Visual Question Answering as a Meta Learning Task

Damien Teney   Anton van den Hengel

Australian Centre for Visual Technologies

The University of Adelaide

{damien.teney,anton.vandenhengel}@adelaide.edu.au

Abstract

The predominant approach to Visual Question Answering (VQA) demands that the model represents within its weights all of the information required to answer any question about any image. Learning this information from any real training set seems unlikely, and representing it in a reasonable number of weights doubly so. We propose instead to approach VQA as a meta learning task, thus separating the question answering method from the information required. At test time, the method is provided with a support set of example questions/answers, over which it reasons to resolve the given question. The support set is not fixed and can be extended without retraining, thereby expanding the capabilities of the model. To exploit this dynamically provided information, we adapt a state-of-the-art VQA model with two techniques from the recent meta learning literature, namely prototypical networks and meta networks. Experiments demonstrate the capability of the system to learn to produce completely novel answers (i.e. never seen during training) from examples provided at test time. In comparison to the existing state of the art, the proposed method produces qualitatively distinct results with higher recall of rare answers, and a better sample efficiency that allows training with little initial data. More importantly, it represents an important step towards vision-and-language methods that can learn and reason on-the-fly.

1. Introduction

The task of Visual Question Answering (VQA) demands that an agent correctly answer a previously unseen question about a previously unseen image. The fact that neither the question nor the image is specified until test time means that the agent must embody most of the achievements of Computer Vision and Natural Language Processing, and many of those of Artificial Intelligence.

VQA is typically framed in a purely supervised learning setting. A large training set of example questions, images, and their correct answers is used to train a method to map
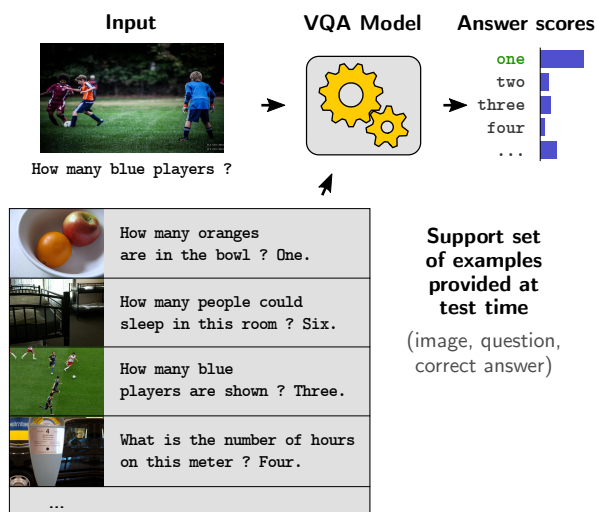

Figure 1. This paper considers Visual Question Answering in a meta learning setting. The model is initially trained on a small set of questions/answers, and is provided with a, possibly large, additional *support set* of examples at test time. The model must *learn to learn*, or to exploit the additional data on-the-fly, without the need for retraining the model. Notably, performance improves as additional and more relevant examples are included.

a question and image to scores over a predetermined, fixed vocabulary of possible answers using the maximum likelihood [39]. This approach has inherent scalability issues, as it attempts to represent all world knowledge within the finite set of parameters of a model such as deep neural network. Consequently, a trained VQA system can only be expected to produce correct answers to questions from a very similar distribution to those in the training set. Extending the model knowledge or expanding its domain coverage is only possible by retraining it from scratch, which is computationally costly, at best. This approach is thus fundamentally incapable of fulfilling the ultimate promise of VQA, which is answering general questions about general images.

As a solution to these issues we propose a meta-learning approach to the problem. The meta learning approach implies that the model *learns to learn*, *i.e.* it learns to use a

set of examples provided at test time to answer the given question (Fig. 1). Those examples are questions and images, each with their correct answer, such as might form part of the training set in a traditional setting. They are referred to here as the *support set*. Importantly, the support set is not fixed. Note also that the support set may be large, and that the majority of its elements may have no relevance to the current question. It is provided to the model at test time, and can be expanded with additional examples to increase the capabilities of the model. The model we propose 'learns to learn' in that it is able to identify and exploit the relevant examples within a potentially large support set dynamically, at test time. Providing the model with more information thus does not require retraining, and the ability to exploit such a support set greatly improves the practicality and scalability of the system. Indeed, it is ultimately desirable for a practical VQA system to be adaptable to new domains and to continuously improve as more data becomes available. That vision is a long term objective and this work takes only a small step in that direction.

Our primary technical contribution is to adapt a state-of-the-art VQA model[34] to the meta learning scenario. Our resulting model is a deep neural network that uses sets of dynamic parameters – also known as fast weights – determined at test time depending on the provided support set. The dynamic parameters allow to modify adaptively the computations performed by the network and adapt its behaviour depending on the support set. We perform a detailed study to evaluate the effectiveness of those techniques under various regimes of training and support set sizes. Those experiments are based on the VQA v2 benchmark, for which we propose data splits appropriate to study a meta learning setting.

A completely new capability demonstrated by the resulting system is to learn to produce completely novel answers (*i.e.* answers not seen during training). Those new answers are only demonstrated by instances of the support set provided at test time. In addition to these new capabilities, the system exhibits a qualitatively distinct behaviour to existing VQA systems in its improved handling of rare answers. Since datasets for VQA exhibit a heavy class imbalance, with a small number of answers being much more frequent than most others, models optimized for current benchmarks are prone to fall back on frequent "safe" answers. In contrast, the proposed model is inherently less likely to fall victim to dataset biases, and exhibits a higher recall over rare answers. The proposed model does *not* surpass existing methods on the common aggregate accuracy metric, as is to be expected given that it does not overfit to the dataset bias, but it nonetheless exhibits desirable traits overall.

The contributions of this paper are summarized as follows.

1. We re-frame VQA as a meta learning task, in which the model is provided a test time with a support set of super-

vised examples (questions and images with their correct answers).

2. We describe a neural network architecture and training procedure able to leverage the meta learning scenario. The model is based on a state-of-the-art VQA system and takes inspiration in techniques from the recent meta learning literature, namely prototypical networks [33] and meta networks [24].

3. We provide an experimental evaluation of the proposed model in different regimes of training and support set sizes and across variations in design choices.

4. Our results demonstrate the unique capability of the model to produce novel answers, *i.e.* answers never seen during training, by learning from support instances, an improved recall of rare answers, and a better sample efficiency than existing models.

2. Related Work

Visual question answering Visual question answering has gathered significant interest from the computer vision community [6], as it constitutes a practical setting to evaluate deep visual understanding. In addition to visual parsing, VQA requires the comprehension of a text question, and combined reasoning over vision and language, sometimes on the basis of external or common-sense knowledge. See [39] for a recent survey of methods and datasets.

