

RelationNet2: Deep Comparison Network for Few-Shot Learning

Xueting Zhang
School of Informatics
University of Edinburgh
Edinburgh, UK
Xueting.Zhang@ed.ac.uk

Yuting Qiang
National Key Laboratory for Software Technology
Nanjing University
Nanjing, China
qiangyuting.new@gmail.com

Flood Sung
Independent Researcher
Beijing, China
floodsung@gmail.com

Yongxin Yang
School of Informatics
University of Edinburgh
Edinburgh, UK
yongxin.yang@ed.ac.uk

Timothy Hospedales
Samsung AI Research Centre School of Informatics
University of Edinburgh
Cambridge, UK
Edinburgh, UK
t.hospedales@samsung.com t.hospedales@ed.ac.uk

Abstract—Few-shot deep learning is a topical challenge area for scaling visual recognition to open ended growth of unseen new classes with limited labeled examples. A promising approach is based on metric learning, which trains a deep embedding to support image similarity matching. Our insight is that effective general purpose matching requires non-linear comparison of features at multiple abstraction levels. We thus propose a new deep comparison network comprised of embedding and relation modules that learn multiple non-linear distance metrics based on different levels of features simultaneously. Furthermore, to reduce over-fitting and enable the use of deeper embeddings, we represent images as distributions rather than vectors via learning parameterized Gaussian noise regularization. The resulting network achieves excellent performance on both miniImageNet and tieredImageNet.

I. INTRODUCTION

The ability to learn from one or few examples is an important property of human learning to function effectively in the real world. In contrast, our most successful deep learning-based approaches to recognition [1–3] treat each learning problem as tabula-rasa, limiting their application to open-ended learning with rare data and expensive annotation (e.g., endangered species and medical images).

These observations have motivated a resurgence of interest in FSL (few-shot learning) for visual recognition [4–7] and beyond. Contemporary deep networks overfit in the few-shot regime – even when exploiting fine-tuning [8], data augmentation [1], or regularization [9] techniques. In contrast, ‘Meta-learning’ techniques extract transferable task agnostic knowledge from historical tasks and benefit sparse data learning of specific new target tasks. These take several forms: Fast adaptation methods enable sparse-data adaptation without overfitting – via good initial conditions [5] or learned optimizers [10]. Weight synthesis approaches learn a meta-network that synthesizes recognition weights given a training set [11, 12]. Deep metric learning approaches support representation [13] and comparison [4, 6] of instances, allowing new categories to

be recognized with nearest-neighbour comparison. However, existing approaches have several drawbacks including inference complexity [14, 15], architectural complexity [16], the need to fine-tune on the target problem [5], or reliance on a simple linear comparison [4, 6, 15].

We build on deep metric learning methods due to their architectural simplicity and instantaneous training of new categories. These methods use auxiliary training tasks to learn a deep image-embedding such that the embedded data becomes linearly separable [4, 6, 13]. Thus the decision is non-linear in image-space, but linear in the embedding space. For learning the target task, images are simply memorized during few-shot training. But for testing the target task, query images are matched to training examples by deep embedding and similarity comparison function. Within this paradigm, the recent Relation Network [17] achieved excellent performance by learning a non-linear comparison function. Learning the embedding and non-linear comparison module jointly alleviates the reliance on the embedding’s ability to generate linearly separable features.

We extend this idea of jointly learning an embedding and a non-linear distance metric with the following further insights. First, we introduce the notion of multiple meta-learners operating at multiple abstraction levels. Concretely we train non-linear distance metrics corresponding to each embedding module in a feature hierarchy - thus covering features from simple textures to complex parts [18]. Secondly, prior studies only use a single linear [6] or non-linear comparison [17]. To provide the inductive bias that each layer of representation should be potentially discriminative for matching, and enable better gradient propagation [19] to each relation module, we deeply supervise [20] all the relation modules. Finally, to enable deeper embedding architectures to be used without overfitting, we design each embedding module to output a feature *distribution*, thus representing each image as a distribution rather than a vector. This can be seen as an end-to-end

learnable noise regularizer that performs data augmentation in semantic feature space rather than image space.

Overall RelationNet2 implements a Deep Comparison Network (DCN) that can be seen as jointly learning embedding and comparison as task agnostic meta knowledge [4, 6, 17, 21]. It makes full use of deep networks by making comparisons with the full feature hierarchy extracted by the embedding network, and learning Gaussian noise to improve generalization. The resulting framework maintains the architecture simplicity and efficiency of other methods in this line, while providing excellent performance on both *miniImageNet* and the more challenging *tieredImageNet* few shot learning benchmarks.

II. RELATED WORK

Contemporary approaches to deep-network few-shot learning have exploited the learning-to-learn paradigm [21]. Auxiliary tasks are used to meta-learn some task agnostic knowledge, before exploiting this to learn the target few-sample more effectively problem. The learning-to-learn idea has a long history [14, 22, 23], but contemporary approaches typically cluster into three categories: Fast adaptation, weight synthesis, and metric-learning approaches.

Fast Adaptation These approaches aim to meta-learn an optimisation process that enables base models to be fine-tuned quickly and robustly. So that a base model can be updated for sparse data target problems without extensive overfitting. Effective ideas include the simply meta-learning an effective initial condition [5, 24], and learning a recurrent neural network optimizer to replace the standard SGD learning approach [10]. Recent extensions also include learning per-parameter learning rates [25], and accelerating fine-tuning through solving some layers in closed form [26]. Nevertheless, these methods suffer from needing to be fine-tuned for the target problem, often generating costly higher-order gradients during meta-learning process [5], and failing to scale to deeper network architectures as shown in [12]. They also suffer from a fixed parametric architecture. For example, once you train MAML [5] for 5-way auxiliary classification problems, it is restricted to the same for target problems without being straightforwardly generalizable to a different cardinality of classification.

