Episodic Learning in Nonparametric Few-shot Classification: A Case Study

Anonymous Author(s)
Affiliation
Address
email

Abstract

Episodic learning is a popular practice among researchers and practitioners interested in few-shot learning. It consists of organising training in a series of learning problems (or *episodes*), each divided into a small training and validation subset to mimic the circumstances encountered during evaluation. But is this always necessary? In this paper, we investigate the usefulness of episodic learning in methods which use nonparametric approaches, such as nearest neighbours, at the level of the episode. For these methods, we not only show how the constraints imposed by episodic learning are not necessary, but that they in fact lead to a data-inefficient way of exploiting training batches. We conduct a wide range of ablative experiments with Matching and Prototypical Networks, two of the most popular methods that use nonparametric approaches at the level of the episode. Their "non-episodic" counterparts are considerably simpler, and significantly improve their performance in multiple few-shot classification datasets.

1 Introduction

The problem of few-shot learning (FSL) – classifying examples from previously unseen classes given only a handful of training data – has considerably grown in popularity within the machine learning community in the last few years. The reason is likely twofold. First, being able to perform well on FSL problems is important for several applications, from learning new symbols [20] to drug discovery [2]. Second, since the aim of researchers interested in meta-learning is to design systems that can quickly learn novel concepts by generalising from previously encountered learning tasks, FSL benchmarks are often adopted as a practical way to empirically validate meta-learning algorithms.

To the best of our knowledge, there is not a widely recognised definition of meta-learning. In a recent survey, Hospedales et al. [19] informally describe it as *"the process of improving a learning algorithm over multiple learning episodes"*. In practical terms, following the compelling rationale that *"test and train conditions should match"* [41, 11], several seminal meta-learning papers (e.g. [41, 28, 12]) have emphasised the importance of organising training into *episodes*, i.e. learning problems with a limited amount of "training" (the *support* set) and "test" examples (the *query* set) to mimic the test-time scenario presented by FSL benchmarks.

However, several recent works (e.g. [7, 42, 9, 38]) showed that simple baselines can outperform established FSL meta-learning methods by using embeddings pre-trained with the standard cross-entropy loss, thus casting a doubt on the importance of episodes in FSL. Inspired by these results, we aim at understanding the practical usefulness of episodic learning in popular FSL methods relying on metric-based nonparametric classifiers such as Matching and Prototypical Networks [41, 35]. We chose this family of methods because they do not perform any adaptation at test time. This allows us to test the efficacy of episodic training without having to significantly change the baseline algorithms, which could potentially introduce confounding factors.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

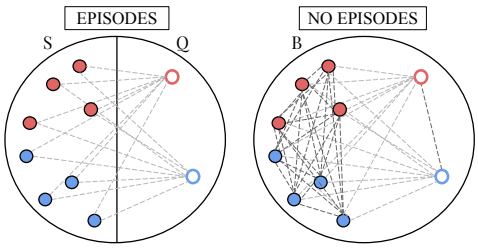

Figure 1: Difference in *batch exploitation* for metric-based methods between adopting or not adopting the concept of *episodes* during training, on an illustrative few-shot learning problem with 2 *ways* (classes), and 4 *shots* (examples) and 1 *query* per class.

	POSITIVES	NEGATIVES
NO EPISODES	$\binom{m+n}{2}w$	$\binom{w}{2}(m+n)^2$
EPISODES	wmn	$w(w-1)mn$
PAIRS LOST	$\frac{w}{2}(m^2+n^2-m-n)$	$\frac{w}{2}(w-1)(m^2+n^2)$

Table 1: The extra number of gradients that, on the *same batch*, a non-episodic method can exploit with respects its episodic counterpart grows quadratically as $O(w^2(m^2+n^2))$, where w is the number of ways, and n and m are the number of shots and queries per class.

In this work we perform a case study focussed on Matching Networks [41] and Prototypical Networks [35], and we show that within this family of methods episodic learning *a)* is detrimental for performance, *b)* is analogous to randomly discarding examples from a batch and *c)* introduces a set of superfluous hyperparameters that require careful tuning. Without episodic learning, these methods are closely related to the classic *Neighbourhood Component Analysis* (NCA) [16, 31] on deep embeddings and achieve, without bells and whistles, an accuracy that is competitive with recent methods on multiple FSL benchmarks: *mini*ImageNet, CIFAR-FS and *tiered*ImageNet.

2 Background and method

This section is divided as follows: Sec. 2.1 introduces episodic learning and illustrates a data efficiency issue encountered with nonparametric few-shot learners based on episodes; Sec. 2.2 introduces the losses from Snell et al. [35], Vinyals et al. [41] and Goldberger et al. [16] which we use throughout our experiments; and Sec. 2.3 explains the three options we explored to perform FSL classification with previously-trained feature embeddings.

2.1 Episodic learning

A common strategy to train FSL methods is to consider a distribution $\hat{\mathcal{E}}$ over possible subsets of labels that is as close as possible to the one encountered during evaluation $\mathcal{E}$ [1] [41]. Each *episodic batch* $B_E = \{S, Q\}$ is obtained by first sampling a subset of labels L from $\hat{\mathcal{E}}$, and then sampling images constituting both *support set* S and *query set* Q from the set of images with labels in L, where $S = \{(\mathbf{s}_1, y_1), \ldots, (\mathbf{s}_n, y_n)\}$, $Q = \{(\mathbf{q}_1, y'_1), \ldots, (\mathbf{q}_m, y_m)\}$, and S_k and Q_k denote the sets of images with label $y = k$ in the support set and query set respectively.

