

Towards Contextual Learning in Few-shot Object Classification

Mathieu Pagé Fortin
Laval University, QC, CA

mathieu.page-fortin.1@ulaval.ca

Brahim Chaib-dra
Laval University, QC, CA

brahim.chaib-draa.1@ift.ulaval.ca

Abstract

Few-shot Learning (FSL) aims to classify new concepts from a small number of examples. While there have been an increasing amount of work on few-shot object classification in the last few years, most current approaches are limited to images with only one centered object. On the opposite, humans are able to leverage prior knowledge to quickly learn new concepts, such as semantic relations with contextual elements.

Inspired by the concept of contextual learning in educational sciences, we propose to make a step towards adopting this principle in FSL by studying the contribution that context can have in object classification in a low-data regime. To this end, we first propose an approach to perform FSL on images of complex scenes. We develop two plug-and-play modules that can be incorporated into existing FSL methods to enable them to leverage contextual learning. More specifically, these modules are trained to weight the most important context elements while learning a particular concept, and then use this knowledge to ground visual class representations in context semantics. Extensive experiments on Visual Genome and Open Images show the superiority of contextual learning over learning individual objects in isolation.

1. Introduction

Whereas Convolutional Neural Networks are currently the state-of-the-art models for object recognition tasks, they generally require a large number of examples from each class to perform well. In the last few years an increasing amount of effort has been done in zero-shot learning (ZSL) and few-shot learning (FSL) to develop approaches that reduce the number of examples required to train efficient models. Progress in this direction is important for solving many real-world problems for which labeling is hard, and for enabling new applications, such as robots that actively learn new concepts on the fly from their environment [29].

However, most current FSL methods focus on visual features and tend to only consider objects in isolation [1, 6, 19,

33, 34, 38, 40, 43, 45]. These methods are therefore primarily evaluated on datasets of images with only one centered object (e.g. *mini*Imagenet [38], Omniglot [18], CUB-200 [41]). On the other hand, images in real-world applications can be more complex, containing many different objects. This scenario has been neglected so far, and while it can be more relevant for real applications, we argue that complex scenes also offer an interesting opportunity to supplement visual information with contextual and semantic relations between concepts. This idea is motivated by the principle of *contextual learning* [15] in educational sciences, and more specifically by the first functional feature of context defined by Dohn *et al.* [7]:

Supplementary role of context. [*Context*] is brought in, or added to, the understanding of a phenomenon—the focal object—that would not have been adequately understood had it been considered in isolation. A context thus completes the conditions for understanding the focal object. [7]

In the few-shot setting, the model has access to a set of *base* classes with many examples and is evaluated on its ability to learn *novel* classes from few examples. This means that when a *novel* class is presented in its context to the model, some *base* classes can also appear in the scene. This is similar to when humans see an unknown object in a familiar scene, for instance a corkscrew in a kitchen. To learn this new label, our brain will not only process the appearance of the object, but also the contextual and semantic relations with other objects [3], for instance a bottle of wine. This is fundamental to learn quickly and for continual incremental learning in humans, as new concepts are generally not learned in isolation, but often in relation to already known concepts [3, 27]. We apply this principle to FSL by proposing a method to add contextual semantic information in visual representations: our model learns to refine class prototypes according to the context in which training examples appear.

However, this low-data setting poses the challenge of context generalization: *How to model the context such that this can generalize to novel classes with only few examples of scenes?* Intuitively, one could represent the context of a

class i by co-occurrence counts, i.e. by counting how many times the class i occurs in the same scene as all other classes in the training set. However, this is restrictive for at least two reasons. First, co-occurrence counts assume that the training set is a sufficient sample to estimate the large variety of co-occurrences that can happen, which is certainly not the case in a few-shot setting. Second, co-occurrence counts are limited to statistical regularities and they ignore the potential semantic relations within each class, which could help building powerful representations.

Instead, we make use of transfer learning from a text model pretrained on large corpora to represent the semantics of classes. Previous work has shown that word embeddings implicitly encode high-level relations between entities [11, 25, 44], which could help context generalization in a few-shot setting by enabling the model to capture higher-order relations. For instance, if in the small training set plates co-occur with forks, a semantic-aware model could infer that plates are also likely to co-occur with spoons because of the semantic similarity between forks and spoons. On the other hand, if there are irrelevant entities in the scene, such as “wall”, which is generally not very informative, or out-of-context objects that could induce confusion [31], an ideal context-aware model should ignore them. These aspects of *contextual learning* are defined by the second functional feature of context of Dohn *et al.* [7]:

Relative role of context. *The context is centered around the object. The context is not a neutral layout of things or properties near the focal object, nor is it a set of circumstances or an indefinite “background”. It is ordered and organized by its relations to the focal object, which co-determines what properties of the surroundings are relevant and thus part of the context.* [7]

We integrate this principle in our approach by proposing a Class-conditioned Context Attention Module (CCAM) such that our model can learn to attend to context elements that are relevant to the focal object. Psychological studies showed that contextual cueing in humans improves object classification in scenes, by capitalizing on the fact that most objects co-occur more often with certain objects and not others [27]. But additionally, the relative role of context specifies that not all co-occurrences are equal [7]. Our experiments show that CCAM effectively weights discriminative contextual objects more strongly.