Classifier Synthesis Another line of work focuses on synthesising a classifier based on the provided few-shot training data [27]. An early method in this line learned a transferrable ‘LearnNet’ that generated convolutional weights for the base recognition network given a one-shot training example [11]. However, this was limited to binary classification. Conditional Neural Processes [28] exploited a similar idea, but in a Bayesian framework. SNAIL obtained excellent results by embedding the training set with temporal convolutions and attention [12]. The PPA model predicts classification parameters given neuron activations [7]. In this case, the global parameter prediction network is the task agnostic knowledge that is transferred from auxiliary categories. Compared to the fast adaptation approaches, these methods generally synthesize their classifier in a single pass, making them faster to train

on the target problem. However, learning to synthesize a full classifier does entail some complexity. This process can overfit and generalize poorly to novel target problem.

Deep Metric Learning These approaches aim to learn a deep embedding that extracts robust features, allowing them to be classified directly with nearest neighbour type strategies in the embedding space. The deep embedding forms the task agnostic knowledge transferred from auxiliary to target tasks. Early work simply used Siamese networks [13] to embed images, such that images of the same class are placed near each other. Matching networks [4] defined a differentiable nearest-neighbour loss based on cosine similarity between the support set and query embedding. Prototypical Networks [6] provide a simpler but more effective variant of this idea where the support set instances for one class are embedded as a single prototype. Their analysis showed that this leads to a linear classifier in the embedding space. RelationNet [17] extended this line of work to use a separate non-linear comparison module instead of relying entirely on the embedding networks to make the data linearly separable [4, 6, 13]. This division of labour between a deep embedding and a deep relation module improved performance in practice [17]. Our approach builds on this line of work in general and RelationNet in particular. RelationNet relied on the embedding networks to produce a *single* embedding for the relation module to compare. We argue that a general purpose comparison function should use any or all of the full feature hierarchy [18] to make matching decisions. For example, matching based on colors, textures, or parts – which may be represented at different layers in a embedding network. To this end, we modularise the embedding networks, and pair every embedding module with its own relation module.

Use of Feature Hierarchies The general strategy of simultaneously exploiting multiple layers of a feature hierarchy has been exploited in conventional many-shot classification network [19, 29], instance recognition [30], and semantic segmentation networks [31]. However, in the context of deep-metric learning, the conventional pipeline is to extract a complete feature [32, 33]. Importantly, in contrast to prior approaches single ‘short-cut’ connection of deeper features to a classifier [30, 31], we uniquely learn a hierarchy of relation modules: One non-linear comparison function for each block of the embedding modules. Our approach is also reminiscent of classic techniques such as spatial pyramids [34] (since each module in the hierarchy operates at different spatial resolutions) and multi-kernel learning [35] (since we learn multiple relation modules for each feature in the hierarchy). This can also be seen as the first multiple meta-learner approach for few shot learning problems.

Leaned Noise and Regularisation Many previous FSL models struggle with deeper backbones [5, 12]. For best performance, we would like to exploit a state-of-the-art embedding module architecture (we use SENet [3]), and also benefit from the array of comparison modules mentioned above. To enable DCN to benefit from deep backbones without

overfitting, we modify the embedding modules to output a feature distribution at each layer. Rather than generating deterministic features at a module output, we generate means and variances which are sampled in the forward pass, with back propagation relying on the reparameterization trick. Unlike density networks [36] where such distributions are only generated at the output layer, or VAEs [37] here they are generated only once by the generator, we generate such stochastic features at *each* embedding module’s output. This can be seen as an end-to-end learnable data augmentation strategy in semantic feature rather than image space. It is also complementary to standard L2/weight decay and image space augmentation techniques.

III. METHODOLOGY

A. Problem Definition

We consider a C -way K -shot classification problem for few shot learning. There are some labelled source tasks with sufficient data, denoted meta-train $\mathcal{D}_{m\text{-train}}$, and we ultimately want to solve a new set of target tasks denoted meta-test $\mathcal{D}_{m\text{-test}}$, for which the label space is disjoint. Within meta-train and meta-test, we denote each task as being composed of a support set of training examples and a query set of testing examples. The meta-test tasks are assumed to be few-shot, so $\mathcal{D}_{m\text{-test}}$ contains a support set with C categories and K examples each. We want to learn a model on meta-train that can generalize out of the box, without fine-tuning, to learning the new categories in meta-test.




Fig. 1: Few-shot learning: Problem Setup.

Episodic Training We adopt an episodic training paradigm for few-shot meta-learning. During meta-training, an episode is formed as follows: (i) Randomly select C classes from $\mathcal{D}_{m\text{-train}}$, (ii) Sample K images each class, which serve as *support set* $\mathcal{D}_{m\text{-train}}^S = \{(x_i, y_i)\}_{i=1}^m$, where $m = K * C$, (iii) For the same C classes, sample K' images each class serving as the *query set* $\mathcal{D}_{m\text{-train}}^Q = \{(\tilde{x}_j, \tilde{y}_j)\}_{j=1}^n$, where $n = K' * C$, $\mathcal{D}_{m\text{-train}}^S \cap \mathcal{D}_{m\text{-train}}^Q = \emptyset$. The support/query distinction mimics the $\mathcal{D}_{m\text{-test}}$ /real-time testing. Our few-shot DCN will be trained for instance comparison using episodes constructed in this manner.