For most methods, this corresponds to training on a series of mini-batches in which each image belongs to *either* the support *or* the query set. Support and query sets are constructed such that they both contain all the classes of L, and a fixed number of images per class. Therefore, episodes are defined by three variables: the number of classes $w = |L|$ (the "ways"), the number of examples per class in the support set $n = |S_k|$ (the "shots"), and the number of examples per class in the query set $m = |Q_k|$. During evaluation, the set $\{w, n, m\}$ defines the problem setup. Instead, at training time $\{w, n, m\}$ can be seen as a set of hyperparameters controlling the batch creation, and that (as we will see in Sec. 3.2) requires careful tuning.

In a Maximum Likelihood Estimation framework, training on these episodes can be written as

$$\arg\max_{\theta} \mathop{\mathbb{E}}_{L\sim\hat{\mathcal{E}}} \mathop{\mathbb{E}}_{\substack{S\sim L \\ Q\sim L}} \left(\sum_{(q_i, y_i)\in Q} \log P_\theta \left(y_i | q_i, S, \rho \right) \right). \tag{1}$$

For the sake of brevity, with a slight abuse of notation we omit the function f_θ (e.g. a deep neural network) which is used to obtain a representation for the images in S and Q, and whose parameters θ

[1] Note that, in FSL, the sets of classes encountered during training and evaluation are disjoint.

are optimised during the training process. Note that the optimisation of Eq. 1 depends on an *optional* set of parameters ρ. This is obtained by an "inner" optimisation procedure, whose scope is limited to the current episode [19]. The idea is that the "outer" optimisation loop, by attending to a distribution of episodes, will appropriately shift the *inductive bias* of the algorithm located in the inner loop, thus learning how to learn [40]. In recent years, many interesting proposals have been made about what form ρ should have, and how it should be computed. For instance, in MAML [12] ρ takes the form of an update of the global parameters θ, while Ravi and Larochelle [28] learn to optimise by considering ρ as set of the hyper-parameters of the optimiser's update rule.

Other methods, such as Matching and Prototypical Networks [41, 35], avoid learning a separate set of parameters ρ altogether, and utilise a nonparametric learner (such as nearest neighbour classifiers) at the inner level. We chose to focus our case study on these methods not only because they have been seminal for the community, but also for ease of experimental design. Having $\rho = \varnothing$ considerably reduces the design complexity of the algorithm, thus allowing precise ablations to understand the efficacy of episodic learning without considerably changing the nature of the original algorithms.

Considerations on data efficiency. The constraints imposed by episodic learning on the role each image has in a training batch has subtle but important implications, illustrated in Fig. 1 by highlighting the number of distances contributing to the loss. By dividing batches between support and query set (S and Q) during training, *episodes* have the side effect of disregarding many of the distances between labelled examples that would constitute useful training signal for nonparametric FSL methods. More specifically, for metric-based nonparametric methods, the number of training distances that are lost in a batch because of the episodic strategy grows quadratically as $O(w^2(m^2+n^2))$ (derivation shown in Appendix A). Table 1 breaks down this difference in terms of gradients from positives and negatives distance pairs (which we simply refer to as *positives* and *negatives* throughout the rest of the paper). In a typical training batch with $w = 20$, $m = 15$ and $n = 5$ [35], ignoring the episodic constraints increases the number of both positives and negatives by more than 150%.

In the remainder of this paper, we conduct a case study to illustrate how this issue affects two of the most popular FSL algorithms relying on nonparametric approaches at the inner level: Prototypical Networks [35] and Matching Networks [41].

2.2 Loss functions

Prototypical Networks (PNs) [35] are one of the most popular and effective approaches in the few-shot learning literature, and they are at the core of several recent methods (e.g. [24, 15, 1, 43, 6]).

During *training*, episodes consisting of a support set S and a query set Q are sampled as described in Sec. 2.1. Then, a *prototype* for each class k is computed as the mean embedding of the samples from the support set belonging to that class: $\mathbf{c}_k = (1/|S_k|) \cdot \sum_{(\mathbf{s}_i, y_k) \in S_k} f_\theta(\mathbf{s}_i)$, where f_θ is a deep neural network with parameters θ learned via Eq. 1.

Let $C = \{(\mathbf{c}_1, y_1), \ldots, (\mathbf{c}_k, y_k)\}$ be the set of prototypes and corresponding labels. The loss can be written as follows:

$$\mathcal{L}_{\text{PNs}} = \frac{-1}{|Q|} \sum_{(\mathbf{q}_i, y_i) \in Q} \log \left(\frac{\exp -\|f_\theta(\mathbf{q}_i) - \mathbf{c}_{y_i}\|^2}{\sum_{k'} \exp -\|f_\theta(\mathbf{q}_i) - \mathbf{c}_{k'}\|^2} \right),$$

where k' is an index that goes over all classes.

Matching Networks (MNs) [41] are closely related to PNs in the multi-shot case and equivalent in the 1-shot case. Rather than aggregating the embeddings of the same class into prototypes, this loss directly computes a softmax over individual embeddings of the support set, as:

$$\mathcal{L}_{\text{MNs}} = \frac{-1}{|Q|} \sum_{\substack{(\mathbf{q}_i, y) \\ \in Q}} \log \left(\frac{\sum\limits_{\substack{\mathbf{s}_j \\ \in S_y}} \exp -\|f_\theta(\mathbf{q}_i) - f_\theta(\mathbf{s}_j)\|^2}{\sum\limits_{\substack{\mathbf{s}_k \\ \in S}} \exp -\|f_\theta(\mathbf{q}_i) - f_\theta(\mathbf{s}_k)\|^2} \right).$$

In their work, Vinyals et al. [41] use the cosine rather than the Euclidean distance. However, (as [35]) we observed that the Euclidean distance is a better choice for FSL problems, and thus we use it in all the losses of our experiments. Note that Vinyals et al. [41] also suggest a variant to $\mathcal{L}_{\text{MNs}}$ (MNs with "Full Context Embeddings"), where an LSTM (with an extra set of parameters) is used to condition the way the inputs are embedded in the current support set. In our experiments, we did not consider this variant as it falls in the category of *adaptive* episodic learning approaches ($\rho \neq \varnothing$, see Sec. 2.1).