In summary, our main contributions are the following:

- We propose a few-shot model that learns class representations grounded in contextual semantics. To this end, we propose a gated visuo-semantic unit (GVSU), a flexible module to combine visual prototypes with contextual semantic information.
- We propose CCAM, an attention module applied on context elements that automatically learns to attend to

the most important entities in scenes relatively to the focal object.

- We conduct extensive experiments on Visual Genome [16] and Open Images [17], which are large-scale datasets of complex scenes with hundreds of classes. Our results support that using context is valuable in a few-shot setting.
- As an auxiliary result, we observe that our model implicitly learns semantic word embeddings grounded in scene context.

2. Related Work

FSL approaches can be divided into *gradient-based* methods [8, 9, 13] and *metric learning based* methods [4, 10, 20, 33, 34]. Gradient-based methods aim to improve the training procedure, which MAML [8] is a typical example. MAML is a meta-learning algorithm that aims to generalize such that new tasks can be learned with few update steps. On the opposite, metric learning approaches aim to learn a metric space where support (train) examples are embedded such that query (test) examples can be classified based on a distance metric (e.g. euclidean distance [33], cosine similarity [10], or Mahalanobis distance [4]) without requiring any parameter update to learn novel classes. Our work is more closely related to metric learning and that is why we focus on this family of approaches for the rest of this section.

Few-shot image classification. Several FSL approaches build on Prototypical Networks [33]. This method learns a metric space by computing class centroids from the examples in the support set. It then compares query image embeddings with these prototypes and assigns a class by performing nearest neighbor search. Other approaches consider different ways to compare support and query embeddings, such as Relation Networks [34] that automatically learn the distance function with a neural network, or the approach of Li *et al.* [20] that compares support and query images based on several descriptors.

Other approaches leverage relations between classes to transfer visual features of *base* classes to *novel* ones. For instance, Wang *et al.* [39] proposed to use a Graph Convolutional Network to transfer features between classes based on a knowledge base that encodes relations between these categories. Li *et al.* [19] developed an approach that learns from predicting class hierarchies, which facilitates feature transfer. A similar idea has been proposed to transfer explicit attributes between classes [1]. In our work, we use another form of transfer learning between classes. We leverage the presence of *base* classes when *novel* classes are presented in a scene to adapt their representation in metric space. By

doing so, our model can benefit from contextual cueing [27] when a new instance is seen in a complex scene.

Auxiliary semantics in FSL. Recently, additional cues that were only considered in ZSL have proved to be also useful in FSL, especially when the quantity of examples for each class is very low. Xing *et al.* [43] built on Prototypical Networks [33] by adding word embeddings in the formation of class prototypes in their approach called AM3. This improves the accuracy in 1-shot by almost 10% on *mini*Imagenet [38] and by 5% on CUB-200 [41]. Furthermore, Schwartz *et al.* [32] built on AM3 by additionally using text descriptions of classes extracted from WordNet, and thereby improved 1-shot accuracy by an additional 2% on *mini*Imagenet.

Context semantics. All the FSL work cited above focus on visual information. Indeed, even the attributes, word embeddings and text descriptions that are used in these work need to encode features that can be detected visually from the appearance of a new concept. However, this is not the only form of semantic information that can be available in images. The presence of other objects in a scene can also inform which classes are more or less likely to appear [3, 27, 44]. The context has been used recently to improve object detection within deep learning models in a standard setting with large datasets [5, 22, 26, 42].

Recently Zablocki *et al.* [44] introduced the use of scene context in ZSL. They showed that their model, with the use of Word2vec [25] embeddings, could learn to rank unseen classes according to their likelihood of appearing in an image given the presence of other objects. This suggests that word embeddings implicitly encode co-occurrences of other classes in real visual scenes, even if they have been trained on text corpora [25]. This is closely related to the distributional hypothesis [12], which states that *words that appear in similar contexts often have similar meanings*. This is exploited by skip-gram and continuous bag-of-words (CBOW) models, and the results from Zablocki *et al.* [44] suggest that this principle could also generalize to visual scenes: *items denoted by words that have similar meanings tend to occur in similar scenes*. This idea has been explored by Lüddecke *et al.* [24], where they proposed a method to learn semantic word embeddings explicitly from images showing objects in context.

To the best of our knowledge, scene context has not been used in previous FSL work. In this paper, we argue that FSL would benefit from considering objects with their context, because it places new concepts in relation with previously acquired knowledge. FSL models could then leverage these relations to benefit from contextual cueing [27] in new scenes. We introduce this idea of using scene context in FSL to perform object classification in complex images

by building on Prototypical Networks [33] and by learning class prototypes grounded in context. Unlike Zablocki *et al.* [44], our model jointly learns to embed visual information with context semantics and word embeddings of class labels. Furthermore, whereas some recent semantic-based approaches use relations between classes to share common visual features, our use of scene context with word embeddings exploit a different form of semantic relation which is complementary and orthogonal to visual features.

