

Regression Networks for Meta-Learning Few-Shot Classification

Arnout Devos

ARNOUT.DEVOS@EPFL.CH

Matthias Grossglauser

MATTHIAS.GROSSGLAUSER@EPFL.CH

École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Abstract

We propose *regression networks* for the problem of few-shot classification, where a classifier must generalize to new classes not seen in the training set, given only a small number of examples of each class. In high dimensional embedding spaces the direction of data generally contains richer information than magnitude. Next to this, state-of-the-art few-shot metric methods that compare distances with aggregated class representations, have shown superior performance. Combining these two insights, we propose to meta-learn classification of embedded points by regressing the closest approximation in every class subspace while using the regression error as a distance metric. Similarly to recent approaches for few-shot learning, regression networks reflect a simple inductive bias that is beneficial in this limited-data regime and they achieve excellent results, especially when more aggregate class representations can be formed with multiple shots. The code for this project is publicly available at <https://github.com/ArnoutDevos/RegressionNet>

1. Introduction

The ability to adapt quickly to new situations is a cornerstone of human intelligence. Artificial learning methods have been shown to be very effective for specific tasks, often surpassing human performance (Silver et al., 2016; Esteva et al., 2017). However, by relying on standard training paradigms for supervised learning or reinforcement learning, these artificial methods still require much training data and training time to adapt to a new task.

An area of machine learning that learns and adapts from a small amount of data is called *few-shot learning* (Fei-Fei et al., 2006). A *shot* corresponds to a single example, e.g., an image and its label. In few-shot learning the learning scope is expanded from the classic setting of a single task with many shots to a variety of tasks with a few shots each. Several machine learning approaches have been used for this, including metric-learning (Vinyals et al., 2016; Snell et al., 2017; Sung et al., 2018), meta-learning (Finn et al., 2017; Ravi and Larochelle, 2017), and generative models (Fei-Fei et al., 2006; Lake et al., 2015).

Chen et al. (2019) show that a metric-learning based method called *prototypical networks* (Snell et al., 2017), although simple in nature, achieves competitive performance with state-of-the-art meta-learning approaches such as MAML (Finn et al., 2017) and other metric-learning methods. Metric-learning methods approach the few-shot classification problem by "*learning to compare*". To learn high capacity nonlinear comparison models, most modern few-shot metric-learning methods use a neural embedding space to measure distance.

In image classification with neural embeddings, the embedding dimensions are usually high: e.g., 1600 for a *Conv-4* backbone (used later). In high dimensional image vector spaces "direction" of data generally contains richer information than "magnitude" (Zhe

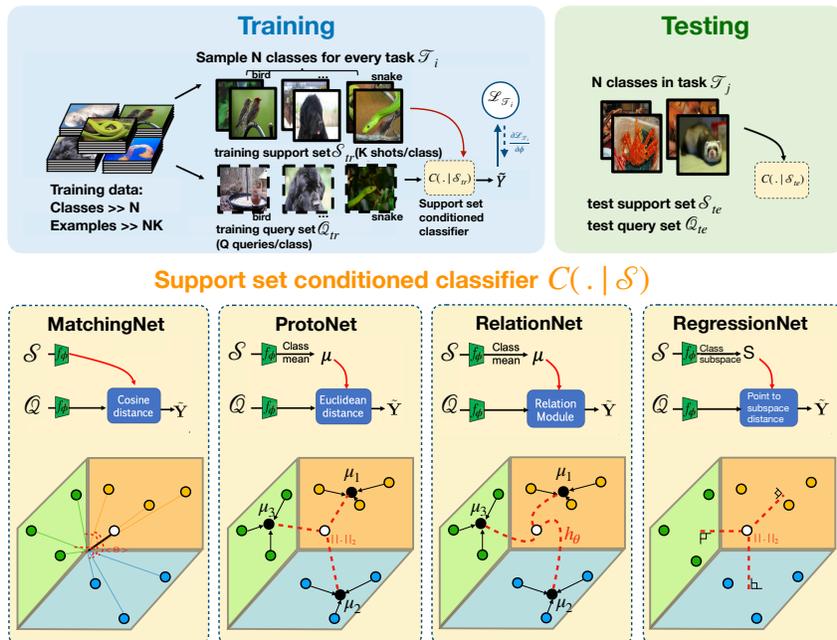


Figure 1: Few-shot learning process (top) and metric-learning based methods (bottom), with \mathcal{S} the support set (colored circles), \mathcal{Q} the query set (white circle), \tilde{Y} the output distribution over classes for the query points, f_ϕ a neural embedding function, and h_θ a neural network based distance function. Figure inspired by Chen et al. (2019).

et al., 2019). MatchingNet (Vinyals et al., 2016) leverages purely directional information whereas ProtoNet (Snell et al., 2017) and RelationNet (Sung et al., 2018) mostly improve on this by comparing with aggregated class representations.