VQA is always approached in a supervised setting, using large datasets [6, 15, 22, 44] of human-proposed questions with their correct answers to train a machine learning model. The *VQA-real* and *VQA v2* datasets [6, 15] have served as popular benchmarks by which to evaluate and compare methods. Despite the large scale of those datasets, *e.g.* more than 650,000 questions in *VQA v2*, several limitations have been recognized. These relate to the dataset bias (*i.e.* the non-uniform, long-tailed distribution of answers) and the question-conditioned bias (making answers easy to guess given a question without the image). For example, the answer *Yes* is particularly prominent in [6] compared to *no*, and questions starting with *How many* can be answered correctly with the answer *two* more than 30% of the time [15]. These issues plague development in the field by encouraging methods which fare well on common questions and concepts, rather than on rare answers or more complicated questions. The aggregate accuracy metric used to compare methods is thus a poor indication of method capabilities for visual understanding. Improvements to datasets have been introduced [1, 15, 43], including the *VQA v2*, but they only partially solve the evaluation problems. An increased interest has appeared in the handling of rare words and answers [29, 35]. The model proposed in this paper is inherently less prone to incorporate dataset biases than existing methods, and shows superior performance for handling rare answers. It accomplishes this by keeping a memory

made up of explicit representations of training and support instances.

VQA with additional data In the classical supervised setting, a fixed set of questions and answers is used to train a model once and for all. With few exceptions, the performance of such a model is fixed as it cannot use additional information at test time. Among those exceptions, [40, 38] use an external knowledge base to gather non-visual information related to the input question. In [35], the authors use visual information from web searches in the form of exemplar images of question words, and better handle rare and novel words appearing in questions as a result. In [34], the same authors use similar images from web searches to obtain visual representations of candidate answers.

Those methods use ad-hoc engineered techniques to incorporate external knowledge in the VQA model. In comparison, this paper presents a much more general approach. We expands the model knowledge with data provided in the form of additional supervised examples (questions and images with their correct answer). A demonstration of the broader generality of our framework over the works above is its ability to produce novel answers, *i.e.* never observed during initial training and learned only from test-time examples.

Recent works on text-based question answering have investigated the retrieval of external information with reinforcement learning [26, 25, 8]. Those works are tangentially related and complementary to the approach explored in this paper.

Meta learning and few shot learning The term *meta learning* broadly refers to methods that *learn to learn*, *i.e.* that train models to make better use of training data. It applies to approaches including the learning of gradient descent-like algorithms such as [5, 13, 17, 30] for faster training or fine-tuning of neural networks, and the learning of models that can be directly fed training examples at test time [7, 33, 36]. The method we propose falls into the latter category. Most works on meta learning are motivated by the challenge of one-shot and few-shot visual recognition, where the task is to classify an image into categories defined by a few examples each. Our meta learning setting for VQA bears many similarities. VQA is treated as a classification task, and we are provided, at test time, with examples that illustrate the possible answers – possibly a small number per answer. Most existing methods are, however, not directly applicable to our setting, due to the large number of classes (*i.e.* possible answers), the heavy class imbalance, and the need to integrate into an architecture suitable to VQA. For example, recent works such as [36] propose efficient training procedures that are only suitable for a small number of classes.

Our model uses a set of memories within a neural network to store the activations computed over the support set. Similarly, Kaiser *et al.* [19] store past activations to remember "rare events", which was notably evaluated on machine translation. Our model also uses network layers parametrized by dynamic weights, also known as fast weights. Those are determined at test time depending on the actual input to the network. Dynamic parameters have a long history in neural networks [32] and have been used previously for few-shot recognition [7] and for VQA [27]. One of the memories within our network stores the gradient of the loss with respect to static weights of the network, which is similar to the Meta Networks model proposed by Munkhdalai *et al.* [24]. Finally, our output stage produces scores over possible answers by similarity to prototypes representing the output classes (answers). This follows a similar idea to the Prototypical Networks [33].

Continuum learning An important outcome of framing VQA in a meta learning setting is to develop models capable of improving as more data becomes available. This touches the fields of incremental [12, 31] and continuum learning [2, 23, 42]. Those works focus on the fine-tuning of a network with new training data, output classes and/or tasks. In comparison, our model does not modify itself over time and cannot experience negative domain shift or catastrophic forgetting, which are a central concern of continuum learning [21]. Our approach is rather to use such additional data on-the-fly, at test time, *i.e.* without an iterative retraining. An important motivation for our framework is its potential to apply to support data of a different nature than question/answer examples. We consider this to be an important direction for future work. This would allow to leverage general, non VQA-specific data, *e.g.* from knowledge bases or web searches.

3. VQA in a Meta Learning Setting

The traditional approach to VQA is in a supervised setting described as follows. A model is trained to map an input question Q and image I to scores over candidate answers [39]. The model is trained to maximize the likelihood of correct answers over a training set $\mathcal{T}$ of triplets $(Q, I, \hat{s})$, where $\hat{s} \in [0,1]^A$ represents the vector of ground truth scores of the predefined set of A possible answers. At test time, the model is evaluated on another triplet $(Q', I', \hat{s}')$ from an evaluation or test set $\mathcal{E}$. The model predicts scores s' over the set of candidate answers, which can be compared to the ground truth $\hat{s}'$ for evaluation purposes.

We extend the formulation above to a meta learning setting by introducing an additional *support set* $\mathcal{S}$ of similar triplets $(Q'', I'', \hat{s}'')$. These are provided to the model at test time. At a minimum, we define the support set to include the training examples themselves, *i.e.* $\mathcal{S} = \mathcal{T}$, but more interestingly, the support set can include novel examples $\mathcal{S}'$ provided at test time. They constitute additional data to learn from, such that $\mathcal{S} = \mathcal{T} \cup \mathcal{S}'$. The triplets $(Q, I, \hat{s})$ in the
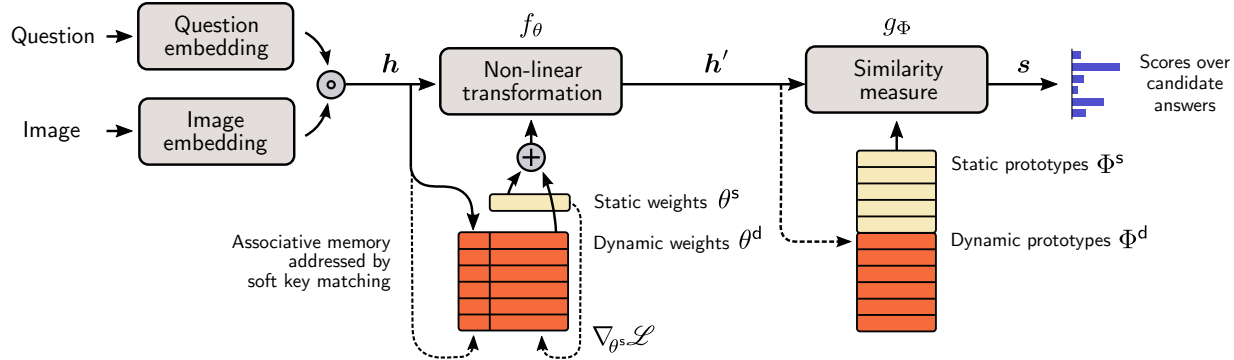
Figure 2. Overview of the proposed model. We obtain an embedding the input question and image following [34] and our contributions concern the mapping of this embedding to scores over a set of candidate answers. First, a **non-linear transformation** (implemented as a gated hyperbolic tangent layer) is parametrized by static and dynamic weights. Static ones are learned like traditional weights by gradient descent, while dynamic ones are determined based on the actual input and a memory of candidate dynamic weights filled by processing the support set. Second, a **similarity measure** compares the resulting feature vector to a set of prototypes, each representing a specific candidate answer. Static prototypes are learned like traditional weights, while dynamic prototypes are determined by processing the support set. Dashed lines indicate data flow during the processing of the support set. See Section 4 for details.