3. Our model

3.1. Preliminaries

FSL aims to solve the problem of M -way K -shot classification, where M is the number of classes in a given task and K is a small number of examples for each class. Generally, few-shot models are trained on a large dataset \mathcal{D}_{train} with a set of *base* classes that is disjoint from the *novel* categories in \mathcal{D}_{test} . The goal is to learn a representation model f_θ on \mathcal{D}_{train} such that it can learn to recognize novel categories with only K examples. This is generally done by simulating the episodic test scenario of M -way K -shot classification during training. That is, even if a large number of examples are available for each class at train time, f_θ is trained by sampling at each episode e (1) a *support* set $\mathcal{S}_e = \{(x_i, y_i)\}_{i=1}^{M \times K}$ that contains K examples for each M class and (2) a *query* set $\mathcal{Q}_e = \{(q_j, y_j)\}_{j=1}^{n_q}$ containing n_q images of the same classes sampled in the support set. The model is then trained according to the cross-entropy loss:

$$\mathcal{L}(\theta) = -\frac{1}{n_q} \sum_{t=1}^{n_q} \log p_\theta(y_t | q_t, \mathcal{S}_e) \quad (1)$$

Prototypical networks [33] offer a simple and efficient way to model $p_\theta(y|q, \mathcal{S}_e)$. Each of the images in the support set are embedded by a CNN denoted by $f_\theta : \mathbb{R}^D \rightarrow \mathbb{R}^{d_x}$. Then a prototype is built for each class by averaging the K vector embeddings from the same class:

$$\mathbf{c}_k = \frac{1}{|\mathcal{S}_k|} \sum_{(x_i, y_i) \in \mathcal{S}_k} f_\theta(x_i) \quad (2)$$

Finally, the class distribution of a query image q is assigned by computing the softmax over the euclidean distances d of its embedding $f_\theta(q)$ and all class prototypes \mathbf{c}_k :

$$p_\theta(y = k | q, \mathcal{S}_e) = \frac{\exp(-d(f_\theta(q), \mathbf{c}_k))}{\sum_{k'} \exp(-d(f_\theta(q), \mathbf{c}_{k'}))} \quad (3)$$

3.2. Context-Aware prototypes learning

Following previous work on the supplementary role of context in learning [7] and classification [3, 27], we propose to learn class prototypes that embed knowledge about

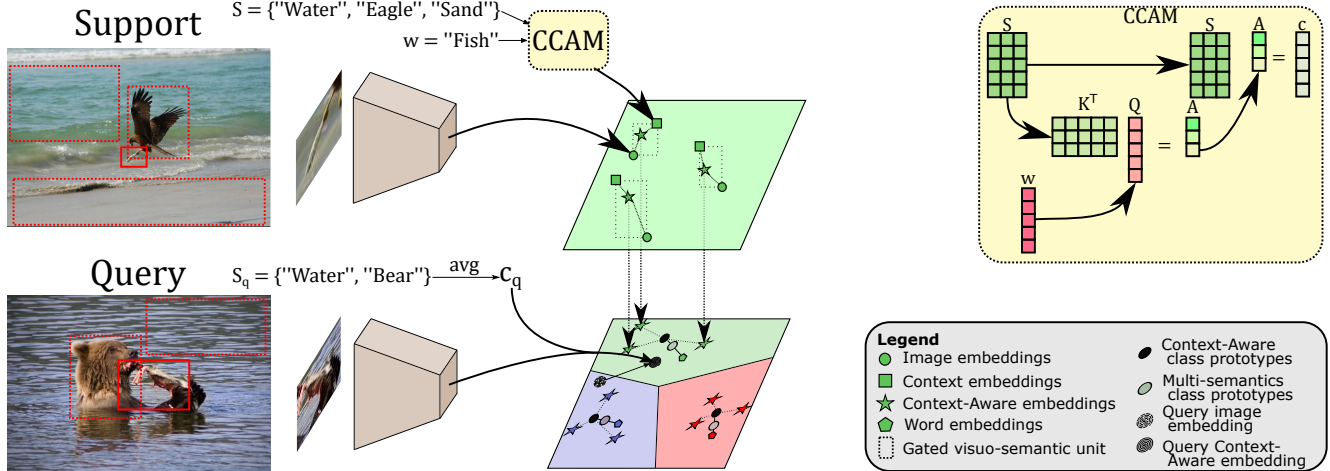


Figure 1: Overview of our model. The focal object (red box with solid lines) is cropped and embedded by a CNN. CCAM computes the relative role of support context elements according to the word embedding w of the class label. Image and context embeddings are projected in a same space (represented by a circle and a square, respectively), and our GVSU produces context-aware embeddings (represented by a star). All context-aware embeddings for a given class are averaged to produce context-aware class prototypes (black dots). Finally multi-semantics class prototypes (light dots) are obtained by refining them with their word embeddings.

their context. To achieve that, we augment the support and query sets with scene context S , $\mathcal{S}_e = \{(x_i, S_i, y_i)\}_{i=1}^{M \times K}$, $\mathcal{Q}_e = \{(q_j, S_j, y_j)\}_{j=1}^{n_q}$, and we adapt the formulation of each prototype c_k as:

$$\hat{c}_k = \frac{1}{|\mathcal{S}_k|} \sum_{(x_i, y_i) \in \mathcal{S}_k} \phi(f(x_i), g(c_i)), \quad (4)$$

where c_i is the context representation of object i obtained from CCAM (see below for more details), $g(c_i)$ is a small neural network projecting c_i in the same space than the image embedding, and $\phi(\cdot, \cdot)$ is a function that adapts the image embedding according to the scene context (see section 3.4). An overview of our approach is shown in Fig. 1. We now describe these components in more details.