We propose *regression networks* which combine the good properties of directional information in high-dimensional vector spaces with rich aggregate class information. The proposed method is based on the idea that there exists an embedding in which every point from the same class can be approximated by a linear combination of other points in that same class. In order to do this, we learn a nonlinear mapping of the input space to an embedding space by using a neural network and regress the best embedded approximation, for each example. Classification of an embedded query point is then performed by simply finding the nearest class subspace by comparing regression errors.

Subspaces have been used to model images for decades in computer vision and machine learning (Fitzgibbon and Zisserman, 2003; Naseem et al., 2010). For example, the linear regression classification (LRC) method (Naseem et al., 2010) relies on the fact that the set of all reflectance functions produced by Lambertian objects, which parts of natural images are composed of, lie near a low-dimensional vector subspace (Basri and Jacobs, 2003). More recently, Simon et al. (2018) have explored few-shot learning with affine subspaces. Unlike our approach with vector subspaces, affine subspaces cannot be constructed with 1-shot learning, a key few-shot learning problem.

Figure 1 shows an overview of the few-shot learning process and a comparison of the proposed method with comparable state-of-the-art metric-learning based approaches.

2. Regression Networks

2.1 Notation

We formulate the N -way K -shot classification problem in an episodic way. Every episode has a small support set of N classes with K labeled examples $\mathcal{S} = \{(\mathbf{x}_{11}, y_{11}), \dots, (\mathbf{x}_{NK}, y_{NK})\}$, and a query set of Q different examples $\mathcal{Q} = \{(\mathbf{x}_{1(K+1)}, y_{1(K+1)}), \dots, (\mathbf{x}_{N(K+Q)}, y_{N(K+Q)})\}$. Note that the query set contains labels only during training, and the goal is to predict the labels of the query set during testing. In \mathcal{S} and \mathcal{Q} each $\mathbf{x}_{ij} \in \mathbb{R}^D$ is the D -dimensional feature vector of an example and $y_{ij} \in \{1, \dots, N\}$ is the corresponding label.

2.2 Model

Regression networks perform classification by regressing the best approximation to an embedding point in each aggregated class representation and subsequently using the regression error as a measure of distance. For this we start by constructing a K -dimensional embedded subspace \mathbf{S}_n of each class n , given K shots per class, through an embedding function $f_\phi : \mathbb{R}^D \rightarrow \mathbb{R}^M$ with learnable parameters ϕ . With slight abuse of notation, every class is represented by its subspace matrix $\mathbf{S}_n \in \mathbb{R}^{M \times K}$, that contains the K vectors of the embedded class support points:

$$\mathbf{S}_n = [f_\phi(\mathbf{x}_{n1}) \ \dots \ f_\phi(\mathbf{x}_{nK})] \quad (1)$$

The *regression-error distance* $\tilde{d}(\mathbf{e}_i, \mathbf{S}_n)$ of a point $\mathbf{e}_i \in \mathbb{R}^M$ in the embedding space to a class subspace \mathbf{S}_n can be measured by regressing the closest point in the subspace in terms of a certain distance metric. The closest point can be constructed with a linear combination (represented by vector $\mathbf{a} \in \mathbb{R}^{K \times 1}$) of the embedded support examples spanning that space. By using a Euclidean distance metric, this can be formulated as a quadratic optimization problem in \mathbf{a} of the following form:

$$\tilde{d}(\mathbf{e}_i, \mathbf{S}_n) = \min_{\mathbf{a}} d(\mathbf{e}_i, \mathbf{S}_n \mathbf{a}) = \min_{\mathbf{a}} \|\mathbf{e}_i - \mathbf{S}_n \mathbf{a}\|_2 \quad (2)$$