support set can also include novel answers, never seen in the training set. In that case, the ground truth score vectors $\hat{bs}$ of the other elements in the support are simply padded with zeros to match the larger size A' of the extended set of answers.

The following sections describe a deep neural network that can take advantage of the support set at test time. To leverage the information contained in the support set, the model must learn to utilize these examples on-the-fly at test time, without retraining of the whole model.

4. Proposed Model

The proposed model (Fig. 3) is a deep neural network that extends the state-of-art VQA system of Teney *et al.* [34]. Their system implements the joint embedding approach common to most modern VQA models [39, 41, 18, 20], which is followed by a multi-label classifier over candidate answers. Conceptually, we separate the architecture into (1) the embedding part that encodes the input question and image, and (2) the classifier part that handles the reasoning and actual question answering[1]. The contributions of this paper address only the second part. Our contributions are orthogonal to developments on the embedding part, which could also benefit *e.g.* from advanced attention mechanisms or other computer vision techniques [3, 37, 39]. We follow the implementation of [34] for the embedding part. For concreteness, let us mention that the question embedding uses *GloVe* word vectors [28] and a Recurrent Gated Unit (GRU [10]). The image embedding uses features from

---

[1]The separation of the network into an embedding and a classifier parts is conceptual. The division is arbitrarily placed after the fusion of the question and image embeddings. Computational requirements aside, the concept of dynamic parameters is in principle applicable to earlier layers as in [7].

a CNN (Convolutional Neural Network) with bottom-up attention [3] and question-guided attention over those features. See [34] for details.

For the remainder of this paper, we abstract the embedding to modules that produce respectively the question and image vectors $\boldsymbol{q}$ and $\boldsymbol{v} \in \mathbb{R}^D$. They are combined with a Hadamard (element-wise) product into $\boldsymbol{h} = \boldsymbol{q} \circ \boldsymbol{v}$, which forms the input to the classifier on which we now focus on. The role of the classifier is to map $\boldsymbol{h}$ to a vector of scores $\boldsymbol{s} \in [0, 1]^A$ over the candidate answers. We propose a definition of the classifier that generalizes the implementation of traditional models such as [34]. The input to the classifier $\boldsymbol{h} \in \mathbb{R}^D$ is first passed through a non-linear transformation $f_{\boldsymbol{\theta}} : \mathbb{R}^D \to \mathbb{R}^D$, then through a mapping to scores over the set of candidate answers $g_{\Phi} : \mathbb{R}^D \to [0, 1]^A$. This produces a vector of predicted scores $\boldsymbol{s} = f_{\boldsymbol{\theta}}(g_{\Phi}(\boldsymbol{h}))$. In traditional models, the two functions correspond to a stack of non-linear layers for $f_{\boldsymbol{\theta}}$, and a linear layer followed by a softmax or sigmoid for g_{Φ}. We now show how to extend $f_{\boldsymbol{\theta}}$ and g_{Φ} to take advantage of the meta learning setting.

4.1. Non-linear Transformation $f_{\boldsymbol{\theta}}(\cdot)$

The role of the non-linear transformation $f_{\boldsymbol{\theta}}(\boldsymbol{h})$ is to map the embedding of the question/image $\boldsymbol{h}$ to a representation suitable for the following (typically linear) classifier. This transformation can be implemented in a neural network with any type of non-linear layers. Our contributions are agnostic to this implementation choice. We follow [34] and use a gated hyperbolic tangent layer [11], defined as

$$f_{\boldsymbol{\theta}}(\boldsymbol{h}) = \sigma(W\boldsymbol{h} + \boldsymbol{b}) \circ \tanh(W'\boldsymbol{h} + \boldsymbol{b}') \quad (1)$$

where σ is the logistic activation function, $W, W' \in \mathbb{R}^{D \times D}$ are learned weights, $\boldsymbol{b}, \boldsymbol{b}' \in \mathbb{R}^D$ are learned biases, and $\circ$ is the Hadamard (element-wise) product. For notation pur-

poses, we define the parameters $\boldsymbol{\theta}$ as the concatenation of the vectorized weights and biases, *i.e.* $\boldsymbol{\theta} = [W_{:}; W'_{:}; \boldsymbol{b}; \boldsymbol{b}']$. This vector thus contains all of the weights and biases used by the non-linear transformation. A traditional model would learn the weights $\boldsymbol{\theta}$ by backpropagation and gradient descent on the training set, and they would be held static during test time. We propose instead to adaptively adjust the weights at test time, depending on the input $\boldsymbol{h}$ and the available support set. Concretely, we use a combination of static parameters $\boldsymbol{\theta}^{\mathsf{s}}$ learned in the traditional manner, and dynamic ones $\boldsymbol{\theta}^{\mathsf{d}}$ determined at test time. They are combined as $\boldsymbol{\theta} = \boldsymbol{\theta}^{\mathsf{s}} + \boldsymbol{w}\boldsymbol{\theta}^{\mathsf{d}}$, with $\boldsymbol{w}'' \in \mathbb{R}^D$ a vector of learned weights. The dynamic weights can therefore be seen as an adjustment made to the static ones depending on the input $\boldsymbol{h}$.