3.3. Class-conditioned Context Attention Module (CCAM)

Scene context. We model the scene context by converting surrounding objects from *base* classes into semantic vector representations. This is done by leveraging word embeddings learned from a semantic model such as Word2vec [25] pre-trained on Wikipedia. Therefore, the scene context of an object is represented by the matrix $S \in \mathbb{R}^{d_w \times n_s}$, where d_w is the word embeddings dimension and n_s is the number of surrounding objects.

Class-conditioned Context Attention. The relative role of context [7] suggests that some elements are more important than others when understanding a particular object. For instance, the concept of *bathroom* might be important with

respect to a *toilet*, but could dupe the model while learning the concept of *cat's paw*, even if in the few support examples cats are observed in bathrooms (e.g. see Fig. 2c). To respect this phenomenon, we propose a Class-conditioned Context Attention Module (CCAM) that enables our model to weight the importance of each context elements in S while learning a particular concept w (see CCAM in Fig. 1). This is done by computing a scaled dot-product attention score A [37] between the word embedding w of the class label and each element in S after linear transformations:

$$\begin{aligned} K &= W_K S \\ Q &= W_Q w \\ A &= \text{softmax}\left(\frac{K^T Q}{\sqrt{d_c}}\right) \\ c &= SA, \end{aligned} \quad (5)$$

where $W_K, W_Q \in \mathbb{R}^{d_c \times d_w}$ are weights matrices, and $\frac{1}{\sqrt{d_c}}$ is a scaling factor proposed in [37] to obtain smoother scores. A reflects the relative role of each object in S , which is used to weight the contribution of context entities with respect to the focal object.

Context averaging C_{avg} . Note that the attention mechanism in CCAM is exclusively applied on context from the support set since it depends on the class category w , which is unknown in queries. For query instances, the context representation c_q is simply obtained by averaging all class embeddings w_q in S_q :

$$c_q = \frac{1}{n_s} \sum_{w_q \in S_q} w_q \quad (6)$$

3.4. Gated visuo-semantic unit (GVSU)

To combine visual embeddings with context semantics according to the supplementary role of context [7], we propose a gated visuo-semantic unit, a module to adaptively combine each feature individually from both representations based on a gating mechanism. This is modeled by:

$$\begin{aligned} \phi(f(x), g(c)) &:= z \cdot f(x) + (1 - z) \cdot g(c) \\ h_v &= \tanh(W_v \cdot f(x)) \\ h_c &= \tanh(W_c \cdot g(c)) \\ z &= \sigma(W_z \cdot [h_v, h_c]), \end{aligned} \quad (7)$$

where $W_v \in \mathbb{R}^{d_x \times d_x}$, $W_c \in \mathbb{R}^{d_c \times d_c}$, $W_z \in \mathbb{R}^{d_x \times 2d_z}$ are weights matrices, and σ is the sigmoid function.

Our fusion mechanism is in the same vein than the Gated Multimodal Unit (GMU) [2] that proved to be successful with multimodal representations. Our module slightly differs in that the output representation of the original GMU is $h = z \cdot h_v + (1 - z) \cdot h_c$, whereas in our formulation the intermediate representations h_v and h_c are used to compute the weighting factors to apply on each dimension of $f(x)$ and $g(c)$. The goal of our GVSU is to move the image embeddings in the semantic space according to the context representation, such that objects of the same class will cluster together based on their appearance *and* their contextual semantics.

3.5. Multi-semantics prototypes

Lastly, we add the word embedding of the class label to further refine each prototype. We adopt a similar mechanism that Xing *et al.* [43] proposed to combine visual prototypes with their word embeddings, since it proved to be particularly useful in settings with less data. Our context-aware prototypes are thus refined by:

$$\mathbf{c}'_k = \lambda \cdot \hat{\mathbf{c}}_k + (1 - \lambda) \cdot \hat{w}_k \quad (8)$$

where \hat{w}_k is a transformation of the word embedding w_k and λ is a coefficient between 0 and 1. Both \hat{w}_k and λ are obtained with a two-layer neural network that uses w_k as input.

Finally, the class distribution of a query image q is computed as:

$$p(y = k | q, S_q, \mathcal{S}_e, w) \propto \exp(-d[\phi(f(q), g(c_q)), \mathbf{c}'_k]) \quad (9)$$

In summary, the intuition of our approach is to model the prototypes \mathbf{c}'_k according to the appearance of support examples, the context in which they appear, *and* the semantics

of the class. Moreover, the context of support examples is represented as a weighted sum of the word embeddings of the *base* classes that co-occur in the scene. To perform inference, a query q is compared to these prototypes \mathbf{c}'_k based on their euclidean distance. The appearance of the query object is fused with the context of the scene, which is modeled as the average of the word embeddings of *base* classes in the scene.

3.6. Assumptions

Note that our approach assumes that the context is known, similar to [44] did in a zero-shot setting, as the detection of objects is another task upstream to the FSL problem on which our work focuses. One could remove this assumption by replacing ground-truth annotations with an object detection module. Indeed, note that we form the context with *base* classes only, for which training examples can be in large number during the metric-learning phase. Our approach could also be complementary to the growing number of work on few-shot object detection [14, 21, 30], where classification is generally performed for each object individually without considering the context.