Neighbourhood Component Analysis (NCA). $\mathcal{L}_{\text{MNs}}$ and $\mathcal{L}_{\text{PNs}}$ sum over the likelihoods that a query image belongs to the same class of a certain sample (or prototype) from the support set by computing the softmax over the distances between the query and the support samples (or prototypes). This is closely related to the *Neighbourhood Component Analysis* approach by Goldberger et al. [16] (and expanded to the non-linear case by Salakhutdinov et al. [31] and Frosst et al. [13]), except for a few important differences which we discuss at the end of this section.

Let $i \in [1, b]$ be the indices of the images within a batch B. The NCA loss can be written as:

$$\mathcal{L}_{\text{NCA}} = \frac{-1}{|B|} \sum_{i \in 1,\ldots,b} \log \left(\frac{\sum\limits_{\substack{j \in 1,\ldots,b \\ j \neq i \\ y_i = y_j}} \exp - \left\| \mathbf{z}_i - \mathbf{z}_j \right\|^2}{\sum\limits_{\substack{k \in 1,\ldots,b \\ k \neq i}} \exp - \left\| \mathbf{z}_i - \mathbf{z}_k \right\|^2} \right),$$

where $\mathbf{z}_i = f_\theta(\mathbf{x}_i)$ is an image embedding and y_i its corresponding label. By minimising this loss, distances between embeddings from the same class will be minimised, while distances between embeddings from different classes will be maximised. Importantly, note how the concepts of support set and query set here do not exist. More simply, the images (and respective labels) constituting the batch $B = \{(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_b, y_b)\}$ are sampled uniformly.

Given the similarity between these three losses, and considering that PNs and MNs do not perform episode-specific parameter adaptation, $\{w, m, n\}$ can be simply interpreted as the set of hyper-parameters controlling the sampling of mini-batches during training. More specifically, PNs, MNs and NCA differ under three aspects:

I. First and foremost, due to the nature of episodic learning, PNs and MNs only consider pairwise distances between the query and the support set (Fig. 1 left); NCA instead uses *all* the distances within a batch and treats each example in the same way (Fig. 1 right).
II. Only PNs rely on the creation of prototypes.
III. Because of how L, S and Q are sampled in episodic learning (Eq. 1), for PNs and MNs some images might be sampled more frequently than others (sampling "with replacement"). NCA instead visits every image of the dataset once for each epoch (sampling "without replacement").

To investigate the effects of these three differences, in Sec. 3 we conduct a wide range of experiments.

2.3 Few-shot classification during evaluation

Once f_θ has been trained, there are many possible ways to perform few-shot classification during evaluation. In this paper we consider three simple approaches that are particularly intuitive for embeddings learned via metric-based losses like the ones described in Sec. 2.2. Note that, in the 1-shot case, all the evaluation methods considered coincide.

k-NN. To classify an image $\mathbf{q}_i \in Q$, we first compute the Euclidean distance to each support point $\mathbf{s}_j \in S$: $d_{ij} = \| f_\theta(\mathbf{q}_i) - f_\theta(\mathbf{s}_j)) \|^2$. Then, we simply assign $y(\mathbf{q}_i)$ to be majority label of the k nearest neighbours. A downside here is that k is a hyper-parameter that has to be chosen, although a reasonable choice in the FSL setup is to set it equal to the number of "shots" n.

Nearest centroid. Similar to k-NN, we can perform classification by inheriting the label of the closest class centroid, i.e. $y(\mathbf{q}_i) = \arg\min_{j \in \{1,\ldots,k\}} \| f_\theta(\mathbf{x}_i) - \mathbf{c}_j \|$. This is the approach used by Prototypical Networks [35] and SimpleShot [42].

Soft assignments. To classify an image $\mathbf{q}_i \in Q$, we compute the values

$$p_{ij} = \frac{\exp(-\|f_\theta(\mathbf{q}_i) - f_\theta(\mathbf{s}_j))\|^2)}{\sum_{\mathbf{s}_k \in S} \exp(-\|f_\theta(\mathbf{q}_i) - f_\theta(\mathbf{s}_k)\|^2)}$$

for all $\mathbf{s}_j \in S$, which is the probability that i inherits its class from j. We then compute the likelihood for each class k: $\sum_{s_j \in S_k} p_{ij}$, and choose the class with the highest likelihood $y(\mathbf{q}_i) = \arg\max_k \sum_{s_j \in S_k} p_{ij}$. This is the approach for classification adopted by the original NCA paper [16] and Matching Networks [41].

We experiment with all three alternatives and observe that the *nearest centroid* approach is the most effective (details available in Appendix D). For this reason, unless differently specified, we use it as default in our experiments.

3 Experiments

In the following, Sec. 3.1 describes our experimental setup; Sec. 3.2 shows the important effect of the hyperparameters controlling the creation of episodes; in Sec. 3.3 we compare the episodic strategy to randomly discarding pairwise distances within a batch; in Sec. 3.4 we perform a set of ablations to better illustrate the relationship between PNs, MNs and NCA; finally, in Sec. 3.5 we compare our version of the NCA to several recent methods.

