

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository:<https://orca.cardiff.ac.uk/id/eprint/175190/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Yan, Kun, Bouraoui, Zied, Wei, Fangyun, Xu, Chang, Wang, Ping, Jameel, Shoaib and Schockaert, Steven 2025. Modelling multi-modal cross-interaction for multi-label few-shot image classification based on local feature selection. *ACM Transactions on Multimedia Computing, Communications and Applications* 10.1145/3711867

Publishers page: <https://doi.org/10.1145/3711867>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Modelling Multi-modal Cross-interaction for Multi-label Few-shot Image Classification Based on Local Feature Selection

KUN YAN, School of Computer Science, Peking University, China

ZIED BOURAOUI, CRIL - University of Artois & CNRS, France

FANGYUN WEI, School of Computer Science, University of Sydney, Australia

CHANG XU*, School of Computer Science, University of Sydney, Australia

PING WANG^{†*}, National Engineering Research Center for Software Engineering, Peking University, China

SHOAIB JAMEEL, Electronics and Computer Science, University of Southampton, UK

STEVEN SCHOCKAERT, School of Computer Science and Informatics, Cardiff University, UK

The aim of multi-label few-shot image classification (ML-FSIC) is to assign semantic labels to images, in settings where only a small number of training examples are available for each label. A key feature of the multi-label setting is that an image often has several labels, which typically refer to objects appearing in different regions of the image. When estimating label prototypes, in a metric-based setting, it is thus important to determine which regions are relevant for which labels, but the limited amount of training data and the noisy nature of local features make this highly challenging. As a solution, we propose a strategy in which label prototypes are gradually refined. First, we initialize the prototypes using word embeddings, which allows us to leverage prior knowledge about the meaning of the labels. Second, taking advantage of these initial prototypes, we then use a Loss Change Measurement (LCM) strategy to select the local features from the training images (i.e. the support set) that are most likely to be representative of a given label. Third, we construct the final prototype of the label by aggregating these representative local features using a multi-modal cross-interaction mechanism, which again relies on the initial word embedding-based prototypes. Experiments on COCO, PASCAL VOC, NUS-WIDE, and iMaterialist show that our model substantially improves the current state-of-the-art.

CCS Concepts: • **Computing methodologies** → **Computer vision**.

Additional Key Words and Phrases: Multi-label classification, few-shot learning, multi-modal interaction, prototypes, word embeddings.

*Corresponding authors.

[†]Also with the School of Software and Microelectronics, Peking University and the Key Laboratory of High Confidence Software Technologies (PKU), Ministry of Education.

Authors' Contact Information: Kun Yan, School of Computer Science, Peking University, Beijing, China, kyan2018@pku.edu.cn; Zied Bouraoui, CRIL - University of Artois & CNRS, France, zied.bouraoui@cril.fr; Fangyun Wei, School of Computer Science, University of Sydney, Australia, fwei8714@uni.sydney.edu.au; Chang Xu, School of Computer Science, University of Sydney, Australia, c.xu@sydney.edu.au; Ping Wang, National Engineering Research Center for Software Engineering, Peking University, Beijing, China, pwang@pku.edu.cn; Shoaib Jameel, Electronics and Computer Science, University of Southampton, UK, M.S.Jameel@southampton.ac.uk; Steven Schockaert, School of Computer Science and Informatics, Cardiff University, UK, schockaerts1@cardiff.ac.uk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 1557-735X/2018/8-ART111

<https://doi.org/XXXXXXX.XXXXXXX>

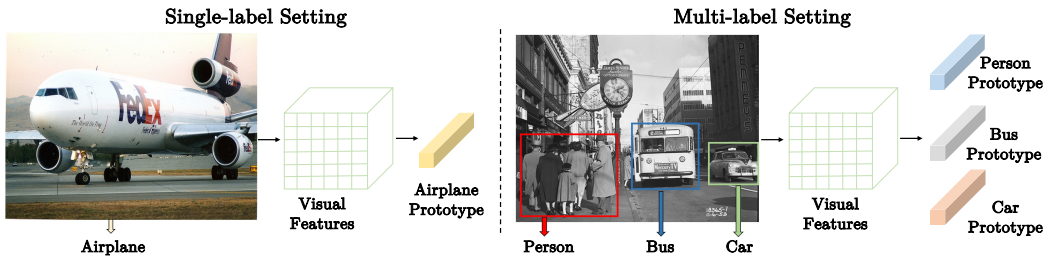


Fig. 1. In the 1-shot single-label setting, a given training image can be interpreted as a prototype for the considered label (left). In the multi-label setting, labels are related to different regions of the image, and these regions need to be identified before meaningful prototypes can be obtained (right).

ACM Reference Format:

Kun Yan, Zied Bouraoui, Fangyun Wei, Chang Xu, Ping Wang, Shoaib Jameel, and Steven Schockaert. 2018. Modelling Multi-modal Cross-interaction for Multi-label Few-shot Image Classification Based on Local Feature Selection. *J. ACM* 37, 4, Article 111 (August 2018), 28 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Multi-label image classification (ML-IC) has received considerable attention in recent years [9, 62, 65, 73]. This task aims to assign descriptive labels to images, where images can have multiple labels. Each label typically describes an object that is shown in the image. As a baseline strategy, we can simply train a binary classifier for each label, to predict the presence of a given type of object. To improve on this strategy, various approaches have been proposed to exploit label dependencies, e.g. taking advantage of the fact that the presence of one label makes the presence of another label more (or less) likely [6, 9, 62, 73]. In this paper, we are specifically interested in the few-shot setting, i.e. the case where only a limited number of training images are available for the labels (e.g., 1 or 5 images per label). Clearly, in this setting, we cannot rely on standard strategies for training binary image classifiers, nor on label co-occurrence statistics.

The challenge of few-shot image classification (FSIC) has garnered significant attention, particularly in the single-label (SL) setting [15, 20, 45, 53, 61, 70, 79]. Metric-based approaches hold a prominent position in SL-FSIC due to their strong performance and fast adaptability to new categories. However, it is worth noting that such methods cannot be directly extended to the multi-label setting. Metric-based approaches usually involve learning a prototype for each image category and subsequently assigning images to the category whose prototype is closest. These prototypes are typically derived by averaging a representation of the training images. In the seminal ProtoNet model [53], for instance, prototypes are simply defined as the average of the global feature maps of the available training examples. This strategy crucially relies on the assumption that most of the image is somehow relevant to its category. In the multi-label setting, such an assumption is highly questionable, given that different labels tend to refer to different parts of the image, as illustrated in Fig. 1. For instance, suppose the figure on the right is the only training example with the label *bus*. If we simply use the global feature map of this image as the corresponding prototype, the model is likely to mistakenly assume that other images of cityscapes should also be assigned this label. To obtain a high-quality prototype, it is thus imperative that we can estimate which regions of the image correspond to a given label.