4. Experiments

4.1. Dataset and settings

Traditional FSL datasets such as *mini*Imagenet [38] and CUB-200 [41] mainly contain images with only one object and little context. Therefore, we rather experiment on Visual Genome [16] and Open Images [17], which are large datasets of scenes with several objects in each image.

Visual Genome [16]. We randomly split the images in 70%/10%/20% *train*, *validation* and *test* sets, respectively. We start by using the public splits by Zablocki *et al.* [44] that keep 50% of classes for *base+val* classes and 50% for *novel* classes. However, a closer look at those sets showed that some *novel* classes are very similar to *base* classes, which could bias generalization evaluations. For instance, “bottle” and “television” are in the *base* set, but “bottles” and “TV” are *novel* classes. To solve this issue, we filter the *novel* classes whose Word2vec [25] embeddings have a cosine similarity higher than 0.75 with any of the *base* or *val* classes. It effectively removes singular/plural nouns and closely related concepts such as “police” (*base* class) and “policeman” (*novel* class). This will prevent our model from picking on those biases that would overestimate FSL performance.

We use the bounding box annotations to crop image parts that correspond to objects, and we remove examples whose smallest side is less than 25 pixels. Following this, we remove the classes that appear in less than 10 images. We also form a set of validation classes for hyper-parameter search.

Table 1: Average Top-1 accuracy (%) with 95% confidence intervals on Visual Genome and Open Images . Results are averaged over 4000 test episodes for Visual Genome and 1000 test episodes for Open Images. V: Uses visual information; W: Uses word embeddings; C: Uses contextual information.

Dataset	Model	V	W	C	5-way 1-shot	5-way 5-shot	20-way 1-shot	20-way 5-shot
Visual Genome	PN [33]	✓			52.23 ± 0.76	69.37 ± 0.63	25.71 ± 0.29	42.61 ± 0.70
	AM3 [43]	✓	✓		62.50 ± 0.66	72.07 ± 0.74	34.36 ± 0.32	44.84 ± 0.61
	Ours	✓	✓	✓	71.54 ± 0.57	78.50 ± 0.55	46.13 ± 0.47	54.72 ± 0.49
Open Images	PN [33]	✓			64.60 ± 1.43	80.44 ± 1.25	34.93 ± 0.98	52.73 ± 1.02
	AM3 [43]	✓	✓		69.87 ± 1.00	80.43 ± 1.02	40.25 ± 0.47	52.59 ± 0.94
	Ours	✓	✓	✓	77.42 ± 1.22	87.70 ± 0.95	51.61 ± 0.99	67.53 ± 0.88

This finally results in 969 *base* classes, 242 *val* classes and 829 *novel* classe.

Open Images v6 [17]. We start by using the original *train/val/test* splits¹ of images. Then, we randomly sample 400/100/100 *base/val/novel* classes, respectively. Classes that appear less than 10 times in their respective split of images are removed and the *novel* classes are also filtered based on the Word2vec [25] cosine similarity with *base* and *val* classes (see above). This results in 371/38/57 *base/val/novel* classes.

4.2. Implementation details

We employ a ResNet-12 CNN backbone as described in [28] and we train it from scratch. It is made of 4 blocks with 3 layers of 3×3 convolutions and a 2×2 maxpooling operation at the end of each block. The first block has 64 filters in each layer, and this number is doubled after each block. The last feature map is vectorized by Global Average Pooling, which results in an embedding of 512 dimensions. Image crops from bounding box annotations are rescaled to $84 \times 84 \times 3$.

Each model is trained for 30,000 episodes with Adam optimizer initialized with a learning rate of 10^{-3} and is divided by a factor of 10 every 10,000 episodes. The validation set is used every 3000 episodes to evaluate the mean accuracy of 500 random episodes and early stopping is done based on the best validation accuracy.

5. Results

Since we want to study the contribution of using context information and class semantics in addition to the appearance of objects, we compare our approach to the following models that build on a Prototypical Networks (Protonet) backbone [33]. **ProtoNet** is our implementation of Prototypical Networks [33]. **AM3** is our implementation of the *Adaptive Modality Mixture Mechanism* [43], which supplement the support prototypes with the word embeddings of their corresponding class.

¹<https://storage.googleapis.com/openimages/web/download.html>

Table 2: Average Top-5 accuracy (%) for 50-way and 100-way classification on Visual Genome (VG) and for 50-way and 57-way († all *novel* classes) on Open Images (OI).

D	Model	50-way		100-way/57-way [†]	
		1-shot	5-shot	1-shot	5-shot
VG	PN [33]	39.68	59.52	27.98	47.23
	AM3 [43]	51.78	64.14	38.22	51.57
	Ours	64.10	72.86	50.49	61.19
OI	PN [33]	52.60	73.16	49.15	70.69
	AM3 [43]	59.94	73.03	56.10	69.91
	Ours	69.93	86.11	68.99	84.08

Table 1 shows the results of few-shot episodes on Visual Genome and Open Images. We also show Top-5 accuracy for 50-way and 100-way classification on Visual Genome and 50-way and 57-way classification on Open Images in Table 2.