B. Model

Overview RelationNet2’s Deep Comparison Network (DCN) is composed of two module types: *embedding* and

relation modules f_θ and g_ϕ , as shown in Fig. 2. The detailed architecture will be given in Section III-C. A pair of images x_i and x_j in the support and query set are fed to embedding modules respectively. Then the multi-level embedding modules output stochastic features to the corresponding multi-level relation modules, and learn the relation score and weights for different relation modules. Finally, the DCN learns weighted non-linear metric of few shot learning tasks.

Distribution Embedding Modules Conventionally, an embedding module (e.g., a ResNet or SENet block) outputs deterministic features. As a regularisation strategy, we treat each feature output as a random variable drawn from a parameterized Gaussian distribution, for which the embedding module outputs the mean and variance. This design is illustrated in Fig. 2. Each v th-level embedding module predicts a feature mean $f_{\theta,\mu}^v$ and a feature variance $f_{\theta,\sigma}^v$. To generate a module’s output f_θ^v , we use the reparameterization trick to draw one (or more) Gaussian random samples

$$f_\theta^v = f_{\theta,\mu}^v + \varepsilon \odot f_{\theta,\sigma}^v, \quad (1)$$

where ε is a standard Gaussian $\mathcal{N}(0, 1)$ random samples, and \odot denotes element-wise product.

Metric Hierarchy The v th-level of embedding modules produce query and support image feature maps, which are concatenated as $[f_\theta^v(x_i), f_\theta^v(x_j)]$, and then fed into the corresponding v th-level relation module for comparison.

For a pair x_i and x_j at level $v - 1$, the relation module outputs a similarity feature map g_ϕ^{v-1} . The v th-level relation module takes both the v th-level embedding output for query and support, and also the $(v - 1)$ th-level relation module similarity feature map as input:

$$g_\phi^v = g([f_\theta^v(x_i), f_\theta^v(x_j), g_\phi^{v-1}]). \quad (2)$$

The first relation module is special as it does not have a predecessor to input, and we cannot use zero-padding because 0 has a specific meaning in our context. Thus we set $g_\phi^1 = g([f_\theta^1(x_i), f_\theta^1(x_j)])$.

Simultaneously, after an average pooling and fully connected layer denoted $q(\cdot)$, each relation module also outputs a real-valued scalar representing similarity/relation score $r_{i,j}^v$ of two images estimated at the feature level v ,

$$r_{ij}^v = q(g_\phi^v). \quad (3)$$

K-Shot For K -shot with $K > 1$, the embedding module outputs the average pooling of features, along the sample axis, of all samples from the same class to produce *one* feature map. Thus, the number of outputs for the v -level relation module is C , regardless of the value of K .

Objective Function There are 2 steps to train the DCN network. We first train the embedding network, then fix the embedding network parameters and train the relation network (run the whole DCN consisting of embedding and relation modules, but only update the relation modules).




Fig. 2: RelationNet2’s DCN architecture. There are 4 embedding modules f_θ for each embedding branch, and a set of 4 corresponding relation modules g_ϕ . Support set and query set share the same embedding network. Each embedding module outputs a feature distribution $\mathcal{N}(f_{\theta,\mu}(x), f_{\theta,\sigma}(x))$, we then randomly sample a feature $f_\theta(x)$ as the input of corresponding relation module and next embedding module.

We first train the embedding network θ as a conventional multi-class classifier for the data in $\mathcal{D}_{m-train}$ using cross entropy loss ℓ^{CE} . To leverage our distribution-embedding, we add a feature variance regularizer:

$$\theta \leftarrow \operatorname{argmin}_{\theta} \ell^{CE}(\theta) - \lambda \frac{1}{m} \sum_{i=1}^m \sigma_i, \quad (4)$$

where σ_i is the predicted standard deviation of each instance and m is their total number, and λ is the hyperparameter to finetune the influence of the regularizer (here is 0.01). This ensures that feature distributions are learned, and we do not collapse to standard (zero-variance) vector embedding (our mean σ is about 0.5). This pipeline can be seen as a learnable data augmentation strategy at each level of the feature hierarchy for relation modules. Learning with these augmented features improves generalization. After embedding training, the parameters θ of embedding modules are fixed.

We next train the column of relation modules ϕ on $\mathcal{D}_{m-train}$ with an episodic strategy [4] using cross entropy loss ℓ^{CE} at each module (Fig. 2). To weight the V relation modules, we assign a learnable attention weight $w_{c,j}^v$ to the calculated relation similarity score $r_{c,j}^v$ of each module.

$$\phi \leftarrow \operatorname{argmin}_{\phi} \sum_{c=1}^C \sum_{j=1}^n \sum_{v=1}^V \ell^{CE}(w_{c,j}^v r_{c,j}^v, \mathbf{1}(y_c = y_j); \phi), \quad (5)$$

where $j = 1 \dots n$ refers to query samples and c refers to a batch of K support examples of class y_c in a C -way- K -shot problem. $r_{c,j}$ are the relation scores between query image j and the class y_c support images.

Additionally, $w_{c,j}^v = \alpha^v(g_{c,j}^v)$ is a sigmoid-activated fully connected layer that computes a scalar attention weight given relation feature map $g_{c,j}^v$, and the weights of α^v are included in ϕ .

Testing Strategy To evaluate our learned model on C -way- K -shot learning, we calculate the final relation score $r_{c,j}$ of one query image x_j to the images of each support class c :

$$r_{c,j} = \sum_{v=1}^V w_j^v r_{c,j}^v \quad (6)$$

where $r_{c,j}^v$ is the relation score between image j and the support images of class c at module v .