A set of candidate dynamic weights are maintained in an associative memory $\mathcal{M}$. This memory is a large set (as large as the support set, see Section 4.2) of key/value pairs $\mathcal{M} = \{(\tilde{\boldsymbol{h}}_i, \tilde{\boldsymbol{\theta}}_i^{\mathsf{d}})\}_{i \in 1 \dots |\mathcal{S}|}$. The interpretation for $\tilde{\boldsymbol{\theta}}_i^{\mathsf{d}}$ is of dynamic weights suited to an input similar to $\tilde{\boldsymbol{h}}_i$. Therefore, at test time, we retrieve appropriate dynamic weights $\boldsymbol{\theta}^{\mathsf{d}}$ by soft key matching:

$$\boldsymbol{\theta}^{\mathsf{d}} = \sum_i \tilde{\boldsymbol{\theta}}_i^{\mathsf{d}} \operatorname*{softmax}_i \big(d_{\mathsf{cos}}(\boldsymbol{h}, \tilde{\boldsymbol{h}}_i) \big) \qquad (2)$$

where $d_{\mathsf{cos}}(\cdot, \cdot)$ is the cosine similarity function. We therefore retrieve a weighted sum, in which the similarity of $\boldsymbol{h}$ with the memory keys $\tilde{\boldsymbol{h}}_i$ serves to weight the memory values $\tilde{\boldsymbol{\theta}}_i^{\mathsf{d}}$. In practice and for computational reasons, the softmax function cuts off after the top k largest values, with k in the order of a thousand elements (see Section 5). We detail in Section 4.2 how the memory is filled by processing the support set. Note that the above formulation can be made equivalent to the original model in [34] by using only static weights ($\boldsymbol{\theta} = \boldsymbol{\theta}^{\mathsf{s}}$). This serves as a baseline in our experiments (see Section 5).

4.1.1 Mapping to Candidate Answers $g_\Phi(\cdot)$

The function $g_\Phi(\boldsymbol{h})$ maps the output of the non-linear transformation to a vector of scores $s \in [0,1]^A$ over the set of candidate answers. It is traditionally implemented as a simple affine or linear transformation (*i.e.* a matrix multiplication). We generalize the definition of $g_\Phi(\boldsymbol{h})$ by interpreting it as a similarity measure between its input $\boldsymbol{h}$ and prototypes $\Phi = \{\phi_i^a\}_{i,a}$ representing the possible answers. In traditional models, each prototype corresponds to one row of the weight matrix. Our general formulation allows one or several prototypes per possible answer a as $\{\phi_i^a\}_{i=1}^{N^a}$. Intuitively, the prototypes represent the typical expected feature vector when a is a correct answer. The score for a is therefore obtained as the similarity between the provided $\boldsymbol{h}'$ and the corresponding prototypes of a. When multiple

prototypes are available, the similarities are averaged. Concretely, we define

$$g_\Phi^a(\boldsymbol{h}') = \sigma\Big(\frac{1}{N^a} \sum_{i=1}^{N^a} d(\boldsymbol{h}', \boldsymbol{\phi}_i^a) + b'' \Big) \qquad (3)$$

where $d(\cdot, \cdot)$ is a similarity measure, σ is a sigmoid (logistic) activation function to map the similarities to $[0,1]$, and b'' is a learned bias term. The traditional models that use a matrix multiplication [18, 34, 35] correspond to $g_\Phi(\cdot)$ that uses a dot product as the similarity function. In comparison, our definition generalizes to multiple prototypes per answer and to different similarity measures. Our experiments evaluate the dot product and the weighted L-p norm of vector differences:

$$d_{\mathsf{dot}}(\boldsymbol{h}, \boldsymbol{\theta}) = \boldsymbol{h}^\mathsf{T} \boldsymbol{\theta} \qquad (4)$$
$$d_{\mathsf{L1}}(\boldsymbol{h}, \boldsymbol{\theta}) = \boldsymbol{w}'''^\mathsf{T} |\boldsymbol{h} - \boldsymbol{\theta}| \qquad (5)$$
$$d_{\mathsf{L2}}(\boldsymbol{h}, \boldsymbol{\theta}) = \boldsymbol{w}'''^\mathsf{T} (\boldsymbol{h} - \boldsymbol{\theta})^2 \qquad (6)$$

where $\boldsymbol{w}''' \in \mathbb{R}^D$ is a vector of learned weights applied coordinate-wise.

Our model uses two sets of prototypes, the static Φ^{s} and the dynamic Φ^{d}. The static ones are learned during training as traditional weights by backpropagation and gradient descent, and held fixed at test time. The dynamic ones are determined at test time by processing the provided support set (see Section 4.2). Thereafter, all prototypes $\Phi = \Phi^{\mathsf{s}} \cup \Phi^{\mathsf{d}}$ are used indistinctively. Note that our formulation of $g_\Phi(\cdot)$ can be made equivalent to the original model of [34] by using only static prototypes ($\Phi = \Phi^{\mathsf{d}}$) and the dot-product similarity measure $d_{\mathsf{dot}}(\cdot, \cdot)$. This will serve as a baseline in our experiments (Section 5).

Finally, the output of the network is attached to a cross-entropy loss $\mathscr{L}(\boldsymbol{s}, \hat{\boldsymbol{s}})$ between the predicted and ground truth for training the model end-to-end [34].

4.2. Processing of Support Set

Both functions $f_{\boldsymbol{\theta}}(\cdot)$ and $g_\Phi(\cdot)$ defined above use dynamic parameters that are dependent on the support set. Our model processes the entire support set in a forward and backward pass through the network as described below. This step is to be carried out once at test time, prior to making predictions on any instance of the test set. At training time, it is repeated before every epoch to account for the evolving static parameters of the network as training progresses (see Algorithm 1).

We pass all elements of the support set $\mathcal{S}$ through the network in mini-batches for both a forward and backward pass. The evaluation of $f_{\boldsymbol{\theta}}(\cdot)$ and $g_\Phi(\cdot)$ use *only static weights and prototypes*, *i.e.* $\boldsymbol{\theta} = \boldsymbol{\theta}^{\mathsf{s}}$ and $\phi = \phi^{\mathsf{s}}$. To fill the memory $\mathcal{M}$, we collect, for every element of the support set, its feature vector $\boldsymbol{h}$ and the gradient $\nabla_{\boldsymbol{\theta}^{\mathsf{s}}} \mathscr{L}$ of the final loss relative to the static weights $\boldsymbol{\theta}$. This effectively captures the adjustments that would be made by a gradient descent algo-

Algorithm 1: Evaluation of one test or training instance.