We can observe that the use of context information outperforms the visual-only ProtoNet [33] and the multimodal AM3 [43] by large margins. Even in the 1-shot setting, where there is a risk of overfitting a particular context since there is only one example of scene, our results on both datasets show that it is still promising to use the context. Interestingly, our results support the benefits of using word embeddings as Xing *et al.* [43] did with AM3, especially in the 1-shot setting. Indeed, using word embeddings increases the accuracy in 5-way and 20-way by 10.27% and 8.65% on Visual Genome, respectively, and by 5.27% and 5.32% on Open Images, respectively. Our use of context further improves these results by an additional 9.04% and 11.77% on Visual Genome, and by 7.55% and 11.36% on Open Images for 5-way and 20-way classification, respectively.

Our model also performs reasonably well on larger-scale experiments shown in Table 2. With only one example per class in 100-way classification on Visual Genome, our model almost doubles the Top-5 accuracy of ProtoNet [33], with 50.49% compared to 27.98%.

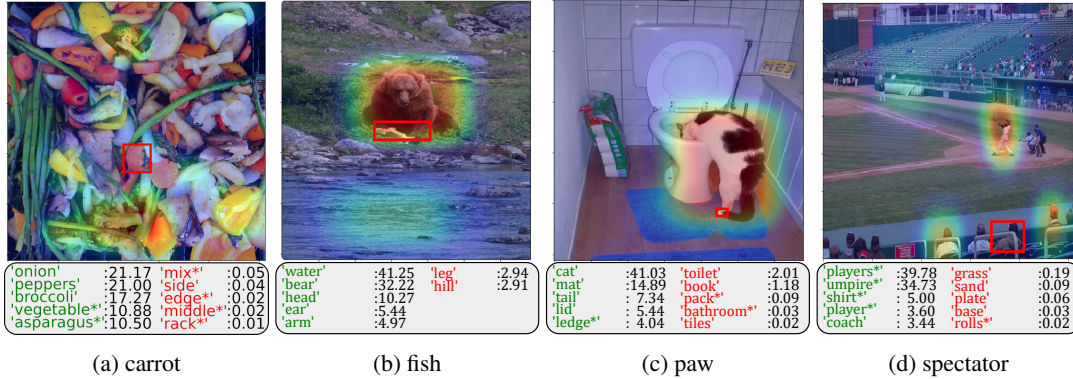


Figure 2: Illustration of the relative role of context estimated by CCAM. Five most important and least important co-occurrent concepts are shown in green and red, respectively, with associated weights (%). Asterisks indicate that the word is a *novel* class. Heatmaps are for visualization purposes. They are computed by summing the weight scores inside their respective bounding box annotations.

Table 3: Average Top-1 accuracy (%) during an ablation study for 5-way classification on Visual Genome. V: Uses visual information; W: Uses word embeddings; C: Uses contextual information.

	V	W	C	C_{avg}	CCAM	1-shot	5-shot
1	✓					52.23	69.37
2	✓	✓				62.50	72.07
3	✓		✓	✓		61.20	76.66
4	✓	✓	✓	✓		69.10	76.63
5	✓		✓		✓	63.56	77.16
6	✓	✓	✓		✓	71.54	78.50

Ablation study. We then performed an ablation study to examine the contribution of each modality and their interactions, shown in Table 3. Interestingly, we observe a synergy between the use of word embeddings and contextual information, which is stronger in the 1-shot setting. For instance, lines 2 and 3 show that the use of word embeddings or context alone perform similarly in 1-shot, with 62.50% and 61.20% respectively; but when word embeddings and context and taken together, the accuracy significantly increases to 69.10% (line 4) and even further to 71.54% when CCAM is used (line 6).

The ablation study also supports the benefits of considering the *relative role of context* [7], as weighting the context elements with CCAM obtains better results than simply considering all of them equally (compare lines 3-4 with lines 5-6). We now investigate this aspect qualitatively.

Relative role of context [7]. The relative role of context states that the importance of context elements is function of the focal object, which is supported by our results (see above). Fig. 2 shows examples of CCAM outputs when our model needed to learn the concepts “carrot”, “fish”, “paw” and “spectator”, respectively. To further examine the abil-

ity of CCAM to meta-learn the importance of co-occurring concepts based on semantic similarity, we also included *novel* classes in the formation of the context for these examples. We can see that CCAM correctly gives more weight to semantically relevant co-occurring concepts, even those that were never encountered during training (e.g. *vegetable*, *players* or *umpire* that are *novel* classes). On the other hand, CCAM also ignores background elements such as “tiles” in Fig. 2c and “grass” in Fig. 2d, as they are unlikely to help recognizing new instances of *paw* and *spectator* in query images.

To further study the contribution of context and CCAM, we show in Fig. 3 a t-SNE visualization [36] of embeddings produced by our model using different amount of information. Visual embeddings (Fig. 3a) seem to produce ambiguous clusters, similar to context averaging (Fig. 3b) which considers each contextual item equally. On the opposite, our CCAM produces good clusters (Fig. 3c), which shows its ability, and the importance, to attend to discriminative elements. Also, some visually different objects seem to share similar contexts, as shown by the mixed cluster in the center of Fig. 3c. This problem is mostly solved by our GVSU that combines visual and contextual information (Fig. 3d).