The associated learning problem with this few-shot (K) high-dimensional (M) overdetermined system is *least-squares linear regression*. Given that $M \geq K$, this is a well conditioned system, and it admits a *differentiable* closed-form solution (Friedman et al., 2001):

$$\tilde{d}(\mathbf{e}_i, \mathbf{S}_n) = \left\| \mathbf{e}_i - \mathbf{S}_n (\mathbf{S}_n^T \mathbf{S}_n)^{-1} \mathbf{S}_n^T \mathbf{e}_i \right\|_2 \quad (3)$$

$$= \|\mathbf{e}_i - \mathbf{P}_n \mathbf{e}_i\|_2 \quad (4)$$

where it is important to note that the matrix to be inverted is of size $K \times K$ and the number of shots K is usually small. The transformation matrix \mathbf{P}_n projects the point \mathbf{e}_i orthogonally onto the subspace spanned by the columns of \mathbf{S}_n (Koç and Barkana, 2014). Note that this transformation matrix \mathbf{P}_n has to be computed only for every class, not every example, which speeds up practical computation. Because the embedding function f_ϕ can output linearly dependent embeddings for different support examples, a small term $\lambda_1 > 0$ is added to avoid singularity:

$$\tilde{d}(\mathbf{e}_i, \mathbf{S}_n) = \left\| \mathbf{e}_i - \mathbf{S}_n (\mathbf{S}_n^T \mathbf{S}_n + \lambda_1 \mathbf{I})^{-1} \mathbf{S}_n^T \mathbf{e}_i \right\|_2 \quad (5)$$

With this point-to-subspace distance function $\tilde{d} : \mathbb{R}^M \times \mathbb{R}^{M \times K} \rightarrow [0, +\infty)$, regression networks give a distribution over classes for a query point \mathbf{x} based on a softmax over distances to each of the class subspaces in the embedding space:

$$p_\phi(y = n \mid \mathbf{x}) = \frac{\exp\left(-\tilde{d}(f_\phi(\mathbf{x}), \mathbf{S}_n)\right)}{\sum_{n'} \exp\left(-\tilde{d}(f_\phi(\mathbf{x}), \mathbf{S}_{n'})\right)} \quad (6)$$

Meta-learning continues by minimizing the negative log-probability function $\mathcal{L}_{base}(\phi) = -\log p_\phi(y = n \mid \mathbf{x})$ of the true class n via stochastic gradient descent (SGD). Training episodes are formed by randomly sampling a subset of N classes from the training set. Then, a subset of K examples within each class is chosen as the support set \mathcal{S} , and a subset of Q examples within each class is chosen as the query set \mathcal{Q} .

2.3 Subspace Orthogonalization

The K -dimensional class vector subspaces live in a much larger M -dimensional space. Therefore, we can exploit this freedom during training by making subspaces as directionally different as possible. Concretely, we propose to add a pairwise subspace orthogonalization term to the loss function:

$$\mathcal{L}_T = \mathcal{L}_{base} + \lambda_2 \sum_{i \neq j}^N \frac{\|S_i^T S_j\|_F^2}{\|S_i\|_F^2 \|S_j\|_F^2} \quad (7)$$

where $\|\cdot\|_F$ is the Frobenius norm and λ_2 is a scaling hyperparameter. Section 3.2 studies the effect of this term. We have also experimented with principal angles between vector subspaces (Bjorck and Golub, 1973), because they only depend on the subspaces, not on the (non-unique) set of points that define the subspaces as in Equation (7). Their results were comparable with our current approach, but come at a higher computational cost with a singular value decomposition.

Algorithm 1 in Appendix A details the complete regression networks training procedure.

3. Experiments

In terms of few-shot learning evaluation, we focus on the natural image-based mini-ImageNet (Vinyals et al., 2016) dataset. To ensure a fair comparison with other methods, we perform experiments under the same conditions using the verified re-implementation (Chen et al., 2019) of MatchingNet, ProtoNet, RelationNet, MAML and extend it with R2D2 (Bertinetto et al., 2019). Compared to our direct approach, R2D2 is a meta-learning technique which leverages the closed-form solution of multinomial regression indirectly for classification (See Section 4). We decide to compare with these methods in particular, because they serve as the basis of many state-of-the-art few-shot classification algorithms (Oreshkin et al., 2018; Xing et al., 2019), and our method is easily interchanged with them. Experimental details can be found in Appendix B.

In this section, next to performance evaluation, we address the following research questions: (i) Can regression networks benefit from richer class representations and higher dimensions? (Section 3.1). (ii) How much effect does subspace orthogonalization have? (Section 3.2).