Input:

Support set $\mathscr{S} = \{(\mathsf{Q}_i, \mathsf{l}_i, \hat{\boldsymbol{s}}_i)\}_i$

Instance to evaluate $(\mathsf{Q}, \mathsf{l}, \hat{\boldsymbol{s}})$ from training or test set

Output: Predicted scores $\boldsymbol{s}$ over candidate answers

$\Phi^{\mathsf{d}} \leftarrow \emptyset$ // *Initialize dynamic prototypes*

$\mathcal{M} \leftarrow \emptyset$ // *Initialize memory of dynamic weights*

for *each element i in support set $\mathscr{S}$* **do**

    if *training, with probability p* **then**

        | **continue** // *Drop random support elements*

    end

    Forward and backward propagation of $(\mathsf{Q}_i, \mathsf{l}_i, \hat{\boldsymbol{s}}_i)$
        using only static weights θ^{s} and prototypes Φ^{s}

    Collect $\boldsymbol{h}_i$

        $\boldsymbol{h}'_i = f_\theta(\boldsymbol{h}_i)$

        $\boldsymbol{s}_i = g_\Phi(\boldsymbol{h}'_i)$

        $\nabla_{\theta^{\mathsf{s}}}\mathscr{L}(\boldsymbol{s}_i, \hat{\boldsymbol{s}}_i)$

    // *Store (key, value) in memory of dynamic weights*

    $\mathcal{M} \leftarrow \mathcal{M} \ \cup \ (\boldsymbol{h}, \nabla_{\theta^{\mathsf{s}}}\mathscr{L}(\boldsymbol{s}_i, \hat{\boldsymbol{s}}_i))$

end

$\phi^a = \frac{1}{N^a} \sum_{i:\hat{s}^a_i=1}^{N^a} \boldsymbol{h}'_i \quad \forall a$ // *One average per answer*

$\Phi^{\mathsf{d}} = \cup_a \phi^a$ // *Store dynamic prototypes*

Forward prop. of (Q, l) with static and dynamic param.

if *test time* **then**

    Return predicted scores $\boldsymbol{s} = g_\Phi\big(f_\theta(\boldsymbol{h})\big)$

else if *training time* **then**

    Backpropagation and gradient descent

    Update static parameters $\theta^{\mathsf{s}}, \phi^{\mathsf{s}}, b'', w'', w'''$, and
    those of the question and image embeddings

end

---

rithm to those weights for that particular example. The pair $(\boldsymbol{h}, \nabla_{\theta^{\mathsf{s}}}\mathscr{L})$ is added to the memory $\mathcal{M}$, which thus holds $|\mathcal{S}|$ elements at the end of the process.

To determine the set of dynamic prototypes ϕ^{d}, we collect the feature vectors $\boldsymbol{h}' = f_{\boldsymbol{\theta}}(\boldsymbol{h})$ over all instances of the support set. We then compute their average over instances having the same correct answer. Concretely, the dynamic prototype for answer a is obtained as $\phi^a = \frac{1}{N^a} \sum_{i:\hat{s}^a_i=1}^{N^a} \boldsymbol{h}'_i$.

During training, we must balance the need for data to train the static parameters of the network, and the need for an "example" support set, such that the network can learn to use novel data. If the network is provided with a fixed, constant support set, it will overfit to that input and be unable to make use of novel examples at test time. Our training procedure uses all available data as the training set $\mathcal{T}$, and we form a different support set $\mathcal{S}$ at each training epoch as a random subset of $\mathcal{T}$. The procedure is summarized in Algorithm 1. Note that in practice, it is parallelized to process instances in mini-batches rather than individually.

5. Experiments

We perform a series of experiments to evaluate (1) how effectively the proposed model and its different components can use the support set, (2) how useful novel support instances are for VQA, (3) whether the model learns different aspects of a dataset from classical VQA methods trained in the classical setting.

Datasets The *VQA v2* dataset [15] serves as the principal current benchmark for VQA. The heavy class imbalance among answers makes it very difficult to draw meaningful conclusions or perform a qualitative evaluation, however. We additionally propose a series of experiments on a subset referred to as *VQA-Numbers*. It includes all questions marked in *VQA v2* as a "number" question, which are further cleaned up to remove answers appearing less than 1,000 times in the training set, and to remove questions that do not have an unambiguous answer (we keep only those with ground truth scores containing a single element equal to 1.0). Questions from the original validation set of *VQA v2* are used for evaluation, and the original training set (45,965 questions after clean up) is used for training, support, and validation. The precise data splits will be available publicly. Most importantly, the resulting set of candidate answers corresponds to the seven numbers from 0 to 6.

Metrics The aggregate metric used for evaluation on *VQA v2* is the accuracy defined as $\frac{1}{|\mathscr{E}|} \sum_i \hat{s}^{a^\star_i}_i$ with ground truth scores $\hat{s}_i$ and $a^\star_i$ the answer of highest predicted score, $\arg\max_a s^a_i$. We also define the recall of an answer a as $\sum_i s^{a^\star}_i \ / \ \sum_i \hat{s}^a_i$. We look at the recall averaged (uniformly) over all possible answers to better reflect performance across a variety of answers, rather than on the most common ones.

Implementation Our implementation is based on the code provided by the authors of [34]. Details non-specific to our contributions can be found there. We initialize all parameters, in particular static weights and static prototypes as if they were those of a linear layer in a traditional architecture, following Glorot and Bengio [14]. During training, the support set is subsampled (Section 4.2) to yield a set of 1,000 elements. We use, per answer, one or two static prototypes, and zero or one dynamic prototype (as noted in the experiments). All experiments use an embedding dimension D=128 and a mini-batches of 256 instances. Experiments with *VQA v2* use a set of candidate answers capped to a minimum number of training occurrences of 16, giving 1,960 possible answers [34]. Past works have shown that small differences in implementation can have noticeable impact on performance. Therefore, to ensure fair comparisons, we repeated all evaluations of the baseline [34] with our code and preprocessing. Results are therefore not directly comparable with those reported in [34]. In particu-

	Average answer recall
(1a) Chance	14.28
(1b) State-of-the-art model [34]	29.72
Equivalent to 1 static prototype per answer, dot prod. similarity, no dynamic param.	
(2b) 1 Static prot./ans., L1 similarity	29.97
(2c) 1 Static prot./ans., L2 similarity	27.80
(2d) 2 Static prot./ans., dot prod. similarity	30.28
(2e) 2 Static prot./ans., L1 similarity	28.34
(2f) 2 Static prot./ans., L2 similarity	31.48
(3a) Dynamic Weights (+2f)	31.81
(3b) **Proposed: dynamic weights and prototypes** (+2f)	**32.32**

Table 1. Ablative evaluation on *VQA-Numbers*, trained and evaluated on all answers. See discussion in Section 5.1.

lar, we do not use the *Visual Genome* dataset [22] for training.

5.1. VQA-Numbers

Ablative evaluation   We first evaluate the components of the proposed model in comparison to the state-of-the-art of [34] which serves as a baseline, being equivalent to our model with 1 static prototype per answer, the dot product similarity, and no dynamic parameters. We train and evaluate on all 7 answers. To provide the baseline with a fair chance[2], we train all models with standard supersampling [9, 16], *i.e.* selecting training examples with equal probability with respect to their correct answer. In these experiments, the support set is equal to the training set.

As reported in Table 1, the proposed dynamic weights improve over the baseline, and the dynamic prototypes bring an additional improvement. We compare different choices for the similarity function. Interestingly, swapping the dot product in the baseline for an L2 distance has a negative impact. When using two static prototypes however, the L2 distances proves superior to the L1 or to the dot product. This is consistent with [33] where a prototypes network also performed best with an L2 distance.