In this paper, we propose a novel metric-based model for multi-label few-shot image classification (ML-FSIC). Given the aforementioned concerns, we need a strategy that is based on local image features, allowing us to focus on those parts of the training images that are most likely to be

relevant. However, as we may only have a single training example for some labels, we cannot implement such a strategy without some kind of prior knowledge about the meaning of the labels. Therefore, we rely on word embeddings, such as those from the GloVe model [46], to provide us with this prior knowledge. Some previous works for the single-label setting have already relied on word embeddings for inferring prototypes directly [67, 68], but as the resulting prototypes are inevitably noisy, they are normally used in combination with prototypes that are derived from visual features. Unfortunately, such visual prototypes are hard to obtain in the ML-FSIC setting.

For this reason, we adopt a different approach to efficiently leverage word embeddings for ML-FSIC. We first construct initial prototypes of the labels based on their word embeddings. Specifically, we introduce a cross-modality loss to learn a joint embedding space, in which the representations of words are aligned with the corresponding visual features. The resulting word representations are too noisy to be used directly. To address this, we propose a multi-modal cross-interaction strategy guided by these noisy prototypes. Concretely, we develop multiple ways of interaction between word embeddings and visual features, including cross-modal attention and cross-modal dynamic convolution, to refine the generation of the final prototype. As an important advantage, this strategy can be applied to previously unseen labels without the need for any fine-tuning of the model. We will refer to the model described above as our Base model.

As a further improvement of this base model, we propose a strategy for identifying the most representative local features prior to prototype generation. To achieve this, we propose a loss change measurement module (LCM), drawing inspiration from [76], to improve the selection of local features by assessing which features exert the greatest influence on the loss function. By filtering noisy local features based on this module, we end up with higher-quality prototypes, but the resulting model has the drawback that it needs to be fine-tuned before it can be applied to unseen labels. The above-described model will be referred to as the LCM model.

Besides, we propose a number of changes to the evaluation methodology for ML-FSIC systems. The most important change is concerned with how support sets are sampled, as part of an episode-based strategy. The standard N -way K -shot framework for evaluating FSIC systems is based on the idea that exactly K training images are available for each category of interest. While earlier work in ML-FSIC has aimed to mimic this N -way K -shot framework as closely as possible, we found this to have significant drawbacks when images can have multiple labels. We also propose some changes related to how the query set is sampled and the choice of evaluation metrics. Finally, we propose three new ML-FSIC datasets based on PASCAL VOC [19], NUS-WIDE [13] and iMaterialist [24], by adapting these ML-IC datasets to the few-shot setting.

The contributions of this work are summarized as follows:

- A novel multi-modal cross-interaction mechanism for prototype generation in ML-FSIC is proposed, featuring channel-wise cross-attention guided by word embeddings and word embedding-based dynamic convolution.
- A loss-change measurement module is introduced to identify representative local features, which helps in mitigating the interference of background features.
- New ML-FSIC benchmarks from NUS-WIDE and iMaterialist are introduced, and extensive experiments on a total of four datasets demonstrate the effectiveness and superiority of our approach over current state-of-the-art methods.

This paper is an extended version of our previous work [71], with several key updates. First, we enhance the overall approach by introducing a novel multi-modal cross-interaction mechanism, accompanied by updated experimental results. Second, we include results for the 5-shot setting and provide a more in-depth analysis. Third, we have introduced the Loss Change Measurement

(LCM) module, which consistently improves the results. Our implementation and associated data are available at <https://github.com/yk-pku/MMCI-ML-FSIC>.

2 RELATED WORK

In this section, we review the related work on *multi-label image classification* (ML-IC), *few-shot image classification* (FSIC), and the combined area of *multi-label few-shot image classification* (ML-FSIC). We also briefly discuss the impact of Vision-Language models on this field.

2.1 Multi-Label Image Classification

Early solutions for ML-IC simply learned a binary classifier for each label [59]. More recently, various methods have been proposed to improve this basic strategy by exploiting label dependencies in some way. For instance, the CNN-RNN architecture [62] learns a joint embedding space for representing both images and labels, which is used to predict image-label relevance. To avoid the need for a predefined label order, as in RNN-based architectures, [73] proposes Minimal Loss Alignment (MLA) and Predicted Label Alignment (PLA) to dynamically align the ground truth labels with the predicted label sequence. Some studies [9, 64, 75, 77, 86] also exploit Graph Convolutional Networks (GCN) [29] to model label dependencies more explicitly. Besides, attention mechanisms play an important role in ML-IC, as they can associate labels with specific image regions [65, 84, 85, 87, 88]. AAMN [85] proposes an attention-augmented memory network that mines contextual information from various categories within the dataset to enhance input feature representation. Recently, MMDSR [30] introduced an approach for dynamically constructing a semantic relationship graph using multi-scale feature reconstruction coupled with channel dual-branch cross-attention. However, the above methods require a large amount of training data and can thus not be directly applied in the few-shot setting.

2.2 Few-Shot Image Classification

Different strategies for single-label few-shot image classification have already been proposed, with metric-based [10, 28, 50, 56, 74] and meta-learning based [20, 39, 48, 49] methods being the most prominent. Methods of the latter kind use a meta-learner to learn aspects like model weight initialization [20] and update step size [39], allowing the adaption of model parameters to new categories in the few-shot regime. It allows adapting architectures to novel tasks with just a few steps of a gradient-based task optimizer. However, our method is more closely related to metric-based methods, which aim to learn a generalizable visual embedding space in which different image categories are spatially separated. For instance, ProtoNet [53] generates a visual prototype for each class by simply averaging the embeddings of the support images. The category of a query image is then determined by its Euclidean distance to these prototypes. Instead of using Euclidean distance, the Relation Network [56] learns to model the distance between query and support images. Other notable models include FEAT [74], which uses a transformer to contextualize the image features relative to the support set, and PSST [8], which introduced a self-supervised learning strategy.

While most metric-based models rely on global features, methods exploiting local features have also been proposed [36, 78], but these methods are designed for single-label classification. For instance, [36] calculates the similarity between all local features of the query image and all local features of the support images. As such, there is no attempt to focus on particular regions of the support images. The use of word vectors has also been considered for estimating visual prototypes [32, 67–69]. However, due to the inevitably noisy nature of the predicted prototypes, such methods are best used in combination with prototypes obtained from visual features.

2.3 Multi-Label Few-Shot Image Classification

The ML-FSIC problem has only received limited attention. LaSO [3] was the first model designed to address this problem. It relies on a data augmentation strategy that generates synthesized feature vectors via label-set operations. KGGR [7] uses a GCN to take label dependencies into account, where labels are modelled as nodes, and two nodes are connected if the corresponding labels tend to co-occur. The strength of these label dependencies is normally estimated from co-occurrence statistics, but for labels with limited training data, dependency strength is instead estimated based on GloVe word vectors [46]. In [52], several FSIC methods are extended to the multi-label setting, and a neural module is introduced to estimate the label count of a given sample by exploiting relational inference. BCR [4] is designed to uncover underlying correlations between instances and labels, incorporating varying levels of importance information. This approach analyzes correlations from both instance-to-label and label-to-instance perspectives. In [38], an ML-FSIC method is proposed which learns compositional embeddings based on weak supervision. However, this method is not directly comparable with our method due to its use of weak supervision.