3.1 Experimental setup

We conduct our experiments on *mini*ImageNet [41], CIFAR-FS [4] and *tiered*ImageNet [30], using the popular ResNet-12 variant first adopted by Lee et al. [21] as embedding function f_θ [2] . A detailed description of benchmarks, architecture and choice of hyperparameters is deferred to Appendix F, while below we discuss the most important choices of the experimental setup.

Like Wang et al. [42], for all our experiments (including those with Prototypical and Matching Networks) we centre and normalise the feature embeddings before performing classification, as it is considerably beneficial for performance. After training, we compute the mean feature vectors of all the images in the training set: $\bar{\mathbf{x}} = \frac{1}{|\mathcal{D}^{\text{train}}|} \sum_{\mathbf{x} \in \mathcal{D}^{train}} \mathbf{x}$. Then, all feature vectors in the test set are updated as $\mathbf{x}_i \leftarrow \mathbf{x}_i - \bar{\mathbf{x}}$, and normalised by $\mathbf{x}_i \leftarrow \frac{\mathbf{x}_i}{\|\mathbf{x}_i\|}$.

As standard [19], performance is assessed on episodes of 5-way, 15-query and 1- or 5-shot. Each model is evaluated on 10,000 episodes sampled from the validation or the test set. To further reduce the variance, we trained each model *three times* with three different random seeds, for a total of 30,000 episodes per configuration, from which 95% confidence intervals are computed.

3.2 Batch size and episodes

Despite Prototypical and Matching Networks being among the simplest FSL methods, the creation of episodes requires the use of several hyperparameters ($\{w, m, n\}$, Sec. 2.1) which can significantly affect performance. Snell et al. [35] state that the number of shots n between training and testing should match and that one should use a higher number of ways w during training. In their experiments, they train 1-shot models with $w = 30$, $n = 1$, $m = 15$ and 5-shot models with $w = 20$, $n = 5$, $m = 15$, with batch sizes of 480 and 400, respectively. Since the corresponding batch sizes of these configurations differ, making direct comparisons between them is difficult.

Instead, to directly compare configurations across batch sizes, we define an episode by its number of shots n, the batch size b and the total number of images per class $m + n$ (the sum of elements across support and query set). For example, if we train a 5-shot model with $m + n = 8$ and $b = 256$, its corresponding training episodes will have $n = 5$, $m = 8 - n = 3$, and $w = 256/(m + n) = 32$. Using this notation, we train configurations of PNs and MNs covering several combinations of these hyperparameters, so that the resulting batch size corresponding to an episode is 128, 256 or 512. Then, we train three configurations of the NCA, where the sole hyperparameter is the batch size b.

[2] Note that, since we do not use a final linear layer for classification, our backbone is in fact a ResNet-11.

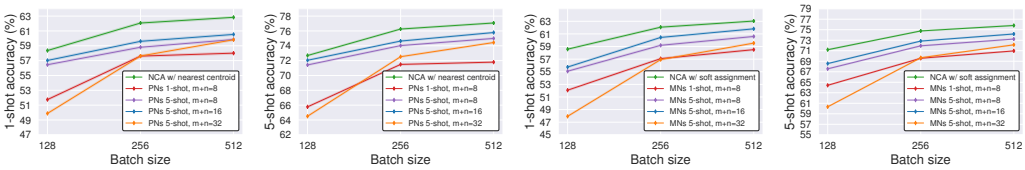

Figure 2: 1-shot and 5-shot accuracies on CIFAR-FS (val. set) for Prototypical and Matching Networks models trained with different episodic configurations: 1-shot with $m + n$=8 and 5-shot with $m + n$=8, 16 or 32. NCA models are trained on batches of size 128, 256 and 512 to match the size of the episodes. Reported values correspond to the mean accuracy of three models trained with different random seeds and shaded areas represent 95% confidence intervals. See Sec. 3.2 for details.

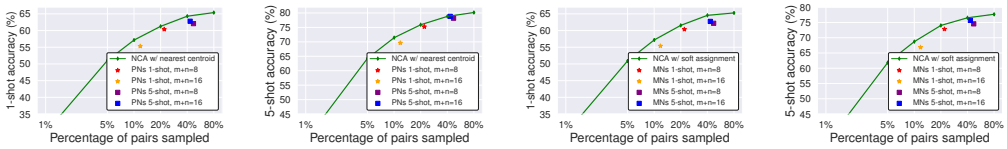

Figure 3: 1-shot and 5-shot accuracies on *mini*ImageNet (val. set) for NCA models trained by only sampling a fraction of the total number of available pairs in the batch (of size 256). Stars and squares represent models trained using the Prototypical or Matching Network loss, and are plotted on the x-axis based on the total number of distance pairs exploited in the loss, so that they can be directly compared with this "sub-sampling" version of the NCA. Reported values correspond to the mean accuracy of three models trained with different random seeds and shaded areas represent 95% confidence intervals. See Sec. 3.3 for details.

Results for CIFAR-FS can be found in Fig. 2, where we report results for NCA, PNs and MNs with $m + n = 8$, 16 or 32. Results for *mini*ImageNet observe the same trend and are deferred to Appendix H. For consistency in our comparisons, we evaluate performance using a *nearest centroid* classifier when comparing against PNs, and *soft assignments* when comparing against MNs (see Sec. 2.3). Note that PNs and MNs results for 1-shot with $m + n = 16$ and $m + n = 32$ are not reported, as they fare significantly worse. The 1-shot $m + n = 16$ is 4% worse in the best case compared to the lowest lines in Fig. 2, and the $m + n = 32$ is 10% worse in the best case. This is likely because these two setups exploit the fewest number of pairs among all the setups, which leads to the least training signal being available. In Appendix E we discuss whether the difference in performance between the different episodic batch setups of Fig. 2 can be solely explained by the differences in the number of distance pairs used in the batch configurations. We indeed find that generally speaking the higher the number of pairs the better. However, one should also consider the positive/negative balance and the number of classes present within a batch.