Finally, the class with the highest relation score r_c is the final predicted classification. We evaluate the approach by the resulting classification accuracy.

C. Network Architecture

RelationNet2’s DCN architecture (Fig. 2) uses 4 embedding modules, each paired with a relation module. We explain our method with SENet for concreteness, but it can be instantiated with any backbone.

Embedding Subnetwork As shown in Fig. 2, first we use a 7×7 convolution followed by a 3×3 max-pooling, which is a common size reduction as [3]. Then, we have 4 embedding modules each composed of a number of SENet blocks. Finally, an avg-pooling and a fully-connected layer are used to produce C logit values, corresponding to C classes in $\mathcal{D}_{m-train}$. More specifically, 4 embedding modules followed the 4 SENet basic blocks composition [3, 4, 6, 3], respectively. In original SENet paper [3], they use SE-ResNet-50, but here we use smaller backbones as SE-ResNet-34, where $(3+4+6+3) * 2+2 = 34$. Otherwise, we follow the other setting in [3], e.g., reduction ratio $r = 16$ as suggested.

Distribution Embedding Conventually, an embedding module outputs deterministic features. As explained in Section III-B, each DCN embedding module’s output is split into two parts: the mean feature $f_{\theta,\mu}$ sized $[b, c, h, w]$ ([batch_size, channel, height, width]), and standard deviation (std) $f_{\theta,\sigma}$ sized $[b, 1, h, w]$.

We assume that every channel shares the same standard deviation (std). This means, in addition to the penultimate-to-output layer (now it is penultimate-to-mean layer), we have a new penultimate-to-std layer (with its own parameters). The motivation behind sharing stds across channels is to reduce the number of parameters in the newly introduced layer. We also control the amount of noise added by applying Sigmoid activation to constrain the std to the range $[0, 1]$. We sample one feature vector per image in a single forward pass, but multiple samples are drawn considering the whole batch.

Relation Subnetwork As illustrated in Fig. 2, the relation column consists of 4 serial modules, each of which has 2 SENet blocks, with a pooling and a fully-connected layer to produce the relation score. Thus the relation modules is designed as $[2, 2, 2, 2]$, where the SENet block architecture is the same as the one used in embedding module.

IV. EXPERIMENTS

We evaluate RelationNet2’s DCN architecture on few-shot classification with *miniImageNet* and *tieredImageNet* datasets. PyTorch code to reproduce results is available at <https://github.com/zhangxueting/DCN>.

Baselines We compare several state-of-the-art baselines for few-shot learning including Matching Nets [4], Meta Nets [16], Meta LSTM [10], MAML [5], Baseline++ [38], Prototypical Nets [6], Graph Neural Nets [39], Meta-SSL [40], Relation Net [17], Meta-SGD [25], TPN [41], CAVIA [42], DynamicFSL [27], SNAIL [12], AdaResNet [43], TADAM [44], MTL [45], TapNet [46], MetaOpt Net [15], PPA [7], LEO [47].

Data Augmentation We follow the standard data augmentation [2, 3, 38, 48] with random-size cropping and random horizontal flipping. Input images are normalized through mean channel subtraction.

A. *miniImageNet*

Dataset *miniImageNet* has 60,000 images in consist of 100 ImageNet classes, each with 600 images [4]. Following the split in [10], the dataset is divided into a 64-class training set, 16-class validation set and a 20-class testing set.

Settings We evaluate both *5-way-1-shot* and *5-way-5-shot*, where each episode contains 5 query images for each sampled class. There are $5 \times 5 + 1 \times 5 = 30$ images per training episode/mini-batch for *5-way-1-shot* experiments, and $5 \times 5 + 5 \times 5 = 50$ images for *5-way-5-shot* experiments. When it comes to *5-shot*, we calculate the class-wise average feature across the support set. Thus we get $5 \times 5 \times 5 \times 1 = 125$ feature pairs as input for the relation module. For embedding and relation module training, optimization uses SGD with momentum 0.9. The initial learning rate is 0.1, decreased by a factor of 5 every 60 epochs, and the training epoch is 200. All models are trained from scratch, using the robust RELU weight initialization [49]. We follow [38] in using 224×224 pixels crops for evaluation on ResNet and SENet, and [10] in using 84×84 images for the smaller Conv-4 backbone.

Results Following the setting of [6], when evaluating testing performance, we batch 15 query images per class in a

testing episode and the accuracy is calculated by averaging over 600 randomly generated testing tasks (for both 1-shot and 5-shot scenarios). In Tab. I, DCN achieves excellent performance with different embedding backbones. Specifically, the accuracy of 5-way *miniImageNet* with SENet is 63.19% and 76.58% for 1-shot and 5-shot respectively. We note that MetaOptNet [15] uses significantly more advanced regularizers than standard among the competitors (which corresponds to about 2% performance according to [15]), also requires an order of magnitude higher dimensionality of embeddings $[64, 160, 320, 640]$ than the other competitors $[64, 96, 128, 256]$. Overall DCN’s 1-shot recognition performance is state-of-the-art among methods that do not require optimisation at meta-test time (unlike, e.g., MAML [5] and MetaOptNet [15]). It is noteworthy that achieving good performance with deeper backbones is not trivially automatic as Dynamic FSL, for example, fails to improve from Conv-4 to ResNet embedding. DCN’s learned noise regularizer helps it to exploit a powerful SENet backbone without overfitting. Direct comparison among models is complicated by the diversity of embedding networks used in different studies, so we show the results of DCN with each commonly used backbone in Tab.I, e.g. Conv-4 and ResNet-12. We can see that DCN performs favorably across a range of architectures.