Additional Support Set and Novel answers   We now evaluate the ability of the model to exploit support data never seen until test time (see Fig. 3). We train the same models designed for 7 candidate answers, but only provide them with training data for a subset of them. The proposed model is additionally provided with a complete support set, covering all 7 answers. Each reported result is averaged over 10 runs. The set of k answers excluded from training is randomized across runs but identical to all models for a given k.

The proposed model proves superior than the baseline and all other ablations (Fig. 3, top). The dynamic prototypes are particularly beneficial. With very little training data, the use of dynamic weights is less effective and sometimes even

---

[2]The *VQA-Numbers* data is still heavily imbalanced, "1" and "2" making up almost 60% of correct answers in equal parts.
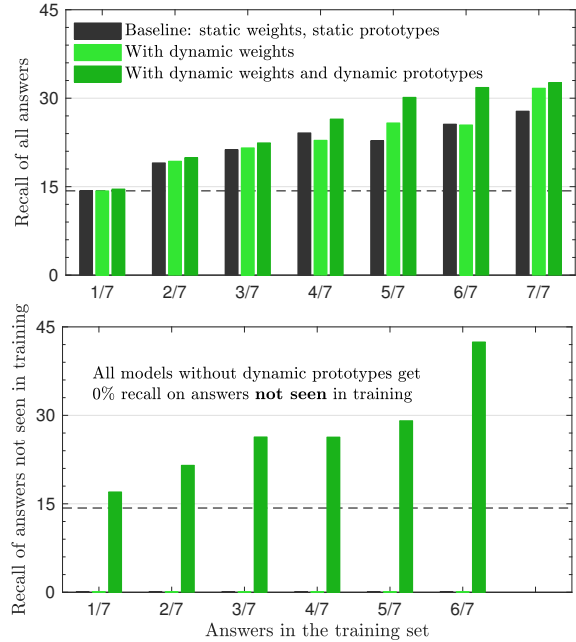

Figure 3. Performance of the proposed model and ablations on *VQA-Numbers* with training data for a subset of the 7 answers. (Top) Performance on all answers. (Bottom) Performance on answers not seen in training. Only the model with dynamic prototypes makes this setting possible. Remarkably, a model trained on two answers (2/7) maintains a capacity to learn about all others. Chance baseline shown as horizontal dashes.

detrimental. We hypothesize that the model may then suffer from overfitting due to the additional learned parameters. When evaluated on novel answers (not seen during training and only present in the test-time support set), the dynamic prototypes provide a remarkable ability to learn those from the support set alone (Fig. 3, bottom). Their efficacy is particularly strong when only a single novel answer has to be learned. Remarkably, a model trained on only two answers maintains some capacity to learn about all others (average recall of 17.05%, versus the chance baseline of 14.28%). Note that we cannot claim the ability of the model to count to those novel numbers, but at the very least it is able to associate those answers with particular images/questions (possibly utilizing question-conditioned biases).

5.2. VQA v2

We performed experiments on the complete *VQA v2* dataset. We report results of different ablations, trained with 50% or 100% of the official training set, evaluated on the validation set as in [34]. The proposed model uses the remaining of the official training set as additional support data at test time. The complexity and varying quality of this dataset do not lead to clear-cut conclusions from the standard accuracy metric (see Table 2). The answer recall leads to more consistent observations that align with those
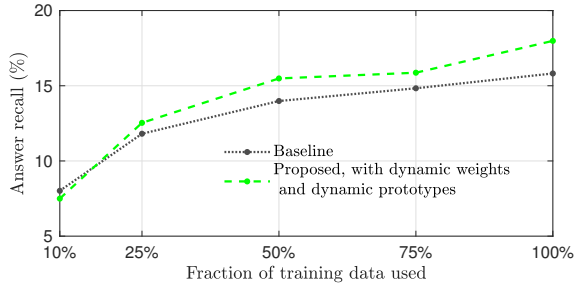
Figure 4. Performance using varying amounts of training data on *VQA v2*. See Section 5.2.

	Question accuracy / Answer recall	
	Trained on 50%	Trained on 100%
Baseline [34]	**57.6** / 14.0	59.8 / 15.8
Proposed model		
With dynamic weights, no dynamic prototypes	**57.6** / 14.1	**60.0** / 16.3
No dynamic weights, **with** dynamic prototypes	**57.6** / 15.2	59.7 / **18.0**
Same, no static prototypes, only dyn. ones	57.2 / 3.6	58.6 / 4.29
With dyn. weights and dyn. prototypes	57.5 / **15.5**	59.9 / **18.0**

Table 2. Evaluation on *VQA v2*. The proposed method exhibits qualitatively different strengths than the classical approach of the baseline [34], producing a generally higher recall (averaged over possible answers) and lower accuracy (averaged over questions).
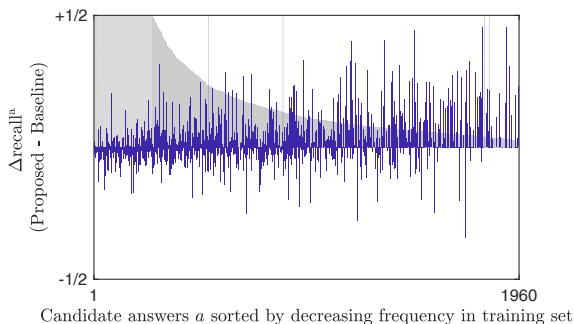

Figure 5. Difference in answer recall between the proposed model (Table 2, last row, last column) and the baseline (Table 2, first row, last column) on *VQA v2*. Each blue bar corresponds to one of the candidate answers, sorted decreasing number of occurrences in the training set (gray background, units not displayed). The two models show qualitatively different behaviour: the baseline is effective with frequent answers, but the proposed model fares better (mostly positive values) in the long tail of rare answers.

made on *VQA-Numbers*. Both dynamic weights and dynamic parameters provide a consistent advantage (Fig. 4). Each technique is beneficial in isolation, but their combination performs generally best. Individually, the dynamic prototypes appear more impactful than the dynamic weights. Note that our experiments on *VQA v2* aimed at quantifying the effect of the contributions in the meta learning setting, and we did not seek to maximize absolute performance in the traditional benchmark setting.

To obtain a better insight into the predictions of the model, we examine the individual recall of possible answers. We compare the values with those obtained by the baseline. The difference (Fig. 5) indicates which of the two models provides the best predictions for every answer. We observe a qualitatively different behaviour between the models. While the baseline is most effective with frequent answers, the proposed model fares better (mostly positive values) in the long tail of rare answers. This corroborates previous discussions on dataset biases [15, 18, 43] which classical models are prone to overfit to. The proposed model is inherently more robust to such behaviour.