Several things can be observed from Fig. 2. First, NCA-trained embeddings perform better than *all* configurations, no matter the batch size. Second, PNs and MNs are very sensitive to different hyperparameter configurations. For instance, with batches of size 128, PNs trained with episodes of 5-shot and m+n=32 perform worse than a PNs trained with 5-shot episodes and m+n=16. Note that, as we will show in Table 2, the best episodic configurations for PNs and MNs found with this hyperparameter search is superior to the setting used in the original papers.

3.3 Episodic batches vs. random sub-sampling

Despite the inferior performance with respect to the NCA, one might posit that, by training on episodes, PNs and MNs can somehow make better use of a smaller number of distances within a batch. This could be useful, for instance, in situations where it is important to train with very large batches. Given the increased conceptual complexity and the extra hyperparameters, the efficacy of episodic learning (in cases where a smaller number of distances should be considered) should be validated against the much simpler approach of random subsampling. We perform an experiment where we train NCA models by randomly discarding a fraction of the total number of distances used in the loss. Then, for comparison, we include different PNs and MNs models, after having computed to which percentage of discarded pairs (in a normal batch) their episodic batches correspond to.

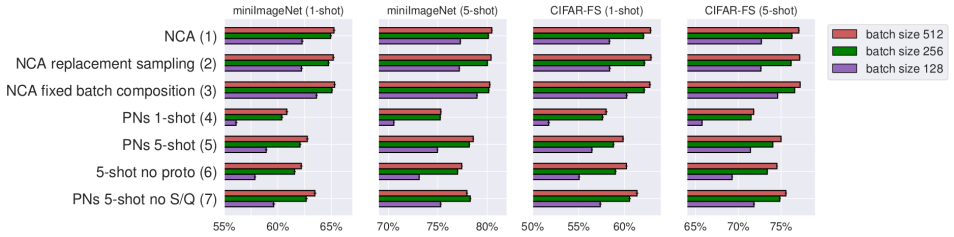

Figure 4: Ablation experiments on NCA and Prototypical Networks, both on batches (or episodes) of size 128, 256, and 512 on *mini*ImageNet and CIFAR-FS (val. set). Reported values correspond to the mean accuracy of three models trained with different random seeds and error bars represent 95% confidence intervals. See Sec. 3.4 for details.

Results can be found in Fig. 3. As expected, we can see how subsampling a fraction of the total available number of pairs within a batch negatively affects performance. More importantly, we can notice that the points representing PNs and MNs models lie very close to the under-sampling version of the NCA. This suggests that the episodic strategy is roughly equivalent, empirically, to only exploiting a fraction of the distances available in a batch. Note how, moving along the x-axis of Fig. 3, variants of PNs and MNs exploiting a higher number of distances perform better.

3.4 Ablation experiments

To better analyse why NCA performs better, in this section we consider the three key differences discussed at the end of Sec. 2.2 by performing a series of ablations on models trained on batches of size 128, 256 and 512. Results are summarised in Fig. 4. We refer the reader to Appendix B to obtain detailed steps describing how these ablations affect the losses of Sec. 2.2.

First, we compare two variants of the NCA: one in which the sampling of the training batches happens sequentially and without replacement, as is standard in supervised learning, and one where batches are sampled with replacement. This modification (row 1 and 2 of Fig. 4) has a negligible effect, meaning that the replacement sampling introduced by episodic learning should not interfere with the other ablations. We then perform a series of ablations on episodic batches, i.e. sampled with the method described in Sec. 2.1. To obtain a reasonably-performing model for both 1- and 5-shot models, we use configurations with $m + n = 8$. This means that, for PNs and MNs, models are trained with 8 images per class, and either 16, 32 or 64 classes (batches of size 128, 256 and 512 respectively). The batch size for NCA is also set to either 128, 256, or 512, allowing direct comparison.

The ablations of Fig. 4 compare PNs to NCA. First, we train standard PNs models (row 4 and 5 of Fig. 4). Next, we train a model where "prototypes" are not computed (row 6). This implies that, similar to what happens in MNs, distances are considered between individual points, but a separation between query and support set remains. This ablation allows us to investigate if the loss in performance by PNs compared to NCA can be attributed to prototype computation during training (which turned out not to be the case). Then, we perform an ablation where we ignore the separation between support and query set, and compute the NCA on the union of the support and query set, while still computing prototypes for the points that would belong to the support set (row 7). Last, we perform an ablation where we consider all the previous points together: we sample with replacement, we ignore the separation between support and query set and we do not compute prototypes (row 3). This amounts to the NCA loss, except that it is computed on batches with a fixed number of classes and a fixed number of images per class (row 3). Notice that in Fig. 4 there is only one row dedicated to 1-shot models. This is because we cannot generate prototypes from 1-shot models, so we cannot have a "no proto" ablation. Furthermore, for 1-shot models the "no S/Q" ablation is equivalent to the NCA with a fixed batch composition.