Cross-way Testing Results Standard procedure in few-shot evaluation is to train models for the desired number of categories to discriminate at testing time. However, unlike alternatives such as MAML [5], our method is not required to match label cardinality between training and testing. We therefore evaluate 5-way trained model on 20-way testing in Tab. II. It shows that our model outperforms the alternatives clearly despite DCN being trained for 5-way, and the others specifically for 20-way, indicating another important aspect of DCN’s flexibility and general applicability.

B. *tieredImageNet*

Dataset *tieredImageNet* is a larger few-shot recognition benchmark containing 608 classes (779,165 images), in which training/validation/testing categories are organized so as to ensure a larger semantic gap than those in *miniImageNet*, thus providing a more rigorous test of generalization. This is achieved by dividing according to 34-higher-level nodes in the ImageNet hierarchy [40], grouped into 20 for training (351 classes), 6 for validation (97 classes) and 8 for testing (160 classes), respectively.

Settings Similar to the setting of *miniImageNet*, we use 5 query images per training episode. Due to the larger data size, we train embedding modules with a larger batch size 512, initial learning rate 0.3 and 100 training epochs. Other settings remain the same as *miniImageNet*.

Results Following the former experiments, we batch 15 query images per class in each testing episode and the accuracy is calculated by averaging over 600 randomly generated testing tasks. From Tab. III, DCN achieves the state-of-the-art performance on the *5-way-1-shot* and *5-shot* tasks with comfortable margins. Again, this is state-of-the-art performance

Model	Embedding	<i>miniImageNet</i> 5-way Acc.	
		1-shot	5-shot
MATCHING NETS [4]	Conv-4	43.56 ± 0.84%	55.31 ± 0.73%
META LSTM [10]	Conv-4	43.44 ± 0.77%	60.60 ± 0.71%
MAML ^O [5]	Conv-4	48.70 ± 1.84%	63.11 ± 0.92%
BASELINE++ [38]	Conv-4	48.24 ± 0.75%	66.43 ± 0.63%
META NETS [16]	Conv-5	49.21 ± 0.96%	-
PROTO NET [6]	Conv-4	49.42 ± 0.78%	68.20 ± 0.66%
GNN [39]	Conv-4	50.33 ± 0.36%	66.41 ± 0.63%
META SSL [40]	Conv-4	50.41 ± 0.31%	64.39 ± 0.24%
RELATION NET [17]	Conv-4	50.44 ± 0.82%	65.32 ± 0.70%
META SGD ^O [25]	Conv-4	50.47 ± 1.87%	64.03 ± 0.94%
TPN [41]	Conv-4	52.78 ± 0.27%	66.59 ± 0.28%
CAVIA [42]	Conv-4	51.82 ± 0.65%	65.85 ± 0.55%
DYNAMIC FSL [†] [27]	Conv-4	56.20 ± 0.86%	72.81 ± 0.62%
RELATIONNET2 (DCN)	Conv-4	53.48 ± 0.78%	67.63 ± 0.59%
BASELINE++ [38]	ResNet-18	51.87 ± 0.77%	75.68 ± 0.63%
RELATIONNET [38]	ResNet-18	52.48 ± 0.86%	69.83 ± 0.68%
PROTONET [38]	ResNet-18	54.16 ± 0.82%	73.68 ± 0.65%
SNAIL [50]	ResNet-12	55.71 ± 0.99%	68.88 ± 0.92%
DYNAMIC FSL [27]	ResNet-12	55.45 ± 0.89%	70.13 ± 0.68%
ADARESNET [43]	ResNet-12	57.10 ± 0.70%	70.04 ± 0.63%
TADAM [44]	ResNet-12	58.50 ± 0.30%	76.70 ± 0.30%
MTL [45]	ResNet-12*	61.20 ± 1.80%	75.50 ± 0.80%
TAP NET [46]	ResNet-12	61.65 ± 0.15%	76.36 ± 0.10%
META OPTNET ^O [15]	ResNet-12*	64.09 ± 0.62%	80.00 ± 0.45%
RELATIONNET2 (DCN)	ResNet-12	63.92 ± 0.98%	77.15 ± 0.59%
PPA [7]	WRN-28-10	59.60 ± 0.41%	73.74 ± 0.19%
LEO ^O [47]	WRN-28-10	61.78 ± 0.05%	77.59 ± 0.12%
MAML	SENet	55.99 ± 0.99%	-
RELATIONNET	SENet	57.39 ± 0.86%	-
PROTO NET	SENet	51.60 ± 0.85%	-
RELATIONNET2 (DCN)	SENet	63.19 ± 0.87%	76.58 ± 0.66%

TABLE I: Few-shot classification results on *miniImageNet*. All accuracies are averaged over 600 test episodes and are reported with 95% confidence intervals. Best-performing method is bold, along with any others whose confidence intervals overlap. From top to bottom: Simple conv block embeddings to other deep embeddings (ResNet, WRN, SENet). ‘-’: not reported. †: use two-step optimization with added attention. ^O: requires gradient-based optimisation at meta-test time. *: Use a wider ResNet than standard and higher dimensional embedding.

for methods that do not require optimisation at meta-testing. We also note that Meta-SSL [40] and TPN [41] are semi-supervised methods that use more information than ours, and have additional requirements such as access to the test set for transduction.