Compared to existing methods such as LaSO [3], KGGR [7], MetaRE [52], and BCR [4], our work introduces a novel multi-modal cross-interaction mechanism tailored for prototype refinement. This mechanism is based on multiple modes of interaction between representative local features and word embeddings. Our study is the first to comprehensively explore the use of various word embeddings, including GloVe [46], BERT [16], and CLIP [47], specifically for multi-label few-shot image classification. Several other studies address multi-label image classification with limited labeled data in different ways, such as leveraging unlabeled data to boost model performance [37], using visual-language models like CLIP to classify images without annotations [1], and training classification models in the presence of ambiguous data [80, 81]. Nevertheless, our focus is on multi-label few-shot image classification, where only a small number of precisely labeled images (e.g., 1 or 5) are available for each label.

2.4 Vision-Language Model

Vision-Language (VL) models use self-supervision strategies to achieve a strong alignment of textual and visual features, demonstrating an impressive ability to learn generic visual representations [2, 23, 27, 33, 47, 66, 72]. For instance, CLIP [47] is trained with 400 million (image, text) pairs and shows strong zero-shot and transfer capabilities on over 30 visual tasks. With the rise of such powerful vision-language models, various studies have proposed approaches to efficiently adapt these models to downstream tasks, such as prompt learning [82, 83], and visual adaptation [22, 51]. Although these methods demonstrate remarkable generalization to various zero-shot or few-shot image classification tasks, due to the way they were trained, they primarily focus on classifying each image into a single label and thus are not well-suited for multi-label settings [55]. Recently, TaI-DPT [25] was proposed to extend CLIP to multi-label image classification by introducing text-as-image prompting. Although it achieves a 59.2 mAP score on the COCO dataset without requiring any task-specific training images, it requires collecting image captions and localized narratives from datasets for training. In contrast, with a single training image per category, our LCM model can achieve a 60.3 mAP score with an extremely shallow backbone Conv-4-64 [4, 52] and off-the-shelf standard word embedding GloVe [46]. In the area of multi-label few-shot image classification, our advantage over VL models comes from the ability to achieve promising performance with a small and thus highly efficient backbone network. Additionally, our approach is compatible with the backbone of VL models, offering the potential for even better performance.

3 PROBLEM SETTING

We consider the following multi-label few-shot image classification (ML-FSIC) setting: we are given a set of base labels C_{base} and a set of novel labels C_{novel} , where $C_{base} \cap C_{novel} = \emptyset$. We also have two sets of labelled images: \mathcal{E}_{base} , containing images with labels from C_{base} , and \mathcal{E}_{novel} , containing images with labels from C_{novel} , where $\mathcal{E}_{base} \cap \mathcal{E}_{novel} = \emptyset$. The images from \mathcal{E}_{base} are used for training the model, while those in \mathcal{E}_{novel} are used for testing. The goal of ML-FSIC is to obtain a model that performs well for the labels in C_{novel} , when given only a few examples of images that have these labels. Models are trained and evaluated using so-called episodes. Each training episode involves a support set and a query set. The support set corresponds to the examples that are available for learning to predict the labels in C_{base} , while the query set is used to assess how well the system has accomplished this goal.

To construct the support set of a given training episode, for every label in C_{base} , we sample K images from \mathcal{E}_{base} which have that label, with $K \in \{1, 5\}$. These images are sampled without replacement, meaning that the total number of images in the support set is given by $K \times |C_{base}|$. The query set is sampled in the same way, except that we sample 4 images per label. Testing episodes are constructed similarly, but with labels from C_{novel} and images from \mathcal{E}_{novel} instead.

Note that there are some differences between our sampling strategy and the standard methodology that is adopted for (single-label) FSIC. In particular, FSIC models are usually evaluated using episodes that contain a sub-sample of N classes, where the support set contains exactly K examples of each class. In ML-FSIC, it is difficult to guarantee that the support set contains exactly K examples of each label because most images have multiple labels. The sampling strategy from [3] nonetheless tries to ensure this, but we argue that doing this has two important drawbacks. First, adhering to the constraint that each label must occur K times often leads to a limited pool of image combinations available for constructing support sets. This limitation persists even when the requirement is somewhat relaxed. This implies that only a limited number of episodes can be sampled, which in turn makes training the model more challenging and results in less stable test outcomes. Second, the total number of images in the support set can vary substantially. For example, if one image contains all labels, then we may have a support set that only contains that one image when $K = 1$. In contrast, our strategy circumvents these drawbacks. We sample K images for each category without repetition, guaranteeing that at least K examples are accessible for each label. Additionally, the support set of each episode is consistently composed of the same number of images, namely $K \times N$. The idea of setting $N = |C_{base}|$ during training and $N = |C_{novel}|$ during testing, rather than selecting a sub-sample of N classes, conforms to the strategy that was used by [3] and [7]. Furthermore, our approach allows for a larger pool of available episodes. In particular, we sample 200 testing episodes, whereas LaSO only samples 10 testing episodes [3]. Moreover, our sampling strategy facilitates episode-based training, closely resembling the evaluation process employed for testing models.

4 METHOD

Our base model consists of two main components. The first component jointly represents label embeddings and visual features in a shared vector space, aiming to predict visual prototypes from the label embeddings. Since such prototypes are noisy, they are not used directly for final label predictions. Instead, this component is merely used to learn a joint representation of visual features and labels, as shown in Fig. 2. The second component is aimed at computing the final prototypes by the proposed multi-modal cross-interaction strategy, which mainly consists of a channel-wise cross-attention module and a word embedding-based dynamic convolution module.

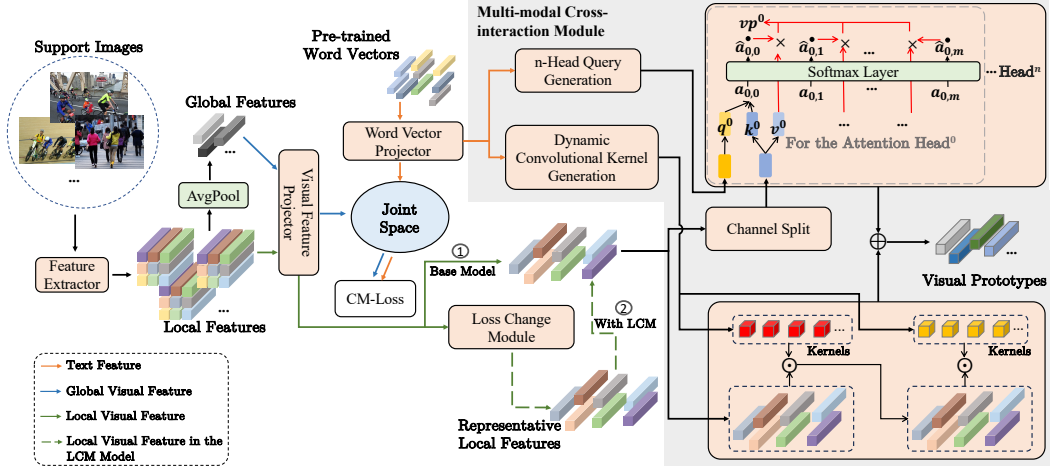


Fig. 2. Overview of our methodology. A joint embedding space is learned in which both labels and images are represented. A customized multi-modal cross-interaction strategy is proposed to calculate label prototypes using local features from the relevant images in the support set, along with word vectors that provide prior knowledge about the considered label. In our visual representation, the green solid line indicates the local feature flow for the base model, while the green dotted line represents the flow for the LCM model.

