree of a density by analyzing the minimal spanning tree of a sample. *Journal of classification* **20**(1), 25–47 (2003)

- [29] Ester, M., Kriegel, H.-P., Sander, J., Xu, X., *et al.*: A density-based algorithm for discovering clusters in large spatial databases with noise. In: Kdd, vol. 96, pp. 226–231 (1996)
- [30] Le, L., Patterson, A., White, M.: Supervised autoencoders: Improving generalization performance with unsupervised regularizers. *Advances in neural information processing systems* **31** (2018)
- [31] Zhao, B., Han, K.: Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. In: *Advances in Neural Information Processing Systems* (2021)
- [32] Xie, J., Girshick, R., Farhadi, A.: Unsupervised deep embedding for clustering analysis. In: ICML, vol. 48, pp. 478–487 (2016)
- [33] Arbelaitz, O., Gurrutxaga, I., Muguerza, J., Pérez, J.M., Perona, I.: An extensive comparative study of cluster validity indices. *Pattern recognition* **46**(1), 243–256 (2013)
- [34] Yang, M., Wang, L., Deng, C., Zhang, H.: Bootstrap your own prior: Towards distribution-agnostic novel class discovery. In: CVPR, pp. 3459–3468 (2023)
- [35] Anguita, D., Ghio, A., Oneto, L., Parra, X., Reyes-Ortiz, J.L.: A public domain dataset for human activity recognition using smartphones. In: ESANN (2013)
- [36] Frey, P.W., Slate, D.J.: Letter recognition using holland-style adaptive classifiers. *Machine Learning* **6**, 161–182 (2005)
- [37] Dua, D., Graff, C.: UCI Machine Learning Repository (2017)
- [38] Kuhn, H.W., Yaw, B.: The hungarian method for the assignment problem. *Naval Res. Logist. Quart.* 83–97 (1955)
- [39] Von Luxburg, U., Williamson, R.C., Guyon, I.: Clustering: Science or art? In: ICML Workshop on Unsupervised and Transfer Learning, pp. 65–79 (2012)
- [40] Satopaa, V., Albrecht, J., Irwin, D., Raghavan, B.: Finding a” kneedle” in a haystack: Detecting knee points in system behavior. In: ICDCS Workshops, pp. 166–171 (2011). IEEE
- [41] Arvai, K.: kneed. Zenodo (2023). <https://doi.org/10.5281/zenodo.7873825>

Appendix A Additional result metrics

In this section, in addition to the ACC results discussed in Section 7.2.2, we present the NMI (Table A1) and ARI (Table A2) for the NCD methods when the number of novel classes C^u is unknown and has to be estimated. Our results are consistent with those shown in Table 5: the PBN method largely outperforms its competitors in

this realistic scenario. In particular, PBN achieves an average NMI higher than the baseline, NCD SC and TabularNCD by 23.5%, 8.6% and 13.5% respectively. Similarly, the ARI is 24.5%, 12.4% and 17.5% higher on average.

Table A1: Test NMI averaged over 10 runs. With C^u estimated with the Silhouette coefficient.

Dataset	Baseline	NCD SC	TabularNCD	PBN
Human	56.1±8.9	45.8±3.0	75.2±0.0	75.2±0.0
Letter	54.4±3.3	52.6±1.0	37.2±3.2	59.2±2.4
Pendigits	41.0±6.7	79.5±1.4	48.0±6.5	73.4±2.5
Census	59.7±0.6	33.5±1.0	42.6±1.5	60.6±0.5
m feat	50.0±3.9	71.9±2.9	44.8±2.1	79.1±3.3
Optdigits	28.5±6.0	77.3±2.0	93.1±2.0	84.9±2.0
CNAE-9	23.0±1.5	56.5±2.0	42.0±1.3	45.2±13.4
Average	44.7	59.6	54.7	68.2

Table A2: Test ARI averaged over 10 runs. With C^u estimated with the Silhouette coefficient.

Dataset	Baseline	NCD SC	TabularNCD	PBN
Human	48.7±4.0	20.3±3.0	61.4±0.0	61.4±0.0
Letter	44.3±4.1	29.5±2.2	23.4±5.6	48.9±3.1
Pendigits	29.8±5.7	73.7±1.7	37.6±6.5	65.3±3.4
Census	42.6±3.7	23.0±3.8	26.7±0.8	35.5±0.2
m feat	40.9±5.3	64.1±4.5	21.5±0.2	79.1±4.2
Optdigits	16.9±6.5	75.6±3.6	94.1±2.7	84.4±4.4
CNAE-9	11.6±1.9	33.2±4.0	18.8±0.4	31.5±1.7
Average	33.5	45.6	40.5	58.0

Appendix B Hyperparameters

The Table B3 shows the hyperparameters found by the full procedure described in Section 6.