Method	Conv-4		ResNet-10	
	5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
MatchingNet	48.14±0.78	63.28±0.68	54.49±0.81	69.14±0.69
ProtoNet	44.42±0.84	65.15±0.67	51.98±0.84	73.77±0.64
RelationNet	49.31±0.85	65.33±0.70	52.19±0.83	69.97±0.68
MAML	46.47±0.82	62.71±0.71	54.69±0.89	66.62±0.83
R2D2	50.07±0.79	65.66±0.69	55.71±0.78	71.69±0.63
RegressionNet (ours)	47.02±0.77	67.09±0.69	55.44±0.86	76.29±0.59

Table 1: Average accuracies (%) of mini-ImageNet test tasks with 95% confidence intervals.

N -way K -shot	$\lambda_2 = 0$	$\lambda_2 > 0$
5-way 1-shot	54.83±0.83	55.44±0.86
5-way 5-shot	74.03±0.68	76.29±0.59

Table 2: Ablation study of effect of subspace orthogonalization stimulation (Equation (7)) using a ResNet-10 backbone on mini-ImageNet. 1-shot: $\lambda_2 = 10^{-3}$, 5-shot: $\lambda_2 = 10^{-2}$

3.1 Few-shot Image Classification: mini-ImageNet

Table 1 shows the results for 5-way classification for mini-ImageNet for a different number of shots and backbones.

First, because regression networks are expected to benefit more from better subspace representations when more support examples are available per class, we investigate the effect of the number of shots. As expected, when increasing the number of shots K per class from 1 to 5, the classification accuracies increase for all methods. In the 5-shot case, RegressionNet significantly outperforms all other methods, showing the benefit of using rich class representations.

Secondly, as the backbone gets deeper, regression networks and prototypical networks begin to perform significantly better than matching networks and relation networks with R2D2 following close. Although the performance difference is small for mini-Imagenet, given a relatively deep feature extractor ResNet-10, regression networks outperform all other meta-learning and metric-learning methods when enough shots are available.

3.2 Ablation study

In order to evaluate the effect of adding the subspace orthogonalization stimulating term to the loss function discussed in Section 2.3, we conduct an experiment without it. The results, comparing a ResNet-10 model trained with subspace orthogonalization and without, are shown in Table 2. All ablation experiments use ResNet-10 as a backbone.

Under all settings considered, subspace orthogonalization gives a classification accuracy improvement (up to 2%). Note that, even without subspace orthogonalization, our proposed method is still competitive with all other methods in Table 1.

4. Related work

In addition to the reproduced metric (meta-)learning based few-shot methods (Snell et al., 2017; Vinyals et al., 2016; Sung et al., 2018; Bertinetto et al., 2019), there is a large body of work on few-shot learning and metric (meta-)learning. We discuss work that is more closely related to regression networks in particular.

Our approach shows similarities to the linear regression classification (LRC) method (Naseem et al., 2010), where each class is represented by the vector subspace spanned by its examples. LRC was developed for face recognition, where only a few examples are available, however it relies on a linear embedding. Our approach also uses few examples, but it incorporates neural networks in order to nonlinearly embed examples and we couple this with episodic training to handle the meta-learning few-shot scenario.

Simon et al. (2018) have explored affine subspace representations for few-shot learning. In contrast to our closed-form linear regression approach, they make use of a truncated singular value decomposition (SVD) of the support example matrix. Affine subspaces cannot be constructed with 1-shot learning, a key few-shot learning problem. In contrast, our closed-form linear regression approach relies on vector subspaces, which can be constructed with 1-shot learning.

Bertinetto et al. (2019) propose to use regularized linear regression as a classifier on top of the embedding function. Doing so, they directly approach a classification problem with a regression method, but they show competitive results. To achieve this, they introduce learnable scalars that scale and shift the regression outputs for them to be used in the cross-entropy loss. Regression networks rely on the same closed-form solver for linear regression to compute the transformation matrices, but are inherently designed for classification problems because of their similarity to the LRC method (Naseem et al., 2010).

5. Conclusion

We have proposed regression networks for meta-learning few-shot classification. The method assumes that for any embedded point we can regress the closest approximation in every class representation and use the error as a distance measure. Classes are represented by their embedded vector subspaces, which are spanned by their examples. The approach produces better results than other state-of-the-art metric-learning based methods, when rich class representations can be formed with multiple shots. Stimulating subspace orthogonality consistently improves performance. A direction for future work is to study the effect of using a low-rank approximation of the class subspace. Overall, the simplicity and effectiveness of regression networks makes it a promising approach for metric-based few-shot classification.