In addition to the base model, we also consider a variant, referred to as the LCM model, which involves an additional step between the two components. In particular, in this variant, we introduce Loss Change Measurement (LCM) to make a selection of local features. This selection is also based on the label embeddings from the first module. We then only aggregate the selected local features to construct the final prototypes. The LCM module thus makes a hard selection of the local features, aimed at removing those that are not relevant, before the attention mechanism is used to make a soft selection of the remaining local features.

It is noteworthy that we adopt distinct feature extractors for images and texts for two primary reasons. First, prevalent vision-language models such as CLIP [47], BLIP [33] and Flamingo [2] typically also employ separate extractors for images and texts due to their inherent modality distinctions. Second, in this way, our framework can accommodate a variety of image or text extractors, enhancing its adaptability. We now describe these different steps in more detail.

4.1 Joint Embedding of Visual Features and Labels

Given an input image I , we first use a feature extractor to obtain its local feature map $\mathbf{f}_{loc}^I \in \mathbb{R}^{n \times h \times w}$, where n is the number of channels, h is the height and w is the width. In this paper, we use a fully convolutional network such as ResNet [26] for this purpose. The global visual feature vector $\mathbf{f}_{glo}^I \in \mathbb{R}^n$ for image I is obtained by averaging the local feature map across the height and width of the image. In particular let us write $f_{i,j,k}$ for the element at position (i, j, k) in the tensor \mathbf{f}_{loc}^I . Then we define \mathbf{f}_{glo}^I as the vector whose i^{th} component is given by:

$$\mathbf{f}_i^* = \frac{1}{h \cdot w} \sum_{j=1}^h \sum_{k=1}^w f_{i,j,k}. \quad (1)$$

We use pre-trained word embeddings to represent the labels in $C_{base} \cup C_{novel}$. Let us write \mathbf{w}_c for the word vector representing label c , and let d_w be the dimensionality of the word vectors. To represent

images and labels in the same vector space, we learn two linear transformations:

$$\widehat{\mathbf{f}}_{glo}^I = \mathbf{A}_{visual} \mathbf{f}_{glo}^I \quad \widehat{\mathbf{w}}_c = \mathbf{A}_{text} \mathbf{w}_c, \quad (2)$$

where $\mathbf{A}_{visual} \in \mathbb{R}^{d_j \times n}$ is used to project the global feature vector for I onto a space of d_j dimensions. Similarly, $\mathbf{A}_{text} \in \mathbb{R}^{d_j \times d_w}$ is used to project the d_w -dimensional embedding of a label c onto the same d_j -dimensional space. To ensure that the resulting image vectors $\widehat{\mathbf{f}}_{glo}^I$ and label representations $\widehat{\mathbf{w}}_c$ are semantically compatible, we propose to use the following loss, which treats the vectors $\widehat{\mathbf{w}}_c$ as prototypes to classify images:

$$\mathcal{L}_{cm} = - \sum_{I \in \mathcal{S}} \sum_{i=1}^{|C|} y_i^I \cdot \log \sigma(s_i^I) + (1 - y_i^I) \cdot \log(1 - \sigma(s_i^I)), \quad (3)$$

where \mathcal{S} represents the set of images from the support set of the current training episode, $C = \{c_1, \dots, c_{|C|}\}$ is the set of labels, $\sigma(\cdot)$ is the sigmoid function and y_i^I represents the ground truth, i.e. $y_i^I = 1$ if image I has label c_i and $y_i^I = 0$ otherwise. Finally, we have

$$s_i^I = \lambda \cos(\widehat{\mathbf{f}}_{glo}^I, \widehat{\mathbf{w}}_{c_i}), \quad (4)$$

where \cos denotes the cosine similarity and the scalar λ is a hyper-parameter to address the fact that the cosine similarity is bounded between -1 and 1. Note that we will not be using the vectors $\widehat{\mathbf{w}}_c$ as prototypes for the final classification. We only use this classification loss to align the two modalities (word vectors and visual features). For this reason, we will refer to Eq. (3) as the Cross-Modality (CM) loss.

It should be noted that the proposed alignment process is conceptually straightforward, being based on two standard linear transformations. While more intricate methods may be conceived, we found this approach to achieve satisfactory results. Moreover, linear mappings have proven surprisingly capable for the problem of aligning embedding spaces, often outperforming deep learning based approaches [43].

4.2 Identifying Representative Local Features using Loss Change Measurement

In our base model, the global image embedding \mathbf{f}_{glo}^I is obtained by averaging the local features across all regions of the image. However, not all these regions are equally important. Ideally, we only want to include those regions that contain the object being described by a given label, omitting local features that describe the image background. In this section, we consider the optional Loss Change Measurement (LCM) module to make a hard selection, aiming to remove those local features that are not relevant.

Let us write (j, k) for the position of some local feature, i.e. the local feature corresponding to the region on the j^{th} row and k^{th} column. We now introduce a parameter $\rho_{j,k} \in [0, 1]$ to capture the importance of the local feature at position (j, k) of a given image. To learn this parameter, we use the CM loss Eq. (3), with two changes. First, the remaining parameters (i.e. the matrices \mathbf{A}_{visual} and \mathbf{A}_{text}) are frozen during this step. Second, when computing the global image embedding \mathbf{f}_{glo}^I , we first multiply each local feature by its corresponding importance $\rho_{j,k}$. In particular, the i^{th} component of the vector \mathbf{f}_{glo}^I is now defined as follows:

$$f_i^* = \frac{1}{h \cdot w} \sum_{j=1}^h \sum_{k=1}^w \rho_{j,k} f_{i,j,k}. \quad (5)$$

Essentially, the importance weight $\rho_{j,k}$ reflects how compatible the local feature at position (j, k) is with the relevant label embeddings (i.e. the embeddings of the labels assigned to the image). After each iteration, the importance weights are normalised, to ensure that $\rho_{j,k} \in [0, 1]$. In particular, we update each $\rho_{j,k}$ as follows:

$$\rho_{j,k} := \frac{\rho - \rho_{min}}{\rho_{max} - \rho_{min}}, \quad (6)$$

with

$$\rho_{min} = \min_{1 \leq j \leq h} (\min_{1 \leq k \leq w} (\rho_{j,k})), \quad \rho_{max} = \max_{1 \leq j \leq h} (\max_{1 \leq k \leq w} (\rho_{j,k})) \quad (7)$$

If we end up with $\rho_{j,k} = 0$, this would mean that the local feature (j, k) is permanently removed. To prevent this from happening, in such a case, we instead set the value of $\rho_{j,k}$ as the lowest non-zero value of the importance weight $\rho_{j',k'}$, across all local features.