C. Further Analysis

1) *Application to Other Metric Learners*: Our main insight is the value of feature comparison at multiple abstraction levels in metric learning, as well as that of learned noise regularizers for deep networks in the few-shot regime. We now confirm that these ideas can be applied to other base metric learners. Tab VI shows the 5-way-1-shot *miniImageNet* results for both RelationNet [17] and ProtoNet [6] base learners controlling for these features. We can see that both architectures benefit from deep comparisons and regularizers. However, the benefit is greater for RelationNet, which we attribute to the learnable non-linear relation modules. These can learn a different com-

Model	Embedding	<i>miniImageNet</i> 20-way Acc.	
		1-shot	5-shot
MATCHING NETS [25]	Conv-4	17.31 ± 0.22%	22.69 ± 0.86%
META LSTM ^O [25]	Conv-4	16.70 ± 0.23%	26.06 ± 0.25%
MAML ^O [25]	Conv-4	16.49 ± 0.58%	19.29 ± 0.29%
META SGD ^O [25]	Conv-4	17.56 ± 0.64%	28.92 ± 0.35%
RELATIONNET2 (DCN)	Conv-4	27.56 ± 0.24%	39.56 ± 0.81%
RELATIONNET2 (DCN)	ResNet-12	31.65 ± 0.34%	50.25 ± 0.46%
RELATIONNET2 (DCN)	SENet	32.90 ± 0.39%	51.37 ± 0.39%

TABLE II: 20-way classification accuracy on *miniImageNet*. DCN is trained on 5-way with different embeddings and transferred to 20-way. Meta LSTM, MAML, and Meta SGD results are from [25].

Model	Embedding	<i>tieredImageNet</i> 5-way Acc.	
		1-shot	5-shot
REPTILE [41]	Conv-4	48.97%	66.47%
MAML [41]	Conv-4	51.67%	70.30%
META SSL [†] [40]	Conv-4	52.39 ± 0.44%	70.25 ± 0.31%
PROTO NET [41]	Conv-4	53.31%	72.69%
RELATION NET [41]	Conv-4	54.48%	71.31%
TPN [†] [41]	Conv-4	59.91%	73.30%
TAP NET [46]	ResNet-12	63.08 ± 0.15%	80.26 ± 0.12%
META OPTNET ^O [15]	ResNet-12*	65.81 ± 0.74%	81.75 ± 0.53%
RELATIONNET2 (DCN)	Conv-4	60.58 ± 0.72%	72.42 ± 0.69%
RELATIONNET2 (DCN)	ResNet-12	68.58 ± 0.63%	80.65 ± 0.91%
RELATIONNET2 (DCN)	SENet	68.83 ± 0.94%	79.62 ± 0.77%

TABLE III: Few-shot classification results on *tieredImageNet*. All accuracies are averaged over 600 test episodes and reported with 95% confidence intervals. For each task, the best-performing method is bold. †: Make use of additional unlabeled data for semi-supervised learning or transductive inference. ^O: requires gradient-based optimisation at meta-test time. *: Uses a wider ResNet than standard size and higher dimensional embedding.

parison function at each abstraction level, but are also more complex so benefit more from the additional regularisation.

2) *Ablation Study*: We further investigate the detailed design parameters of our method with a series of ablation studies reported in Tab. IV. The conclusions are as follows:

Deep Supervision: The DCN-No Deep Sup. Result shows that deep supervision is important to gain full benefit from a column of relation modules. **Module Weighting**: Learning attention weights per module helps somewhat compared to manually tuned module weights. More importantly, it eliminates the need for hand-tuning model weights. **Multiple Non-linear Metrics**: Tab IV also shows the testing accuracy with each DCN relation module output score r_v in isolation (DCN- r_v). Each module performs competitively, but their combination clearly leads to the best overall performance, supporting our argument that multiple levels of the feature hierarchy should be used to make general purpose matching decisions. Multiple meta learner design is a creative contribution of our work.

3) *Architecture*: Our DCN benefits from deeper embedding architectures (Tab. I). It improves when going from simple convolutional blocks (used by early studies [5, 6, 17]), to ResNet [2] and SENet [3]. For fair comparison, when fixing a common ResNet-12, DCN outperforms the others that do




Fig. 3: Illustration of query-support score distribution and the link to ImageNet hierarchy. Colors indicate query images of a $(query, support1, support2)$ class triple matching the specified ImageNet distance relationship $[D(q, s1), D(q, s2)]$.

Model	<i>miniImageNet</i> 5-way-1-shot Acc.
DCN Full model	63.19 \pm 0.87%
DCN-No module weight	62.88 \pm 0.83%
DCN-No deep sup.	58.02 \pm 0.80%
DCN- r_1	52.25 \pm 0.80%
DCN- r_2	58.07 \pm 0.80%
DCN- r_3	60.69 \pm 0.81%
DCN- r_4	58.31 \pm 0.79%

TABLE IV: Ablation study using 5-way-1-shot classification on *miniImageNet* evaluating the impact of regularization techniques and multiple relation modules.

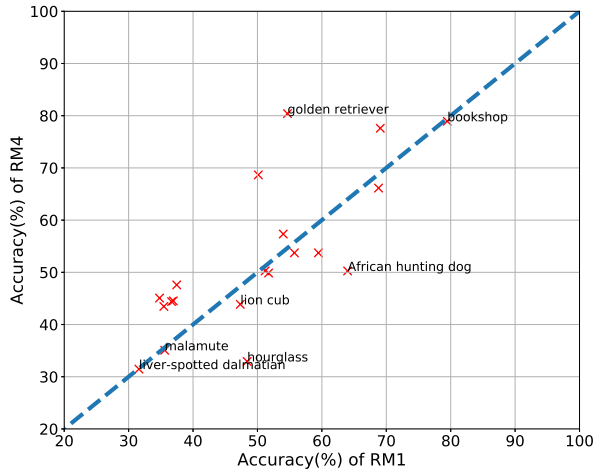


Fig. 4: Category-wise accuracy of RM1 vs RM4. Different relation modules are better at detecting different categories.

not require meta-test optimization. Moreover, when fixing a common SENet, competitors RelationNet/ProtoNet/MAML are improved, but still surpassed by DCN.