Robustness to noise. Since we aimed to study the potential contribution of context in a low-data regime, we used an oracle to annotate the context elements. In real applications, one would need to employ an object detection module to classify *base* classes in query images and ideally in the few support examples also, although these could be manually annotated. While this aspect is left for future work, we evaluated the robustness of our model while adding noise in support and query context annotations to simulate an object detection module that makes a certain percentage of errors. With a probability p_{noise} , we randomly swap context elements by another *base* class. The results for 20-way 5-shot

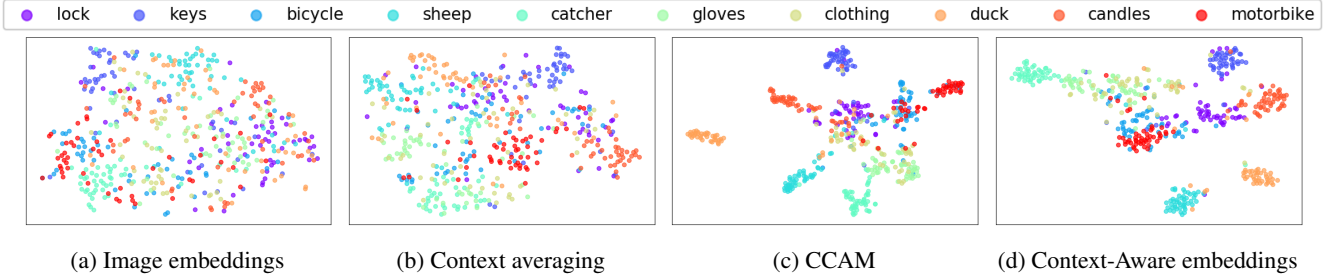


Figure 3: t-SNE visualization of different embeddings.




Figure 4: Robustness to noise in 20-way 5-shot tasks. Context elements in support and query sets are randomly swapped by another *base* class label with a probability p_{noise} .

classification with noise are shown in Fig. 4. Protonet and AM3 are showed as reference but they are not influenced by noise in context annotations since they only use visual and visual+semantic information, respectively. Interestingly, our model seems reasonably tolerant to out-of-context objects. With $p_{noise} = 0.5$, our model still outperforms AM3 [43] and ProtoNet [9].

Learning semantic word embeddings through visual scenes. Our approach also offers an auxiliary result that could be further investigated: CCAM implicitly learns and enriches semantic word embeddings by acting similarly to a CBOW model. To evaluate this aspect, we inputted to CCAM a matrix S that contains all *base* classes and we conditioned its attention on a few words w to see how CCAM would weight S . We show in Table 4 a few examples. Interestingly, the contextual concepts defined by CCAM strongly differ from those obtained with Word2vec [25] embeddings, which shows that CCAM captures different semantic relations between concepts.

6. Conclusion and Future work

In this work, we proposed a few-shot learning model that uses scene context semantics to improve class representa-

Table 4: Examples of words and concepts that received the highest score by CCAM. Underlined words are also in the Top-10 of Word2vec [25] cosine similarities.

Word	Contextual words
bike	<u>cyclists</u> , <u>skateboards</u> , snowboards, guardrail, pedestrians, mopeds, snowmobile, tricycle, kiteboard, <u>motorcycles</u>
rocks	treetops, dunes, thickets, grassy, vegetation, hill, grasslands, sky, <u>cliff</u> , pines
game	referee, softball, scoreboard, basketball, jersey, volleyball, <u>matches</u> , baseball, frisbee, football
rope	boater, surfing, lifeguard, spear, <u>ladders</u> , fisherman, horseback, ropes, cliff, sandals
sandwich	plate, plates, <u>bread</u> s, dough, <u>sandwich</u> es, flatbread, cutlery, pork, pancakes, <u>steak</u> ,

tions. Our approach integrates the supplementary role of context [7] by building context-aware class prototypes, and we apply the relative role of context [7] with our CCAM, a module that proved to be able to focus on discriminative elements in the scene with respect to the class semantics (see Fig. 2, Fig. 3c and Table 4).

Our experiments on Visual Genome and Open Images showed promising results of using context information, by increasing the accuracy of Prototypical Networks [33] and AM3 [43] by large margins. More generally, our multi-semantics model is a step towards holistic approaches of few-shot object classification that can be applied in challenging real scenarios.

As future work, we plan to replace ground-truth class annotations of context by automatic object detection and classification of base classes in queries. Our experiment on robustness to noise suggests that even with many labeling errors, our method could still outperform models that ignore the context. We also plan to explore the ability of our model to *learn semantic word embeddings through visual scenes*. This could be further investigated in line with work on learning multimodal word embeddings [23, 24, 35].