Appendix C Cluster Validity Indices numerical results

An estimate of the number of clusters in the 7 datasets considered in this paper can be found in Table C4. Among the 6 CVIs reported here, the Silhouette coefficient performed the best. Furthermore, compared to the original feature space, its average estimation error significantly decreased in the latent space, validating our approach. For some datasets, the Davies-Bouldin index continued to decrease and the Dunn index continued to increase as the number of clusters increased, resulting in very large overestimations. Note that the estimates of the number of novel classes in Table C4 are

Table B3: Hyperparameters of the methods found when C^u is estimated.

	Parameter	Human	Letter	Pendigits	Census	m features	Optdigits	CNAE-9
NCD Base-line	latent dim	560	9	9	30	34	42	569
	lr	0.0001543	0.0013334	0.0055179	0.0020210	0.0001187	0.0033302	0.0001676
	dropout	0.1680448	0.1400953	0.0525056	0.4678363	0.5648556	0.0751809	0.2577655
	s_{min}	0.293960	0.981371	0.861477	0.635344	0.700053	0.204129	0.831326
SC	u	144	14	18	53	16	26	14
	k neighbors	70	29	73	45	15	48	7
TabularNCD	latent dim	372	14	13	56	98	62	791
	lr	0.009614	0.000365	0.004648	0.001248	0.008400	0.007554	0.001296
	dropout	0.154362	0.066725	0.278239	0.274898	0.013843	0.091586	0.082204
	top k	0.610971	0.165173	0.467394	0.939568	0.590449	0.995869	0.942106
	w_1	0.047453	0.413677	0.139925	0.312476	0.301939	0.785324	0.196314
	w_2	0.702652	0.970951	0.605445	0.882650	0.834274	0.933564	0.641289
	latent dim	504	22	12	13	172	53	579
PBN	learning rate	0.000101	0.000588	0.001076	0.002111	0.004675	0.000364	0.004400
	dropout	0.509847	0.027456	0.011264	0.167778	0.239342	0.066065	0.081317
	w	0.467146	0.722411	0.106714	0.960869	0.670342	0.189179	0.698152

not needed in the experiments of Section 7.2.2, since Algorithm 1 directly incorporates such estimates in the training procedure. This table has only helped us to identify the most appropriate CVI for our problem. The only exception is the TabularNCD method, which requires an a priori estimation of the number of novel classes in the original feature space.

Table C4: An estimation of the number of novel classes with some CVIs in the latent space of PBN.

Dataset	Human	Letter	Pendigits	Census	m feat	Optdigits	CNAE-9
Ground-truth	3	7	5	6	5	5	5
PBN latent space							
Silhouette	2	8	5	3	5	5	5
CH	2	3	5	4	2	2	5
Dunn	2	3	98	3	2	95	2
KM ACC	1	2	3	3	5	6	1
Davies-B.	2	63	6	3	99	4	96
Elbow	9	14	10	7	16	13	9
Original feature space							
Silhouette	2	45	5	3	2	9	2

Appendix D NCD k -means centroids convergence study

In this appendix, we aim to determine how to achieve the best performance with NCD k -means. Specifically, after the centroid initialization described in Section 3.2, we investigate: (1) whether it is more effective to update the centroids of both known and novel classes, or only the centroids of novel classes; (2) whether the centroids need to be updated using data from both known and novel classes, or only using data from novel classes. The results are presented in Table D5 and show that for 5 out of 7 datasets, the best results are obtained when only the centroids of the novel classes are

updated on the unlabeled data. Updating the centroids of the known classes always leads to worse performance, as the class labels are not used in this process. Thus, the centroids of the known classes run the risk of capturing data from the novel classes (and vice versa).

Table D5: ACC of NCD k -means averaged over 10 runs.

Dataset	Converging the centroids...	On unlabeled data only	On labeled and unlabeled data
Human	novel only	75.9±0.0	77.4±0.0
	known and novel	-	75.1±0.5
Letter	novel only	51.9±2.3	39.5±1.9
	known and novel	-	42.3±2.6
Pendigits	novel only	81.7±0.0	72.7±0.9
	known and novel	-	75.3±4.0
Census	novel only	50.4±1.1	50.4±4.8
	known and novel	-	44.6±8.3
m feat	novel only	89.7±0.4	69.1±0.2
	known and novel	-	84.1±7.0
Optdigits	novel only	94.2±0.0	70.8±7.8
	known and novel	-	74.0±14.7
CNAE-9	novel only	61.2±4.5	48.3±8.5
	known and novel	-	68.1±7.5