From Fig. 4, we can see that disabling prototypes (row 6) negatively affects the performance of 5-shot (row 5), albeit slightly. Since for PNs the amount of gradient signal is the same with (row 5, Fig.4) or without (row 6, Fig.4) the computation of prototypes, we believe that this could be motivated by the increased misalignment between the training and test setup present in the ablation of row 6. Nonetheless, enabling the computation between all pairs increases the performance (row 7) and, importantly, enabling *all* the ablations (row 3) completely recovers the performance lost by PNs. Note

	*mini*ImageNet		CIFAR-FS		*tiered*ImageNet	
	1-shot	**5-shot**	**1-shot**	**5-shot**	**1-shot**	**5-shot**
Episodic methods						
adaResNet [23]	56.88 ± 0.62	71.94 ± 0.57	-	-	-	-
TADAM[24]	58.50 ± 0.30	76.70 ± 0.30	-	-	-	-
Shot-Free [29]	60.71± n/a	77.64± n/a	69.2± n/a	84.7± n/a	63.52± n/a	82.59± n/a
TEAM [26]	60.07± n/a	75.90± n/a	-	-	-	-
MTL [36]	61.20 ± 1.80	75.50 ± 0.80	-	-	-	-
TapNet [43]	61.65 ± 0.15	76.36 ± 0.10	-	-	63.08 ± 0.15	80.26 ± 0.12
MetaOptNet-SVM[21]	62.64 ± 0.61	78.63 ± 0.46	$\mathbf{72.0 \pm 0.7}$	84.2 ± 0.5	65.99 ± 0.72	81.56 ± 0.53
Variatonal FSL [44]	61.23 ± 0.26	77.69 ± 0.17	-	-	-	-
Simple cross-entropy baselines						
Transductive finetuning [9]	62.35 ± 0.66	74.53 ± 0.54	70.76 ± 0.74	81.56 ± 0.53	-	-
RFIC-simple [38]	62.02 ± 0.63	$\mathbf{79.64 \pm 0.44}$	71.5 ± 0.8	86.0 ± 0.5	$\mathbf{69.74 \pm 0.72}$	$\mathbf{84.41 \pm 0.55}$
Meta-Baseline [8]	$\mathbf{63.17 \pm 0.23}$	79.26 ± 0.17	-	-	68.62 ± 0.27	83.29 ± 0.18
Our implementations:						
MNs ([35] episodes)	58.91 ± 0.12	72.48 ± 0.10	69.28 ± 0.13	80.79 ± 0.10	65.75 ± 0.13	78.40 ± 0.10
PNs ([35] episodes)	59.78 ± 0.12	75.42 ± 0.09	69.94 ± 0.12	84.01 ± 0.09	65.80 ± 0.13	81.26 ± 0.10
MNs (our episodes)	60.77 ± 0.12	73.82 ± 0.09	71.86 ± 0.13	82.41 ± 0.10	66.53 ± 0.13	79.08 ± 0.10
PNs (our episodes)	61.32 ± 0.12	77.77 ± 0.09	70.41 ± 0.12	84.61 ± 0.10	66.89 ± 0.14	82.20 ± 0.09
SimpleShot [42]	62.16 ± 0.12	78.33 ± 0.09	69.98 ± 0.12	84.40 ± 0.09	66.67 ± 0.14	81.57 ± 0.10
NCA soft assignment (ours)	62.55 ± 0.12	76.93 ± 0.11	$\mathbf{72.49 \pm 0.12}$	83.38 ± 0.09	68.35 ± 0.13	81.04 ± 0.09
NCA nearest centroid (ours)	62.55 ± 0.12	78.27 ± 0.09	$\mathbf{72.49 \pm 0.12}$	$\mathbf{85.15 \pm 0.09}$	68.35 ± 0.13	83.20 ± 0.10

Table 2: Comparison of methods that use ResNet12 as f_θ, on the test set of *mini*ImageNet, CIFAR-FS, and *tiered*ImageNet. Values are reported with 95% confidence intervals. For our methods, reported values correspond to the mean accuracy of three models trained with different random seeds.

the meaningful gap in performance between row 1 and 3 in Fig. 4 for batch size 128, which disappers for batch size 512. This is likely due to the number of positives available in an excessively small batch size. Since our vanilla NCA creates batches by simply sampling images randomly from the dataset, there is a limit to how small a batch can be (which depends on the number of classes of the dataset). As an example, consider the extreme case of a batch of size 4. For the datasets considered, it is very likely that such a batch will contain no positive pairs for some classes. Conversely, the NCA ablation with a fixed batch composition (i.e. with a fixed number of images per class) will have a higher number of positive pairs (at the cost of a reduced number of classes per batch). This can explain the difference, as positive pairs constitute a less frequent (and potentially more informative) training signal. In Appendix E we extend this discussion, commenting on the role of positive and negative distances. In Appendix H we also report the results of a second set of ablations to compare NCA and Matching Networks, which are analogous to the ones with Prototypical Networks we just described and lead to the same conclusions.

These experiments highlight that the separation of roles between the images belonging to support and query set, which is typical of episodic learning [41], is detrimental for the performance of metric-based nonparametric few-shot learners. Instead, using the NCA loss on standard mini-batches allows full exploitation of the training data and significantly improves performance. Moreover, the NCA has the advantage of simplifying the overall training procedure, as the hyperparameters for the creation of episodes $\{w, n, m\}$ no longer need to be considered.

3.5 Comparison with recent methods

We now evaluate our models on three popular FSL datasets to contextualise their performance with respect to the recent literature. When considering which methods to compare against, we chose those *a)* which have been recently published, *b)* that use a ResNet-12 architecture [21] (the most commonly used), and *c)* with a setup that is not significantly more complicated than ours. For example, we only report results for the main approach of Tian et al. [38]. We omit their self-distillation [14] variant, as it can be applied to most methods and involves multiple stages of training.