References

- [1] Ziad Al-Halah, Makarand Tapaswi, and Rainer Stiefelhagen. Recovering the missing link: Predicting class-attribute associations for unsupervised zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5975–5984, 2016.
- [2] John Arevalo, Thamar Solorio, Manuel Montes-y Gómez, and Fabio A González. Gated multimodal units for information fusion. *arXiv preprint arXiv:1702.01992*, 2017.
- [3] Moshe Bar. Visual objects in context. *Nature Reviews Neuroscience*, 5(8):617, 2004.
- [4] Peyman Bateni, Raghav Goyal, Vaden Masrani, Frank Wood, and Leonid Sigal. Improved few-shot visual classification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [5] Xinlei Chen and Abhinav Gupta. Spatial memory for context reasoning in object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4086–4096, 2017.
- [6] Zitian Chen, Yanwei Fu, Yu-Xiong Wang, Lin Ma, Wei Liu, and Martial Hebert. Image deformation meta-networks for one-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8680–8689, 2019.
- [7] Nina Dohn, Stig Hansen, and Søren Klausen. On the concept of context. *Education Sciences*, 8(3):111, 2018.
- [8] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- [9] Chelsea Finn, Kelvin Xu, and Sergey Levine. Probabilistic model-agnostic meta-learning. In *Advances in Neural Information Processing Systems*, pages 9516–9527, 2018.
- [10] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4367–4375, 2018.
- [11] Abhijeet Gupta, Gemma Boleda, and Sebastian Padó. Distributed prediction of relations for entities: The easy, the difficult, and the impossible. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pages 104–109, 2017.
- [12] Zellig S Harris. Distributional structure. *Word*, 10(2-3):146–162, 1954.
- [13] Muhammad Abdullah Jamal and Guo-Jun Qi. Task agnostic meta-learning for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11719–11727, 2019.
- [14] Ho-Deok Jang, Sanghyun Woo, Philipp Benz, Jinsun Park, and In So Kweon. Propose-and-attend single shot detector. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, March 2020.
- [15] Elaine B Johnson. *Contextual teaching and learning: What it is and why it's here to stay*. Corwin Press, 2002.
- [16] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, Michael Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. 2016.
- [17] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020.
- [18] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [19] Aoxue Li, Tiange Luo, Zhiwu Lu, Tao Xiang, and Liwei Wang. Large-scale few-shot learning: Knowledge transfer with class hierarchy. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7212–7220, 2019.
- [20] Wenbin Li, Lei Wang, Jinglin Xu, Jing Huo, Yang Gao, and Jiebo Luo. Revisiting local descriptor based image-to-class measure for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7260–7268, 2019.
- [21] Yazhao Li, Yanwei Pang, Jianbing Shen, Jiale Cao, and Ling Shao. Netnet: Neighbor erasing and transferring network for better single shot object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [22] Yong Liu, Ruiping Wang, Shiguang Shan, and Xilin Chen. Structure inference net: Object detection using scene-level context and instance-level relationships. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6985–6994, 2018.
- [23] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. VILBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, pages 13–23, 2019.
- [24] Timo Lüddecke, Alejandro Agostini, Michael Fauth, Minija Tamosiunaite, and Florentin Wörgötter. Distributional semantics of objects in visual scenes in comparison to text. *Artificial Intelligence*, 274:44–65, 2019.
- [25] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- [26] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 891–898, 2014.
- [27] Aude Oliva and Antonio Torralba. The role of context in object recognition. *Trends in cognitive sciences*, 11(12):520–527, 2007.
- [28] Boris Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. Tadam: Task dependent adaptive metric for improved

- few-shot learning. In *Advances in Neural Information Processing Systems*, pages 721–731, 2018.
- [29] Jose L Part and Oliver Lemon. Incremental online learning of objects for robots operating in real environments. In *2017 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, pages 304–310. IEEE, 2017.
- [30] Juan-Manuel Perez-Rua, Xi Tian Zhu, Timothy M. Hospedales, and Tao Xiang. Incremental few-shot object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [31] Amir Rosenfeld, Richard Zemel, and John K Tsotsos. The elephant in the room. *arXiv preprint arXiv:1808.03305*, 2018.
- [32] Eli Schwartz, Leonid Karlinsky, Rogerio Feris, Raja Giryes, and Alex M Bronstein. Baby steps towards few-shot learning with multiple semantics. *arXiv preprint arXiv:1906.01905*, 2019.
- [33] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- [34] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1199–1208, 2018.
- [35] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5103–5114, 2019.
- [36] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008.
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [38] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.
- [39] Xiaolong Wang, Yufei Ye, and Abhinav Gupta. Zero-shot recognition via semantic embeddings and knowledge graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6857–6866, 2018.
- [40] Xin Wang, Fisher Yu, Ruth Wang, Trevor Darrell, and Joseph E Gonzalez. Tafe-net: Task-aware feature embeddings for low shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1831–1840, 2019.
- [41] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. 2010.
- [42] Sanghyun Woo, Dahun Kim, Donghyeon Cho, and In So Kweon. Linknet: Relational embedding for scene graph. In *Advances in Neural Information Processing Systems*, pages 560–570, 2018.
- [43] Chen Xing, Negar Rostamzadeh, Boris N Oreshkin, and Pedro O Pinheiro. Adaptive cross-modal few-shot learning. *arXiv preprint arXiv:1902.07104*, 2019.
- [44] Eloi Zablocki, Patrick Bordes, Benjamin Piwowarski, Laure Soulier, and Patrick Gallinari. Context-Aware Zero-Shot Learning for Object Recognition. In *Thirty-sixth International Conference on Machine Learning (ICML)*, Long Beach, CA, United States, June 2019.
- [45] Hongguang Zhang, Jing Zhang, and Piotr Koniusz. Few-shot learning via saliency-guided hallucination of samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2770–2779, 2019.