6. Conclusions and Future Work

We have devised a new approach to VQA through framing it as a meta learning task. This approach enables us to provide the model with supervised data at test time, thereby allowing the model to adapt or improve as more data is made available. We believe this view could lead to the development of scalable VQA systems better suited to practical applications. We proposed a deep learning model that takes advantage of the meta learning scenario, and demonstrated a range of benefits: improved recall of rare answers, better sample efficiency, and a unique capability of to learn to produce novel answers, *i.e.* those never seen during training, and learned only from support instances.

The learning-to-learn approach we propose here enables a far greater separation of the questions answering method from the information used in the process than has previously been possible. Our contention is that this separation is essential if vision-and-language methods are to move beyond benchmarks to tackle real problems, because embedding all of the information a method needs to answer real questions in the model weights is impractical.

Even though the proposed model is able to use novel support data, the experiments showed room for improvement, since a model trained initially from the same amount of data still shows superior performance. Practical considerations should also be addressed to apply this model to a larger scale, in particular for handling the memory of dynamic weights that currently grows linearly with the support set. Clustering schemes could be envisioned to reduce its size [33] and hashing methods [4, 19] could improve the efficiency of the content-based retrieval.

Generally, the handling of additional data at test time opens the door to VQA systems that interact with other sources of information. While the proposed model was demonstrated with a support set of questions/answers, the principles extend to any type of data obtained at test time *e.g.* from knowledge bases or web searches. This would drastically enhance the scalability of VQA systems.

References

[1] A. Agrawal, A. Kembhavi, D. Batra, and D. Parikh. C-vqa: A compositional split of the visual question answering (vqa) v1. 0 dataset. *arXiv preprint arXiv:1704.08243*, 2017. 2

[2] R. Aljundi, P. Chakravarty, and T. Tuytelaars. Expert gate: Lifelong learning with a network of experts. *arXiv preprint arXiv:1611.06194*, 2016. 3

[3] P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, and L. Zhang. Bottom-up and top-down attention for image captioning and vqa. *arXiv preprint arXiv:1707.07998*, 2017. 4

[4] A. Andoni and P. Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 459–468. IEEE, 2006. 8

[5] M. Andrychowicz, M. Denil, S. Gomez, M. W. Hoffman, D. Pfau, T. Schaul, and N. de Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, pages 3981–3989, 2016. 3

[6] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. VQA: Visual Question Answering. In *Proc. IEEE Int. Conf. Comp. Vis.*, 2015. 2

[7] L. Bertinetto, J. F. Henriques, J. Valmadre, P. H. S. Torr, and A. Vedaldi. Learning feed-forward one-shot learners. In *NIPS*, pages 523–531, 2016. 3, 4, 11

[8] C. Buck, J. Bulian, M. Ciaramita, A. Gesmundo, N. Houlsby, W. Gajewski, and W. Wang. Ask the right questions: Active question reformulation with reinforcement learning. *arXiv preprint arXiv:1705.07830*, 2017. 3

[9] M. Buda, A. Maki, and M. A. Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *arXiv preprint arXiv:1710.05381*, 2017. 7

[10] K. Cho, B. van Merrienboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proc. Conf. Empirical Methods in Natural Language Processing*, 2014. 4

[11] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier. Language modeling with gated convolutional networks. *arXiv preprint arXiv:1612.08083*, 2016. 4

[12] C. Fernando, D. Banarse, C. Blundell, Y. Zwols, D. Ha, A. A. Rusu, A. Pritzel, and D. Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017. 3

[13] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017. 3

[14] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proc. Int. Conf. Artificial Intell. & Stat.*, pages 249–256, 2010. 6

[15] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. *arXiv preprint arXiv:1612.00837*, 2016. 2, 6, 8

[16] H. Guo, Y. Li, J. Shang, G. Mingyun, H. Yuanyue, and G. Bing. Learning from class-imbalanced data: Review of methods and applications. *Expert Syst. Appl.*, 73:220–239, 2017. 7

[17] S. Hochreiter, A. S. Younger, and P. R. Conwell. Learning to learn using gradient descent. In *International Conference on Artificial Neural Networks*, pages 87–94. Springer, 2001. 3

[18] A. Jabri, A. Joulin, and L. van der Maaten. Revisiting visual question answering baselines. 2016. 4, 5, 8

[19] L. Kaiser, O. Nachum, A. Roy, and S. Bengio. Learning to remember rare events. *CoRR*, 2017. 3, 8

[20] V. Kazemi and A. Elqursh. Show, ask, attend, and answer: A strong baseline for visual question answering. *arXiv preprint arXiv:1704.03162*, 2017. 4

[21] J. Kirkpatrick, R. Pascanu, N. C. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell. Overcoming catastrophic forgetting in neural networks. *arXiv preprint arXiv:1612.00796*, 2016. 3

[22] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. Bernstein, and L. Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *arXiv preprint arXiv:1602.07332*, 2016. 2, 6

[23] D. Lopez-Paz and M. Ranzato. Gradient episodic memory for continuum learning. *arXiv preprint arXiv:1706.08840*, 2017. 3

[24] T. Munkhdalai and H. Yu. Meta networks. In *International Conference on Machine Learning (ICML)*, pages 2554–2563, 2017. 2, 3

[25] K. Narasimhan, A. Yala, and R. Barzilay. Improving information extraction by acquiring external evidence with reinforcement learning. *arXiv preprint arXiv:1603.07954*, 2016. 3

[26] R. Nogueira and K. Cho. Task-oriented query reformulation with reinforcement learning. *arXiv preprint arXiv:1704.04572*, 2017. 3

[27] H. Noh, P. H. Seo, and B. Han. Image Question Answering using Convolutional Neural Network with Dynamic Parameter Prediction. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 3

[28] J. Pennington, R. Socher, and C. Manning. Glove: Global Vectors for Word Representation. In *Conference on Empirical Methods in Natural Language Processing*, 2014. 4

[29] S. K. Ramakrishnan, A. Pal, G. Sharma, and A. Mittal. An empirical evaluation of visual question answering for novel objects. *arXiv preprint arXiv:1704.02516*, 2017. 2

[30] S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. 2017. 3

[31] S. Rebuffi, A. Kolesnikov, and C. H. Lampert. icarl: Incremental classifier and representation learning. *arXiv preprint arXiv:1611.07725*, 2016. 3

[32] J. Schmidhuber. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Learning*, 4(1), 2008. 3

[33] J. Snell, K. Swersky, and R. S. Zemel. Prototypical networks for few-shot learning. *arXiv preprint arXiv:1703.05175*, 2017. 2, 3, 7, 8

[34] D. Teney, P. Anderson, X. He, and A. van den Hengel. Tips and tricks for visual question answering: Learnings from the 2017 challenge. *arXiv preprint arXiv:1708.02711*, 2017. 2, 3, 4, 5, 6, 7, 8