Rather than directly using the weights $\rho_{j,k}$ to make the selection of local features, taking inspiration from [70, 76], we measure the impact on the classification loss when a given feature is (partially) removed. Intuitively, if removing a given local feature has a large impact on the classification loss, it is likely to be representative of the object we are trying to model (i.e. the object described by the target label). In contrast, if the impact is minimal, it is more likely that this feature describes the image background or a different object. The impact of removing the local feature (j, k) on the loss function is computed as follows:

$$g(j, k) = |\mathcal{L}_{\Omega}^{(j,k)}(\rho_{j,k}) - \mathcal{L}_{\Omega}^{(j,k)}(0)|, \quad (8)$$

where \mathcal{L}_{Ω} represents a classification loss with parameters Ω . We again use the CM loss \mathcal{L}_{cm} as the classification loss, where we assume that \mathbf{f}_{glo}^l is computed based on the importance weights, according to Eq. (5). The argument of $\mathcal{L}_{\Omega}^{(j,k)}$ represents the importance weight of the local feature at position (j, k) . Note that $\mathcal{L}_{\Omega}^{(j,k)}(0)$ thus represents the value of the loss function we obtain after setting $\rho_{j,k}$ to 0, i.e. if we completely disregard the region of the image at position (j, k) . In other words, $g(j, k)$ captures the difference in the loss function between (i) using $\rho_{j,k}$ as the importance of the local feature at position (j, k) and (ii) setting this weight to 0. Intuitively, if $g(j, k)$ is above some threshold (for the value of $\rho_{j,k}$ upon convergence), we keep the local feature at position (j, k) ; otherwise, we remove it. However, in practice, we do not use $g(j, k)$ directly. In particular, (i) we use an approximation of $g(j, k)$ to make the computation more efficient and (ii) rather than only evaluating $g(j, k)$ after the value of $\rho_{j,k}$ has converged, we use a momentum based strategy to improve stability. We now explain these two steps in more detail.

To simplify the implementation, we take advantage of the Taylor approximation of $\mathcal{L}_{\Omega}^{j,k}(x)$ at x_0 :

$$\mathcal{L}_{\Omega}^{j,k}(x) = \mathcal{L}_{\Omega}^{j,k}(x_0) + \frac{\mathcal{L}_{\Omega}^{j,k(1)}(x_0)}{1!}(x - x_0) + \dots + \frac{\mathcal{L}_{\Omega}^{j,k(n)}(x_0)}{n!}(x - x_0)^n + R_n(x), \quad (9)$$

where we write $\mathcal{L}_{\Omega}^{j,k(i)}$ for the i^{th} derivative of $\mathcal{L}_{\Omega}^{j,k}$. We find that $\mathcal{L}_{\Omega}^{j,k}(0)$ can be estimated as $\mathcal{L}_{\Omega}^{j,k}(\rho) - \rho \mathcal{L}_{\Omega}^{j,k(1)}(\rho)$ by setting $x = 0$ and $x_0 = \rho$. Therefore, Eq. 8 can be estimated as:

$$g(j, k) \approx |\rho_{j,k} \mathcal{L}_{\Omega}^{j,k(1)}(\rho_{j,k})|. \quad (10)$$

Note that the derivatives $\mathcal{L}_{\Omega}^{j,k(1)}(\rho_{j,k})$ can be computed in parallel for all local features, which makes using this approximation much faster than evaluating Eq. (8) for every local feature.

As already mentioned, the importance weights $\rho_{j,k}$ are iteratively learned by optimising the CM loss. These weights are initialized as $\rho_{j,k} = 1$. Rather than only using the converged values of

$\rho_{j,k}$, we use a momentum-based strategy to improve the stability of the estimation. In particular, drawing on the experience of [58], we compute the following estimate in iteration i :

$$f_i(j, k) = \alpha_i f_{i-1}(j, k) + (1 - \alpha_i) g_i(j, k), \quad \alpha_i = \min\left(1 - \frac{1}{(i+1)}, \alpha\right), \quad (11)$$

where we set $f_0(\rho) = 0$ and $\alpha = 0.95$ following [58]. The value of $g_i(j, k)$ is computed according to Eq. (10), using the importance weights $\rho_{j,k}$ that were obtained in the previous iteration.

After the final iteration n , we make the following decision: the local feature at position (j, k) is retained if $\sigma(f_n(j, k)) \geq \theta$, where $\sigma(\cdot)$ denotes the sigmoid function. Otherwise, this local feature is removed.

4.3 Multi-modal Cross-interaction for Prototype Construction

We now explain how the final label prototypes are constructed. The multi-modal cross-interaction module, depicted in Fig. 2, allows two distinct types of information interaction between visual and text features. The first is to aggregate various visual patterns using word vectors as *query* within a customized cross-attention mechanism. The second is to employ word embeddings to generate convolutional kernels that interact with image features.

4.3.1 Channel-wise Cross-attention Guided by Word Embedding. As pointed out in [11], features learned from different channels correspond to distinct visual patterns. Utilizing the word embedding-based prototypes from Sec. 4.1 as prior knowledge, we extract various visual patterns from local features. Concretely, consider a label c and suppose this label has been assigned to m images from the support set. We can derive a total of $l = h \cdot w \cdot m$ local feature vectors from these m images. Let us write these local feature vectors as $\mathbf{u}_1, \dots, \mathbf{u}_l$. These feature vectors are first projected into the joint embedding space: $\widehat{\mathbf{u}}_i = \mathbf{A}_{\text{visual}} \mathbf{u}_i$. Subsequently, the feature vectors are split evenly along the channel dimension to serve as *value* and *key* inputs for n_a attention heads within the attention mechanism [60]. The corresponding label embedding $\widehat{\mathbf{w}}_c$ provides the *query* for each attention head through a fully connected (FC) layer. Specifically, we have:

$$\mathbf{q}_c^j = \mathbf{Q}_j \widehat{\mathbf{w}}_c, \quad \mathbf{k}_i^j = \mathbf{v}_i^j = \widehat{\mathbf{u}}_i[(j-1) \cdot d_a : j \cdot d_a], \quad (12)$$

$$\mathbf{p}_c^j = \sum_{i=1}^l \mu_i^j \mathbf{v}_i^j, \quad (\mu_1^j, \dots, \mu_l^j) = \text{Softmax}\left(\frac{\mathbf{q}_c^j \cdot \mathbf{k}_1^j}{\sqrt{d_a}}, \dots, \frac{\mathbf{q}_c^j \cdot \mathbf{k}_l^j}{\sqrt{d_a}}\right) \quad (13)$$

where $d_a = \frac{d_j}{n_a}$ is the dimensionality of the vectors \mathbf{q}_c^j , \mathbf{k}_i^j , and \mathbf{v}_i^j . The vector \mathbf{p}_c^j represents the contribution of the j^{th} attention head to the prototype of label c . As there are n_a attention heads in total, the prototype of label c in this cross-attention mechanism is given by:

$$\mathbf{p}_c^{\text{att}} = \text{MLP}(\mathbf{p}_c^1 \oplus \dots \oplus \mathbf{p}_c^{n_a}), \quad (14)$$

where we write \oplus for vector concatenation and MLP consists of two linear layers with GeLU activation and dropout.