Acknowledgements

We acknowledge fruitful discussions with Andreas Loukas, William Trouleau, Gergely Odor, Farnood Salehi, Negar Foroutan and anonymous reviewers. This project is supported by the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 754354.

References

- Ronen Basri and David W Jacobs. Lambertian reflectance and linear subspaces. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (2):218–233, 2003.
- Luca Bertinetto, Joo F. Henriques, Philip H. S. Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. In *ICLR 2019 : 7th International Conference on Learning Representations*, 2019. URL <https://academic.microsoft.com/paper/2964206659>.
- Ake Bjorck and Gene H Golub. Numerical methods for computing angles between linear subspaces. *Mathematics of computation*, 27(123):579–594, 1973.
- Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International Conference on Learning Representations*, 2019.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- Andre Esteva, Brett Kuprel, Roberto A Novoa, Justin Ko, Susan M Swetter, Helen M Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115, 2017.
- Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE transactions on pattern analysis and machine intelligence*, 28(4):594–611, 2006.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML’17 Proceedings of the 34th International Conference on Machine Learning - Volume 70*, pages 1126–1135, 2017. URL <https://academic.microsoft.com/paper/2604763608>.
- Andrew W Fitzgibbon and Andrew Zisserman. Joint manifold distance: a new approach to appearance based clustering. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 1, pages I–I. IEEE, 2003.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics New York, 2001.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Mehmet Koç and Atalay Barkana. Application of linear regression classification to low-dimensional datasets. *Neurocomputing*, 131:331–335, 2014.

- Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Imran Naseem, Roberto Togneri, and Mohammed Bennamoun. Linear regression for face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 32(11): 2106–2112, 2010.
- Boris Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. In *Advances in Neural Information Processing Systems*, pages 721–731, 2018.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan L Yuille. Few-shot image recognition by predicting parameters from activations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7229–7238, 2018.
- Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR 2017 : International Conference on Learning Representations 2017*, 2017. URL <https://academic.microsoft.com/paper/2753160622>.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- Christian Simon, Piotr Koniusz, and Mehrtash Harandi. Projective subspace networks for few-shot learning, 2018. URL <https://openreview.net/forum?id=rkzfuiA9F7>.
- Jake Snell, Kevin Swersky, and Richard S. Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017. URL <https://academic.microsoft.com/paper/2601450892>.
- Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1199–1208, 2018.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.
- Chen Xing, Negar Rostamzadeh, Boris N. Oreshkin, and Pedro O. Pinheiro. Adaptive cross-modal few-shot learning. *CoRR*, abs/1902.07104, 2019. URL <http://arxiv.org/abs/1902.07104>.
- Xuefei Zhe, Shifeng Chen, and Hong Yan. Directional statistics-based deep metric learning for image classification and retrieval. *Pattern Recognition*, 93:113–123, 2019.

Appendix A. Algorithm

Algorithm 1 Regression Networks

Require: Training set with task batches of N_T N -way K -shot episodes/tasks $\{\mathcal{T}_1, \dots, \mathcal{T}_{N_T}\}$, with every \mathcal{T}_i containing sets $\mathcal{S}_{\mathcal{T}_i} = \{(\mathbf{x}_{i11}, y_{i11}), \dots, (\mathbf{x}_{iNK}, y_{iNK})\}$ and $\mathcal{Q}_{\mathcal{T}_i} = \{(\mathbf{x}_{i1(K+1)}, y_{i1(K+1)}), \dots, (\mathbf{x}_{iN(K+Q)}, y_{iN(K+Q)})\}$. $\mathcal{S}_{\mathcal{T}_in}$ and $\mathcal{Q}_{\mathcal{T}_in}$ denote the class n subsets of support and query sets, respectively, of episode \mathcal{T}_i .