[35] D. Teney and A. van den Hengel. Zero-shot visual question answering. 2016. 2, 3, 5

[36] E. Triantafillou, R. Zemel, and R. Urtasun. Few-shot learning through an information retrieval lens. *arXiv preprint arXiv:1707.02610*, 2017. 3

[37] P. Wang, Q. Wu, C. Shen, and A. van den Hengel. The VQA-Machine: Learning how to use existing vision algorithms to answer new questions. *arXiv preprint arXiv:1612.05386*, 2016. 4

[38] P. Wang, Q. Wu, C. Shen, A. van den Hengel, and A. Dick. Explicit knowledge-based reasoning for visual question answering. *arXiv preprint arXiv:1511.02570*, 2015. 3

[39] Q. Wu, D. Teney, P. Wang, C. Shen, A. Dick, and A. van den Hengel. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding*, 2017. 1, 2, 3, 4

[40] Q. Wu, P. Wang, C. Shen, A. Dick, and A. v. d. Hengel. Ask Me Anything: Free-form Visual Question Answering Based on Knowledge from External Sources. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 3

[41] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola. Stacked Attention Networks for Image Question Answering. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 4

[42] J. Yoon, E. Yang, J. Lee, and S. ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017. 3

[43] P. Zhang, Y. Goyal, D. Summers-Stay, D. Batra, and D. Parikh. Yin and yang: Balancing and answering binary visual questions. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 2, 8

[44] Y. Zhu, O. Groth, M. Bernstein, and L. Fei-Fei. Visual7W: Grounded Question Answering in Images. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 2

Supplementary material

A. Factorized Non-Linear Transformation

We defined in Eq. 1 our non-linear transformations using weights (static or dynamic) $\boldsymbol{\theta}$ containing all the parameters of a gated tanh layer. Although this is sound in principle, it is computationally costly to handle a memory containing dynamic weights of such a large dimensionality $(2D^2 + 2D)$. To alleviate this, we follow [7] and factorize the parameters of the gated tanh layer, rewriting Eq. 1 as follows:

$$f_{\boldsymbol{\theta}}(\boldsymbol{h}) = \sigma(\boldsymbol{w}_a W_b \boldsymbol{h} + \boldsymbol{b}) \circ \tanh(\boldsymbol{w}_a' W_b' \boldsymbol{h} + \boldsymbol{b}') \quad (7)$$

with vectors $\boldsymbol{w}_a, \boldsymbol{w}_a' \in \mathbb{R}^D$ and matrices $W_b, W_b' \in \mathbb{R}^{D \times D}$. The matrices W_b and W_b' are learned like traditional weights, and only $\boldsymbol{w}_a$ and $\boldsymbol{w}_a'$ are those incorporated into the vector of weights $\boldsymbol{\theta}$ (static or dynamic). This reduces the dimensionality of $\boldsymbol{\theta}$ from $(2D^2 + 2D)$ to $4D$ (accounting for $\boldsymbol{w}_a$, $\boldsymbol{w}_b$, $\boldsymbol{b}$, and $\boldsymbol{b}'$).

B. Bias in the Output Mapping

In Eq. 3, the output mapping to answer scores uses a *scalar* bias term. This is different to the vector bias in a classical model that use a linear (affine) layer to implement $g_{\Phi}(\cdot)$. A vector contains a value for each output class (*i.e.* each candidate answer), whereas our formulation uses a single value shared among all of them. Our formulation helps avoid the model incorporating biases towards frequent training answers, as discussed in Section 1 and 2. This feature is essential to enable the capability of our model to produce novel answers (unseen during training) only demonstrated by instances in the support set. A vector of answer-specific biases would prevent this capability, as the bias for the novel answers can not be learned from training data.

Completely removing the bias term is another option. At test time, it is without effect compared to a scalar bias, since it is only followed by a (monotonic) sigmoid. Removing the bias term however renders the training by gradient descent numerically unstable, because the chosen similarity function can map to saturating regions of the domain of the sigmoid.

C. VQA-Numbers Dataset

We provide below statistics of the *VQA-Numbers* dataset.

| | Correct answer | | | | | | Row |
0	1	2	3	4	5	6	sum
Training/validation/support set							
2,529	8,193	7,030	2,485	1,520	579	602	22,938
11.0%	35.7%	30.7%	10.8%	6.6%	2.6%	2.6%	100%
Test set							
858	2,804	2,434	843	495	173	205	7,812
11.0%	35.9%	31.2%	10.8%	6.3%	2.2%	2.6%	100%

Table 3. Number and ratio of questions per type of correct answer in the *VQA-Numbers* dataset.

D. VQA-Numbers Experiments

Our experiments on *VQA-Numbers* (Section 5.1) use supersampling during training to ensure that none of the compared models can be influenced by dataset biases (*i.e.* class imbalance). The supersampling is performed at the epoch level, not at the mini-batch level. Concretely, the training instances of all answers (classes) except the most frequent one are repeated at random such that there are as many of them as instances with the most frequent one. The elements within a mini-batch are selected at random. Each training epoch will thus go through every training instance at least once. We did not try to constrain the sampling within mini-batches.

Note finally that the supersampling strategy is practical on VQA-Numbers thanks to the small number of classes and only mild imbalance. It would not be suitable to *VQA v2* for the opposite reasons.
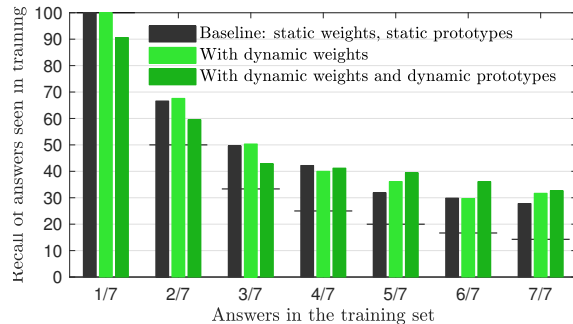

Figure 6. Additional results of the experiments of Fig. 3. Performance is reported here only on answers seen in the training set. See text for discussion.

We provide in Fig. 6 additional results of the experiments of Section 5.1. We report the performance of the same models, now evaluated only on answers present in the training set. The number of those answers is varied from 1 to 7. The chance performance (gray dashes) diminishes as the number of possible answers gets larger. As expected, the

baseline model performs at 100% recall in the trivial case of 1 possible answer. The proposed model however receives a support set containing examples of all other answers (*i.e.* all 7 of them). This explains the non-perfect result in the trivial case. As the number of possible answers is increased, the proposed models (with dynamic weights and dynamic prototypes) shows a growing advantage and finally surpass the baseline by a clear margin on the most interesting cases of 5, 6, and 7 answers.