4.3.2 Word Embedding-based Dynamic Convolution. Inspired by the use of dynamic convolution in [54] to enhance object features, we propose generating convolutional kernels from word embeddings to interact with image features. Specifically, we begin by selecting the top n_d local features from the joint embedding space, denoted as $\widehat{\mathbf{u}}_1, \dots, \widehat{\mathbf{u}}_{n_d}$, based on their cosine similarities to $\widehat{\mathbf{w}}_c$. We then utilize linear layers to produce two sets of convolutional kernels from $\widehat{\mathbf{w}}_c$, i.e. $\Theta_c^1 \in \mathbb{R}^{1 \times 1 \times d_j \times d_c}$ and $\Theta_c^2 \in \mathbb{R}^{1 \times 1 \times d_c \times d_j}$. Finally, the operation of word embedding-based dynamic convolution is

defined as :

$$\bar{\mathbf{u}}_i = \text{ReLU}(\text{Norm}(\Theta_c^1 \widehat{\mathbf{u}}_i)), \quad (15)$$

$$\mathbf{p}_c^{dyn} = \frac{1}{n_d} \sum_{i=1}^{n_d} \text{ReLU}(\text{Norm}(\Theta_c^2 \bar{\mathbf{u}}_i)), \quad (16)$$

where ReLU and Norm denote the ReLU activate function and Layer Normalization [54], respectively. Finally, the prototype of category c is obtained as:

$$\mathbf{p}_c = \mathbf{p}_c^{att} + \mathbf{p}_c^{dyn}. \quad (17)$$

4.3.3 Loss Function. To train our multi-modal cross-interaction module, we use the following loss function, which we refer to as the Query Loss:

$$\mathcal{L}_{query} = - \sum_{I \in \mathcal{Q}} \sum_{i=1}^{|\mathcal{C}|} y_i^I \cdot \log \sigma(q_i^I) + (1 - y_i^I) \cdot \log(1 - \sigma(q_i^I)), \quad (18)$$

where \mathcal{Q} represents the set of images from the query set of the current training episode. As before, \mathcal{C} is the set of labels and y_i^I represents the ground truth. The predictions q_i are obtained as follows:

$$q_i^I = \lambda \cos(\widehat{\mathbf{f}}_{glo}^I, \mathbf{p}_{c_i}), \quad (19)$$

with λ the same scalar as in Eq. (4). Here we assume that the global feature vector $\widehat{\mathbf{f}}_{glo}^I$ is computed as a standard average, according to Eq. (1), as we are not able to compute importance weights for query images. As an alternative to using the global feature vector $\widehat{\mathbf{f}}_{glo}^I$ we could also compare \mathbf{p}_{c_i} to the local image features. However, as we will see in our analysis in Section 5, this variant does not lead to clear improvements while having a considerably higher computational cost. Note that there are two key differences between \mathcal{L}_{cm} and \mathcal{L}_{query} : (i) \mathcal{L}_{cm} is trained using the support images while \mathcal{L}_{query} is trained using the query images; and (ii) prototypes in \mathcal{L}_{cm} are estimated from word vectors while prototypes in \mathcal{L}_{query} are those obtained by aggregating local visual features.

4.4 Model Training and Evaluation

The model is trained by repeatedly sampling training episodes from \mathcal{C}_{base} , as explained in Section 3. Given a training episode with support set \mathcal{S} and query set \mathcal{Q} , the model parameters are updated using the following loss:

$$\mathcal{L}_{all} = \mathcal{L}_{cm} + \gamma \mathcal{L}_{query}, \quad (20)$$

with γ a hyperparameter to control the relative importance of both components.

After the model has been trained, it can be evaluated on test episodes as follows. For the base model, without the LCM step, we first construct the prototypes using the support set, as in Eq. (14). Note that we can do this without fine-tuning any model parameters. For each query image I , the probability that it has label c_i is computed as $\sigma(q_i^I)$ with q_i^I as defined in Eq. (19). For the variant with LCM, we first need to use the Loss Measurement Module to identify the most representative local features in the images in the support set. The labels are then predicted in the same way as for the base model.

5 EXPERIMENTS

5.1 Experimental Setup

5.1.1 Datasets. We have conducted experiments on four datasets, starting with COCO [40]. This dataset was already used for evaluating LaSO in [3], where a split into 64 training labels and 16 test labels was proposed. However, as this split did not include a validation set, we split their 64 training labels into 12 labels for validation (*cow, dining table, zebra, sandwich, bear, toaster, person, laptop, bed, teddy bear, baseball bat, skis*) and 52 labels for training, while keeping the same 16 labels for testing. We follow [3] to include images from the COCO 2014 training and validation sets. The images which do not contain any of the test and validation labels are used as the training set. Similarly, the validation set only contains images that do not contain any training or test labels. Second, we propose a new ML-FSIC dataset based on PASCAL VOC [19], which has 20 labels. To use as many images as possible, we select the following six labels for the novel set C_{novel} : *dog, sofa, cat, potted plant, tv monitor, sheep*. The following six labels were selected for the validation split: *boat, cow, train, aeroplane, bus, bird*. The remaining eight labels are used for training. We use the images from the VOC 2007 training, validation, and test splits, as well as the VOC 2012 training and validation splits (noting that the labels of the VOC 2012 test split are not publicly available). Third, we propose a new ML-FSIC dataset based on NUS-WIDE [12]. NUS-WIDE contains 161789 images for training and 107859 images for testing, which were manually annotated with 81 concepts. In our setting, we use both training and testing images. We have selected 17 labels for the novel set (*bear, cityscape, statue, bridge, sunset, sign, coral, fish, fire, cars, sand, vehicle, police, temple, horses, reflection, street*) and 12 labels for the validation set (*tower, tree, food, sun, town, beach, military, birds, flowers, house, lake, garden*). The training set consists of the remaining labels. The complete NUS-WIDE dataset is no longer available from the official website¹, as many image download links have expired. Therefore, we have instead collected the images from the NUS-WIDE dataset from a Kaggle website². Lastly, we introduce a new ML-FSIC dataset based on iMaterialist [24]. Following a similar methodology as MoFSOD [31], we extract 46 classes. Out of these, 16 labels (*sock, jumpsuit, glove, scarf, coat, belt, zipper, cardigan, jacket, hood, vest, flower, tie, collar, ruffle, shirt_blouse*) are selected for the novel set, while 12 labels (*watch, lapel, shorts, bead, glasses, hat, sweater, tights_stockings, buckle, skirt, applique, rivet*) are allocated to the validation set. The remaining labels constitute the training set. An overview of the labels contained in the four datasets can be found in the Appendix.