Require: $\alpha, \lambda_1, \lambda_2$: step size, conditioning and orthogonalization parameters

```

1: Randomly initialize  $\phi$ 
2: while not done do
3:   for each batch do ▷ Sample batch of tasks
4:     for  $i$  in  $\{1, \dots, N_T\}$  do ▷ Select task
5:       for  $n$  in  $\{1, \dots, N\}$  do ▷ Select class
6:          $\mathbf{S}_n \leftarrow f_\phi(\mathcal{S}_{\mathcal{T}_in})$  ▷ Embed class support set subspace  $\mathbf{S}_n \in \mathbb{R}^{M \times K}$ 
7:          $\mathbf{P}_n \leftarrow \mathbf{S}_n (\mathbf{S}_n^T \mathbf{S}_n + \lambda_1 \mathbf{I})^{-1} \mathbf{S}_n^T$  ▷ Compute transformation matrix
8:          $\mathcal{L}_{\mathcal{T}_i} \leftarrow 0$  ▷ Initialize episode loss
9:         for  $n$  in  $\{1, \dots, N\}$  do
10:          for  $(\mathbf{x}, y = n)$  in  $\mathcal{Q}_{\mathcal{T}_in}$  do
11:            
$$\mathcal{L}_{\mathcal{T}_i} \leftarrow \mathcal{L}_{\mathcal{T}_i} + \frac{1}{N_Q} \left[ \|f_\phi(\mathbf{x}) - \mathbf{P}_n f_\phi(\mathbf{x})\|_2 + \log \sum_{n'} \exp(-\|f_\phi(\mathbf{x}) - \mathbf{P}_{n'} f_\phi(\mathbf{x})\|_2) \right]$$


$$+ \lambda_2 \sum_{i \neq j}^N \frac{\|\mathbf{s}_i^T \mathbf{s}_j\|_F^2}{\|\mathbf{s}_i\|_F^2 \|\mathbf{s}_j\|_F^2}$$

12:           $\phi \leftarrow \phi - \alpha \nabla_\phi \sum_i \mathcal{L}_{\mathcal{T}_i}$  ▷ Update embedding parameters  $\phi$  with gradient descent

```

Appendix B. Experimental details

B.1 Experimental Setup and Datasets

The mini-Imagenet dataset proposed by (Vinyals et al., 2016) contains 100 classes, with 600 84×84 images per class sampled from the larger ImageNet dataset (Deng et al., 2009). Following Ravi and Larochelle (2017), 64 classes are isolated for the training set and, from the remaining classes, the validation and test sets of 16 and 20 classes, respectively, are constructed. We use exactly the same train/validation/test split of classes as the one suggested by Ravi and Larochelle (2017). We implement regression networks using the automatic-differentiation framework PyTorch (Paszke et al., 2017).

Many training optimizations exist, including using more classes in the training episodes than in the testing episodes (Snell et al., 2017) or pre-training the feature extractor with all training classes and a linear output layer (Qiao et al., 2018). However, we train all methods from scratch and construct our training and testing episodes to have the same number of classes N and shots K because we are interested in the relative performance of the methods. For RegressionNet, the conditioning parameter used to ensure a fully invertible matrix in Equation (5) is set to $\lambda_1 = 10^{-3}$. For 1-shot and 5-shot learning, $\lambda_2 = 10^{-3}$ and $\lambda_2 = 10^{-2}$ are used, respectively.

B.2 Architectures and training

Despite the modification of some implementation details of the methods with respect to the original papers, these settings ensure a fair comparison and Chen et al. (2019) report a maximal drop in classification performance of 2% with respect to the original reported performance of each method. For MAML, we reuse the results from Chen et al. (2019).

The *Conv-4* backbone is composed of four convolutional blocks with an input size of 84×84 as in Snell et al. (2017). Each block comprises a 64-filter 3×3 convolution with a padding of 1 and a stride of 1, a batch normalization layer, a ReLU nonlinearity and a 2×2 max-pooling layer.

The *ResNet-10* backbone has an input size of 224×224 and is a simplified version of ResNet-18 in He et al. (2016), by using only one residual building block in each layer.

All methods are trained with a random parameter initialization and use the Adam optimizer (Kingma and Ba, 2014) with an initial learning rate of 10^{-3} . During the training stage, data augmentation is done in the form of: random crop, color jitter, and left-right flip. For MatchingNet, a rather sophisticated long-short term memory (LSTM)-based full context embedding (FCE) classification layer without fine-tuning is used over the support set, and the cosine similarity metric is multiplied by a constant factor of 100. In RelationNet, the L2 norm is replaced with a softmax layer to help the training procedure (Chen et al., 2019). The relation module is composed of 2 convolutional blocks (same as in Conv-4), followed by two fully connected layers of size 8 and 1.

In all settings, we train for 60,000 episodes and use the held-out validation set to select the best model during training. To construct an N -way episode, we sample $N = 5$ classes from the training set of classes. From each sampled class, we sample K examples to construct the support set of an episode and $Q = 16$ examples for the query set. During testing, we average the results over 600 episodes of exactly the same form, but this time they are sampled from the test set classes.