Results can be found in Table 2. Besides the results for the NCA loss, we also report PNs and MNs results with both the episodic setup from Snell et al. [35] and the best one (batch size 512, 5-shot, $m + n$=16 for both PNs and MNs) found from the experiment of Fig. 2, which brings a considerable improvement over the original. Note that our aim is not to improve the state of the art, but rather to shed light on the practice of episodic learning. Nonetheless, our vanilla NCA is competitive and sometimes even superior to recent methods, despite being extremely simple. It fares surprisingly well against methods that use meta-learning (and episodic learning), and also against the high-performing simple baselines based on pre-training with the cross-entropy loss. Moreover, because of the explicit inductive bias that it encodes in terms of relative position in the embedding space of samples from the same class, the NCA loss is a useful tool to consider *alongside* cross-entropy trained baselines.

4 Related work

Pioneered by Utgoff [39], Schmidhuber [33, 34], Bengio et al. [3] and Thrun [37], the general concept of meta-learning is several decades old (for a survey see [40, 19]). However, in the last few years it has experienced a surge in popularity, becoming the most used paradigm for learning from very few examples. Several methods addressing the FSL problem by learning on episodes were proposed. MANN [32] uses a Neural Turing Machine [18] to save and access the information useful to meta-learn; Bertinetto et al. [5] and Munkhdalai et al. [22] propose a deep network in which a "teacher" branch is tasked with predicting the parameters of a "student" branch; Matching Networks [41] and Prototypical Networks [35] are two nonparametric methods in which the contributions of different examples in the support set are weighted by either an LSTM or a softmax over the cosine distances for Matching Networks, and a simple average for Prototypical Networks; Ravi and Larochelle [28] propose instead to use an LSTM to learn the hyperparameters of SGD, while MAML [12] learns to fine-tune an entire deep network by backpropagating through SGD. Despite these works widely differing in nature, they all stress on the importance of organising training in a series of small learning problems (*episodes*) that are similar to those encountered during inference at test time.

In contrast with this trend, a handful of papers have recently shown that simple approaches that forego episodes and meta-learning can perform well on FSL benchmarks. These methods all have in common that they pre-train a feature extractor with the cross-entropy loss on the "meta-training classes" of the dataset. Then, at test time a classifier is adapted to the support set by weight imprinting [25, 9], fine-tuning [7], transductive fine-tuning [9] or logistic regression [38]. Wang et al. [42] suggest performing test-time classification by using the label of the closest centroid to the query image. Unlike these papers, which propose new methods, we are more focussed on shedding some light on the possible causes behind the inefficiency of popular nonparametric few-shot learning algorithms such as Prototypical and Matching Networks.

Despite maintaining a support and a query set, the work of Raghu et al. [27] is similar in spirit to ours, and modifies episodic learning in MAML, showing that performance is almost entirely preserved when only updating the network head during meta-training and meta-testing. In this paper, we focussed on FSL algorithms just as established, and uncovered inefficiencies that not only allow for notable conceptual simplifications, but that also bring a significant boost in performance. Finally, two related but different works are the ones of Goldblum et al. [17] and Fei et al. [10]. The former addresses PNs' poorly representative samples by training on episodic pairs with the same classes (but different instances) and using a regularizer enforcing consistency across them. The latter investigates meta-learning methods with *parametric* base learners, and shows interesting findings on the importance of having tightly clustered classes in feature space, which inspires a regularizer that improves non meta-learning models.

5 Conclusion

Towards the aim of understanding the reasons behind the poor competitiveness of meta-learning methods with respect to simple baselines, in this paper we investigate the role of episodes in popular nonparametric few-shot learning methods. We found that their performance is highly sensitive to the set of hyperparameters used to sample these episodes. By replacing the Prototypical Networks and Matching Networks losses with the closely related (and non-episodic) Neighbourhood Component Analysis, we were able to ignore these hyperparameters, while significantly improving the few-shot classification accuracy. We found out that the performance discrepancy is in large part caused by the separation between support and query set within each episode, which negatively affects the number of pairwise distances contributing to the loss. Moreover, with nonparametric few-shot approaches, the episodic strategy is almost empirically equivalent to randomly discarding a portion of the distances available within a batch. Finally, we showed that our variant of the NCA achieves an accuracy on multiple popular FSL benchmarks that is competitive with recent methods, making it a simple and appealing baseline for future work.

Broader impact. Upon reflection, we believe progress in few-shot learning is beneficial for society, as it can significantly impact important problems such as drug discovery and medical imaging. We also recognise that the capability of leveraging very small datasets might constitute a threat if deployed for surveillance (e.g. identity recognition) by entities hostile to civil liberties.

References

[1] K. R. Allen, E. Shelhamer, H. Shin, and J. B. Tenenbaum. Infinite mixture prototypes for few-shot learning. In *International Conference on Machine Learning*, 2019.

[2] H. Altae-Tran, B. Ramsundar, A. S. Pappu, and V. Pande. Low data drug discovery with one-shot learning. *ACS central science*, 2017.

[3] S. Bengio, Y. Bengio, J. Cloutier, and J. Gecsei. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*. Univ. of Texas, 1992.

[4] L. Bertinetto, J. F. Henriques, P. H. Torr, and A. Vedaldi. Meta-learning with differentiable closed-form solvers. In *International Conference on Learning Representations*, 2019.

[5] L. Bertinetto, J. F. Henriques, J. Valmadre, P. Torr, and A. Vedaldi. Learning feed-forward one-shot learners. In *Advances in Neural Information Processing Systems*, 2016.

[6] T. Cao, M. Law, and S. Fidler. A theoretical analysis of the number of shots in few-shot learning. In *International Conference on Learning Representations*, 2020.