5.1.2 Methodology. Every model is trained for 200 epochs, with the first 10 epochs used as a warm-up. We used the Adam optimizer with an initial learning rate of 0.001. The variant with LCM requires a training step during the test phase, to learn the importance weights of the local features of the support images. We set the number of epochs for this training step to 20. For baseline methods that require a training or fine-tuning step, we set the number of epochs to 40 (noting that 20 epochs are not sufficient for some of these methods to converge). Note that our base model is used without any training or fine-tuning. During the test phase, we sample 200 test episodes. The results are averaged over all test episodes. We report results in terms of macro and micro metrics. The macro metric is evaluated by averaging per-class predictions while the micro metric is an overall measure for all images over all labels. Specifically, we use macro/micro average precision (Ma-AP/Mi-AP) and macro/micro F1-measure (Ma-F1/Mi-F1), following earlier work in ML-IC [87]. Ma-AP is also denoted as mAP in many works [3, 14, 18]. We assume that a label is predicted as positive if its estimated probability is greater than 0.5 following the standard setting [87].

¹<https://lms.comp.nus.edu.sg/wp-content/uploads/2019/research/nuswide/NUS-WIDE.html>

²<https://www.kaggle.com/datasets/xinleili/nuswide?resource=downloadx>

Table 1. Overview of the main results on COCO, in terms of micro/macro AP and F1 scores (in %). The best results are highlighted in bold font. “-” indicates that no word embeddings are used.

	Word-Emb.	1-shot				5-shot			
		Mi-AP	Mi-F1	Ma-AP	Ma-F1	Mi-AP	Mi-F1	Ma-AP	Ma-F1
ResNet-50	-	11.31	12.51	19.67	13.27	19.12	16.08	24.17	17.22
ResNet-101	-	10.22	12.04	18.87	12.18	17.93	16.32	21.67	15.74
ViT	-	10.07	11.02	16.94	11.13	24.73	21.10	28.78	20.93
ResNet-50 + ViT	-	10.13	10.21	16.71	11.17	23.65	20.41	26.65	19.56
PLA	Self-learned	21.33	20.89	30.61	20.96	27.87	29.91	36.38	28.25
PLA (GloVe)	GloVe	23.02	22.15	30.13	20.37	27.11	28.48	35.13	27.47
LaSO	-	16.83	15.09	25.41	16.11	22.24	20.14	28.41	21.59
MAML	-	25.42	23.11	35.30	23.82	28.30	30.50	37.89	29.13
WGA	GloVe	35.30	34.21	42.84	28.91	38.97	38.10	44.91	33.12
Base model	GloVe	39.68	38.24	47.02	32.94	42.84	42.17	48.85	37.47
LCM model	GloVe	41.02	38.67	48.92	33.55	46.35	44.86	52.79	39.01
Base model	mirrorBERT	41.79	39.42	49.49	34.46	45.15	43.26	51.49	39.54
LCM model	mirrorBERT	43.20	40.39	51.55	35.55	47.72	45.81	54.98	42.04

5.1.3 Implementation Details. We use ResNet-50 [26], ViT [17] and GoogleNet-v3 [57] as feature extractors. The dimensionality of the joint embedding space was set as $d_j = 512$. As word embeddings, we mainly consider 300-dimensional GloVe [46] vectors³ as this is the most common choice in the literature for few-shot image classification [32, 67] and multi-label image classification [7, 9]. In addition, we also experimented with the standard pre-trained FastText and Word2Vec embeddings from an online repository⁴. Beyond standard word embeddings, we also experimented with embeddings that were extracted from BERT [16] using a number of different strategies. In particular, we use mirrorBERT [41] and the biencoder model from [21] (referred to as BERT_{biencoder}), both of which are fine-tuned BERT models that can map any word or phrase onto a vector. In addition, we also use mirrorWiC [42], and the strategies from [34] (referred to as BERT_{avg}) and [35] (referred to as BERT_{concn}), all of which learn embeddings of words in context. To use these strategies to learn the embedding of a given word, we randomly select 500 sentences that mention the word from Wikipedia. Then we obtain the contextualised representations of these word mentions, using the aforementioned models, and finally compute the average of these vectors across the 500 sentences. Finally, we also consider word vectors extracted from CLIP [47], using “a photo of a [CLASS]” as the prompt to guide its text encoder. Based on the validation split, for both datasets, the number of attention heads was set to 8, the θ in LCM was set to 0.65, and γ was set to 1. We present our results for two variants: the base model and the LCM model. Note that an important advantage of the base model is its ability to classify novel categories without requiring any parameter adjustments. On the other hand, using the LCM module should lead to better results.

5.2 Baselines

We compare with LaSO [3], KGGR [7], NLC [52], BCR [4], which were designed for the ML-FSIC setting. To put the results in context, we also include comparisons with methods tailored for ML-IC. First, following LaSO [3], we attach a standard multi-label classifier to a number of different feature extractors: *ResNet-50*, *ResNet-101*, *ViT-Base* [17] and *ResNet-50 + ViT-Base*. Second, we report

³Specifically, we use vectors that were trained from Wikipedia 2014 and Gigaword 5, which we obtained from the GloVe project page, at <https://nlp.stanford.edu/projects/glove>.

⁴<https://developer.syn.co.in/tutorial/bot/oscovova/pretrained-vectors.html>

Table 2. Overview of the main results on COCO, in terms of micro/macro AP and F1 scores (in %), following the same data split as in FsPML [81]. The best results are highlighted in bold font. “-” indicates that no word embeddings are used.

	Word-Emb.	1-shot				5-shot			
		Mi-AP	Mi-F1	Ma-AP	Ma-F1	Mi-AP	Mi-F1	Ma-AP	Ma-F1
ResNet-50	-	18.43	18.38	25.03	18.92	26.53	24.62	32.91	25.33
ResNet-101	-	15.85	17.93	24.61	18.77	23.79	23.38	28.37	22.98
ViT	-	14.53	16.29	21.00	17.36	30.46	26.77	35.83	26.41
ResNet-50 + ViT	-	14.44	15.18	22.56	18.28	31.02	26.58	34.61	26.22
PLA	Self-learned	29.34	29.01	37.68	28.30	35.84	37.17	43.75	36.71
PLA (GloVe)	GloVe	30.44	29.16	37.81	27.56	35.06	36.61	42.14	35.02
LaSO	-	24.75	22.21	33.02	23.81	29.59	27.62	35.47	29.35
MAML	-	33.02	29.37	42.03	29.84	35.91	37.72	44.78	36.05
WGA	GloVe	39.26	37.43	47.25	33.60	42.71	43.04	49.55	39.01
Base model	GloVe	46.44	44.82	52.86	39.21	49.68	48.76	55.21	44.75
LCM model	GloVe	47.37	45.17	54.78	40.57	52.99	50.76	58.72	45.92
Base model	mirrorBERT	48.81	45.60	55.94	40.83	51.08	49.43	57.90	45.43
LCM model	mirrorBERT	50.06	47.35	56.75	41.02	53.92	51.84	60.19	47.85

Table 3. Overview of the main results on NUS-WIDE, in terms of micro/macro AP and F1 scores (in %), following the same data split as in FsPML [81]. The best results are highlighted in bold font. “-” indicates that no word embeddings are used.