4) *Relation Module Analysis*: A key contribution in DCN is to perform metric learning at multiple abstraction levels

Module	RM1	RM2	RM3	RM4
RM1	-	-	-	-
RM2	0.75	-	-	-
RM3	0.55	0.73	-	-
RM4	0.34	0.45	0.61	-

TABLE V: Spearman rank-order correlation coefficient between different relation modules: Modules make diverse predictions.

Model	Noise Reg.?	Deep Comparisons?	Acc.
PROTONET [6]	X	X - 1 module	51.04 \pm 0.77%
PROTONET	✓	X - 1 module	51.60 \pm 0.85%
PROTONET	X	✓ - 4 modules	53.62 \pm 0.82%
PROTONET	✓	✓ - 4 modules	54.78 \pm 0.88%
RELATIONNET [17]	X	X - 1 module	52.48 \pm 0.86%
RELATIONNET	✓	X - 1 module	57.39 \pm 0.86%
RELATIONNET2 (DCN)	X	✓ - 4 modules	60.57 \pm 0.86%
RELATIONNET2 (DCN)	✓	✓ - 4 modules	63.19 \pm 0.87%

TABLE VI: Multiple deep comparisons and distribution embedding of features benefit both RelationNet (learnable relation modules) and ProtoNet (fixed linear modules) few-shot architectures. Accuracy is calculated on 5-way-1-shot classification of *miniImageNet*.

simultaneously via a series of paired relation and embedding modules. Relation modules are analyzed to provide insight into the complementarity.

Score-Distance Correlation We first check how the relation module (RM) scores relate to distances in the ImageNet hierarchy. Using *miniImageNet* data, we search for $(support1, support2, query)$ category tuples where the distance $D(query, support1)$ and $D(query, support2)$ match a certain number of links, and then plot instances from these tuples query categories against the relative relation module scores $RM(q, s1)$, $RM(q, s2)$. Fig. 3 presents scatter plots for the four relation modules where points are images and colors indicate category tuples with specified distance from the two support classes. We can see that: (1) The scores generally match ImageNet distances: The most/least similar categories (red/magenta) are usually closer to the top right/bottom left of the plot; while query categories closer to one support class are in the opposite corners (blue/yellow-green). (2) Generally, higher numbered relation modules are more discriminative, separating classes with larger differences in relation score.

Score Correlation We next investigated if relation module predictions are diverse or redundant. We analyzed the correlation in their predictions by randomly picking 10,000 image pairs from *miniImageNet* and computing the Spearman rank-order correlation coefficient [51] between each pair of relation module’s scores. The results in Tab. V, show that: (1) Many correlations are relatively low (down to 0.34), indicating that they are making diverse, non-redundant predictions; and (2) Adjacent RMs have higher correlation than non-adjacent RMs, indicating that prediction diversity is related to RM position

in the feature hierarchy.

Prediction Success by Module We know that RM predictions do not necessarily agree. But to find out if they are complementary, we made a scatter plot of the per-class accuracy of RM-1 vs. RM-4 in Fig. 4. We can see that many categories lie on the diagonal, indicating that RM-1 and-4 get them right equally often. However, there are some categories *below* the diagonal, indicating that RM-1 gets them right more often than RM-4. Examples include both stereotyped and fine-grained categories such as ‘hourglass’ and ‘African hunting dog’. These below diagonal elements confirm the value of using deeper features in metric learning.

V. CONCLUSION

We proposed RelationNet2, a general purpose matching framework for few-shot learning. It implements a Deep Comparison Network architecture that performs effective few-shot learning via learning multiple non-linear comparisons corresponding to multiple levels of feature extraction, while resisting overfitting through stochastic regularisation. The resulting method achieves state-of-the-art results on *miniImageNet* and the more ambitious *tieredImageNet*, while retaining architectural simplicity, and fast training and testing processes.

Acknowledgements This work was supported by EPSRC grant EP/R026173/1.