[7] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. F. Wang, and J.-B. Huang. A closer look at few-shot classification. *arXiv preprint arXiv:1904.04232*, 2019.

[8] Y. Chen, X. Wang, Z. Liu, H. Xu, and T. Darrell. A new meta-baseline for few-shot learning. *arXiv preprint arXiv:2003.04390*, 2020.

[9] G. S. Dhillon, P. Chaudhari, A. Ravichandran, and S. Soatto. A baseline for few-shot image classification. In *International Conference on Learning Representations*, 2020.

[10] N. Fei, Z. Lu, T. Xiang, and S. Huang. Melr: Meta-learning via modeling episode-level relationships for few-shot learning. In *International Conference on Learning Representations*, 2021.

[11] C. Finn. Stanford cs330: Multi-task and meta-learning, 2019 | lecture 4 - non-parametric meta-learners. `https://www.youtube.com/watch?v=bc-6tzTyYcM&list=PLoROMvodv4rMC6zfYmnD7UG3LVvwaITY5&index=4`, 2019. Accessed on 26/01/2021.

[12] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, 2017.

[13] N. Frosst, N. Papernot, and G. Hinton. Analyzing and improving representations with the soft nearest neighbor loss. *arXiv preprint arXiv:1902.01889*, 2019.

[14] T. Furlanello, Z. Lipton, M. Tschannen, L. Itti, and A. Anandkumar. Born again neural networks. In *International Conference on Machine Learning*, 2018.

[15] S. Gidaris, A. Bursuc, N. Komodakis, P. Pérez, and M. Cord. Boosting few-shot visual learning with self-supervision. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[16] J. Goldberger, G. E. Hinton, S. T. Roweis, and R. R. Salakhutdinov. Neighbourhood components analysis. In *Advances in Neural Information Processing Systems*, 2005.

[17] M. Goldblum, S. Reich, L. Fowl, R. Ni, V. Cherepanova, and T. Goldstein. Unraveling meta-learning: Understanding feature representations for few-shot tasks. In *International Conference on Machine Learning*. PMLR, 2020.

[18] A. Graves, G. Wayne, and I. Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.

[19] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey. Meta-learning in neural networks: A survey. *arXiv preprint arXiv:2004.05439*, 2020.

[20] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 2015.

[21] K. Lee, S. Maji, A. Ravichandran, and S. Soatto. Meta-learning with differentiable convex optimization. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[22] T. Munkhdalai and H. Yu. Meta networks. In *International Conference on Machine Learning*, 2017.

[23] T. Munkhdalai, X. Yuan, S. Mehri, and A. Trischler. Rapid adaptation with conditionally shifted neurons. In *International Conference on Machine Learning*, 2018.

[24] B. Oreshkin, P. R. López, and A. Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. In *Advances in Neural Information Processing Systems*, 2018.

[25] H. Qi, M. Brown, and D. G. Lowe. Low-shot learning with imprinted weights. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

[26] L. Qiao, Y. Shi, J. Li, Y. Wang, T. Huang, and Y. Tian. Transductive episodic-wise adaptive metric for few-shot learning. In *IEEE International Conference on Computer Vision*, 2019.

[27] A. Raghu, M. Raghu, S. Bengio, and O. Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of maml. In *International Conference on Learning Representations*, 2020.

[28] S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2017.

[29] A. Ravichandran, R. Bhotika, and S. Soatto. Few-shot learning with embedded class models and shot-free meta training. In *IEEE International Conference on Computer Vision*, 2019.

[30] M. Ren, E. Triantafillou, S. Ravi, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel. Meta-learning for semi-supervised few-shot classification. In *International Conference on Learning Representations*, 2018.

[31] R. Salakhutdinov and G. Hinton. Learning a nonlinear embedding by preserving class neighbourhood structure. In *Artificial Intelligence and Statistics*, 2007.

[32] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. Lillicrap. Meta-learning with memory-augmented neural networks. In *International Conference on Machine Learning*, 2016.

[33] J. Schmidhuber. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. PhD thesis, Technische Universität München, 1987.

[34] J. Schmidhuber. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 1992.

[35] J. Snell, K. Swersky, and R. Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, 2017.

[36] Q. Sun, Y. Liu, T.-S. Chua, and B. Schiele. Meta-transfer learning for few-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[37] S. Thrun. Is learning the n-th thing any easier than learning the first? In *Advances in Neural Information Processing Systems*, 1996.

[38] Y. Tian, Y. Wang, D. Krishnan, J. B. Tenenbaum, and P. Isola. Rethinking few-shot image classification: a good embedding is all you need? *arXiv preprint arXiv:2003.11539*, 2020.

[39] P. E. Utgoff. Shift of bias for inductive concept learning. *Machine learning: An artificial intelligence approach*, 1986.

[40] R. Vilalta and Y. Drissi. A perspective view and survey of meta-learning. *Artificial Intelligence Review*, 2002.

[41] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra, et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, 2016.

[42] Y. Wang, W.-L. Chao, K. Q. Weinberger, and L. van der Maaten. Simpleshot: Revisiting nearest-neighbor classification for few-shot learning. *arXiv preprint arXiv:1911.04623*, 2019.

[43] S. W. Yoon, J. Seo, and J. Moon. Tapnet: Neural network augmented with task-adaptive projection for few-shot learning. In *International Conference on Machine Learning*, 2019.

[44] J. Zhang, C. Zhao, B. Ni, M. Xu, and X. Yang. Variational few-shot learning. In *IEEE International Conference on Computer Vision*, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] , in the supplemental material.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] , in Section 5.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] , in the supplemental material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] , in the supplemental material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]