REFERENCES

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *NIPS*, 2012.
- [2] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016.
- [3] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *CVPR*, 2018.
- [4] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, “Matching networks for one shot learning,” in *NIPS*, 2016.
- [5] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *ICML*, 2017.
- [6] J. Snell, K. Swersky, and R. S. Zemel, “Prototypical networks for few-shot learning,” in *NIPS*, 2017.
- [7] S. Qiao, C. Liu, W. Shen, and A. L. Yuille, “Few-shot image recognition by predicting parameters from activations,” in *CVPR*, 2018.
- [8] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, “How transferable are features in deep neural networks?” in *NIPS*, 2014.
- [9] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *JMLR*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [10] S. Ravi and H. Larochelle, “Optimization as a model for few-shot learning,” in *ICLR*, 2017.
- [11] L. Bertinetto, J. F. Henriques, J. Valmadre, P. H. S. Torr, and A. Vedaldi, “Learning feed-forward one-shot learners,” in *NIPS*, 2016.
- [12] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel, “A simple neural attentive meta-learner,” in *ICLR*, 2018.
- [13] G. Koch, R. Zemel, and R. Salakhutdinov, “Siamese neural networks for one-shot image recognition,” in *ICML Deep Learning Workshop*, 2015.
- [14] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum, “Human-level concept learning through probabilistic program induction,” *Science*, vol. 350, no. 6266, pp. 1332–1338, 2015.
- [15] K. Lee, S. Maji, A. Ravichandran, and S. Soatto, “Meta-learning with differentiable convex optimization,” in *CVPR*, 2019.
- [16] T. Munkhdalai and H. Yu, “Meta networks,” in *ICML*, 2017.
- [17] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales, “Learning to compare: Relation network for few-shot learning,” in *CVPR*, 2018.
- [18] M. D. Zeiler and R. Fergus, “Visualizing and understanding convolutional networks,” in *ECCV*, 2014.
- [19] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *CVPR*, 2017.
- [20] C.-Y. Lee, S. Xie, P. W. Gallagher, Z. Zhang, and Z. Tu, “Deeply-supervised nets,” in *AISTATS*, 2015.
- [21] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, “Meta-learning in neural networks: A survey,” *arXiv preprint arXiv:2004.05439*, 2020.
- [22] S. Thrun, “Is learning the n-th thing any easier than learning the first?” in *NIPS*, 1996.
- [23] L. Fei-Fei, R. Fergus, and P. Perona, “One-shot learning of object categories,” *TPAMI*, vol. 28, no. 4, pp. 594–611, 2006.
- [24] A. Nichol, J. Achiam, and J. Schulman, “On first-order meta-learning algorithms,” *arXiv preprint arXiv:1803.02999*, 2018.
- [25] Z. Li, F. Zhou, F. Chen, and H. Li, “Meta-sgd: Learning to learn quickly for few shot learning,” *arXiv preprint arXiv:1707.09835*, 2017.
- [26] L. Bertinetto, J. F. Henriques, P. H. Torr, and A. Vedaldi, “Meta-learning with differentiable closed-form solvers,” in *ICLR*, 2019.
- [27] S. Gidaris and N. Komodakis, “Dynamic few-shot visual learning without forgetting,” in *CVPR*, 2018.
- [28] M. Garnelo, D. Rosenbaum, C. J. Maddison, T. Ramalho, D. Saxton, M. Shanahan, Y. W. Teh, D. J. Rezende, and S. Eslami, “Conditional neural processes,” *arXiv preprint arXiv:1807.01613*, 2018.
- [29] R. K. Srivastava, K. Greff, and J. Schmidhuber, “Highway networks,” in *ICML Deep Learning Workshop*, 2015.
- [30] X. Chang, T. M. Hospedales, and T. Xiang, “Multi-level factorisation net for person re-identification,” in *CVPR*, 2018.
- [31] B. Hariharan, P. Arbelaez, R. Girshick, and J. Malik, “Hypercolumns for object segmentation and fine-grained localization,” in *CVPR*, 2015.
- [32] W. Ge, W. Huang, D. Dong, and M. R. Scott, “Deep metric learning with hierarchical triplet loss,” in *ECCV*, 2018.
- [33] J. Hu, J. Lu, and Y.-P. Tan, “Discriminative deep metric learning for face verification in the wild,” in *CVPR*, 2014.
- [34] S. Lazebnik, C. Schmid, and J. Ponce, “Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories,” in *CVPR*, 2006.
- [35] A. Vedaldi, V. Gulshan, M. Varma, and A. Zisserman, “Multiple kernels for object detection,” in *ICCV*, 2009.
- [36] C. M. Bishop, “Mixture density networks,” Aston University, Tech. Rep., 1994.
- [37] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” in *ICLR*, 2014.
- [38] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. Wang, and J.-B. Huang, “A closer look at few-shot classification,” in *ICLR*, 2019.
- [39] V. Garcia and J. Bruna, “Few-shot learning with graph neural networks,” in *ICLR*, 2018.
- [40] M. Ren, E. Triantafyllou, S. Ravi, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel, “Meta-learning for semi-supervised few-shot classification,” in *ICLR*, 2018.
- [41] Y. Liu, J. Lee, M. Park, S. Kim, and Y. Yang, “Transductive propagation network for few-shot learning,” in *ICLR*, 2019.
- [42] L. M. Zintgraf, K. Shiari, V. Kurin, K. Hofmann, and S. Whiteson, “Fast context adaptation via meta-learning,” in *ICML*, 2019.
- [43] T. Munkhdalai, X. Yuan, S. Mehri, and A. Trischler, “Rapid adaptation with conditionally shifted neurons,” in *ICML*, 2018.
- [44] B. Oreshkin, P. R. López, and A. Lacoste, “Tadam: Task dependent adaptive metric for improved few-shot learning,” in *NIPS*, 2018.
- [45] Q. Sun, Y. Liu, T.-S. Chua, and B. Schiele, “Meta-transfer learning for few-shot learning,” in *CVPR*, 2019.
- [46] S. W. Yoon, J. Seo, and J. Moon, “Tapnet: Neural network augmented with task-adaptive projection for few-shot learning,” in *ICML*, 2019.
- [47] A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, and R. Hadsell, “Meta-learning with latent embedding optimization,” in *ICLR*, 2019.
- [48] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *CVPR*, 2015.
- [49] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *ICCV*, 2015.
- [50] A. Santoro, D. Raposo, D. G. Barrett, M. Malinowski, R. Pascanu, P. Battaglia, and T. Lillicrap, “A simple neural network module for relational reasoning,” in *NIPS*, 2017.
- [51] C. Spearman, “The proof and measurement of association between two things,” *The American Journal of Psychology*, vol. 15, no. 1, pp. 72–101, 1904.