	Word-Emb.	1-shot				5-shot			
		Mi-AP	Mi-F1	Ma-AP	Ma-F1	Mi-AP	Mi-F1	Ma-AP	Ma-F1
ResNet-50	-	18.76	23.98	22.04	22.91	24.53	28.21	27.53	28.69
ResNet-101	-	16.55	22.76	20.20	21.31	24.91	27.11	26.43	27.81
ViT	-	20.19	23.97	22.44	24.29	27.40	30.29	29.45	28.67
ResNet-50 + ViT	-	19.26	20.45	20.35	18.41	28.86	30.48	30.24	28.79
PLA	Self-learned	26.14	28.09	32.40	27.86	33.72	34.31	38.85	33.15
PLA (GloVe)	GloVe	25.53	27.42	32.01	26.54	32.92	32.69	38.43	33.57
LaSO	-	22.82	26.27	29.70	26.03	28.62	31.87	33.74	31.54
MAML	-	28.32	30.02	33.76	27.95	36.35	35.43	39.31	33.82
WGA	GloVe	36.56	34.39	41.70	33.35	40.87	39.04	45.34	40.62
Base model	GloVe	41.18	38.48	45.26	38.56	45.29	43.32	50.30	44.14
LCM model	GloVe	42.75	39.16	46.78	40.48	48.25	45.47	54.58	46.23
Base model	mirrorBERT	42.01	38.64	46.17	38.68	47.67	44.70	52.74	46.07
LCM model	mirrorBERT	43.32	40.15	48.13	41.05	48.74	47.13	55.56	47.58

results for the recent state-of-the-art CNN-RNN based method PLA [73], which, to the best of our knowledge, has not previously been evaluated in the ML-FSIC setting. Furthermore, we include the seminal meta-learning method MAML [20] designed for FSIC. We also compare our results with our previous work, WGA [71]. It is worth noting that the base model in this work is an enhanced version of WGA, incorporating additional multi-modal interaction operations, as detailed in Sec. 4.3. Since we did not have access to the source code of KGGR, NLC, and BCR, we only compared our method against the published results from the original paper, which followed the experimental setting from [3]. For PLA, LaSO, and our method, we used *ResNet-50* as the feature extractor in the

Table 4. Experimental results (mAP in %) on COCO, when following the data and evaluation setting from LaSO [3]. The best results are indicated by being highlighted in bold font; “-” indicates no word embeddings are used.

Method	Backbone	Word Emb.	1-shot	5-shot
LaSO [3]	GoogleNet-v3	-	45.3	58.1
KGGR [7]	GoogleNet-v3	GloVe	49.4	61.0
ProtoNets + NLC [52]	Conv-4-64	-	50.2	60.4
RelationNets + NLC [52]	Conv-4-64	-	53.3	60.8
LPN + NLC [52]	Conv-4-64	-	56.8	64.8
BCR [4]	GoogleNet-v3	-	61.5	70.1
Base model	Conv-4-64	GloVe	59.8	66.7
LCM model			60.3	68.3
Base model	GoogleNet-v3	GloVe	62.9	71.7
LCM model			64.2	73.9

main experiments. However, to ensure a fair comparison, we adopt the same feature extractors used in KGGR, LaSO, NLC, and BCR, for comparing our method in the LaSO setting.

5.3 Experimental Results

The experimental results for COCO are shown in Tab. 1 for the 1-shot and 5-shot settings. The results show that our proposed method outperforms the other methods by a substantial margin. From Tab. 1, it is evident that large models are prone to overfitting, noting that in terms of model size, we have: *ResNet-50 + ViT-Base* > *ViT-Base* > *ResNet-101* > *ResNet-50*. This is not unexpected given the small number of labeled examples in the 1-shot setting. However, in the 5-shot setting, where more training examples are available, ViT-based models can outperform ResNet-based models. LaSO can improve the ResNet-50 baseline because of its data augmentation strategy. Somewhat surprisingly, PLA performs better than LaSO, which shows that its LSTM component can model label dependencies in a meaningful way. PLA also needs word embeddings, but in the original model, these embeddings are learned from the training data itself. For comparison, we also report results for a variant where these word vectors are initialized using GloVe vectors instead; the results are shown as PLA (GloVe). As can be seen, using pre-trained word vectors does not lead to a meaningful improvement over the original PLA model and can even degrade the performance on macro metrics, as shown in Tab. 1. Interestingly, MAML [20], which was designed for the FSIC setting, performs better than the conventional ML-IC approaches, both in the 1-shot and 5-shot settings. The results also show that adding the LCM module improves the result, with the gains being somewhat clearer in the 5-shot setting compared to the 1-shot setting. Please refer to the Appendix for the experimental results for PASCAL VOC, NUS-WIDE, and iMaterialist. To provide a more thorough evaluation, we conducted additional experiments on both COCO and NUS-WIDE datasets using the data split proposed in FsPML [81]. As shown in Tab. 2 and Tab. 3, we can again see that our proposed method achieves the best performance.

To facilitate a comparison with KGGR, NLC, and BCR, we adopt the 1-shot and 5-shot settings from [3]. For each method, we utilize the results reported in their respective original papers. As shown in Tab. 4, our method significantly outperforms existing approaches. Notably, our method achieves superior performance even with a simple and shallow Conv-4-64 backbone, whereas both LaSO and KGGR rely on the deeper GoogleNet-v3 backbone. Furthermore, when equipped with GoogleNet-v3, our approach establishes a new state-of-the-art.

Table 5. Ablation study (in %) on components in our method. ‘Cro. Att.’ denotes the channel-wise cross-attention guided by word embedding. ‘Dyn. Conv.’ represents the word-embedding-based dynamic convolution. The best results are highlighted in bold.

CM-loss	Cro. Att.	Dyn. Conv.	LCM	1-shot	5-shot	Notes
	✓	✓		44.01	46.59	Base model w/o CM-loss
✓		✓		44.92	46.85	Base model w/o Cro. Att.
✓	✓			45.67	47.03	Base model w/o Dyn. Conv.
✓	✓	✓		47.02	48.85	Base model
✓	✓	✓	✓	48.92	52.79	LCM model

Table 6. Results (in %) on COCO for using “vision and language” model CLIP. The best results are highlighted in bold.

Feature Ext.	Word Emb.	CM Loss	1-shot	5-shot
ViT-B/32 _{CLIP}	Word _{CLIP}		53.17	56.36
		✓	55.88	60.59

5.4 Ablation Study

Here we analyze the importance of the main components of our model. All experiments in this section are conducted on the COCO dataset, employing the *ResNet-50* feature extractor and GloVe vectors, unless stated otherwise. The results are reported by default in terms of Ma-AP, commonly referred to as mAP [3, 14, 18]. For additional ablation study results and qualitative analysis, please refer to the Appendix.

5.4.1 Effectiveness of Each Component. We ablate the effectiveness of each component in our method as shown in Tab. 5. When the component of the CM loss is removed, we only use \mathcal{L}_{query} for training. Note that the alignment between the word vectors and visual features then has to be learned indirectly, together with the parameters of the overall model. As indicated in the first row of Tab. 5, the absence of the CM loss significantly diminishes performance. The second to fourth rows illustrate that removing any multi-modal interaction module also leads to a reduction in performance. Conversely, the inclusion of LCM, as shown in the last row, markedly enhances performance, particularly in the 5-shot setting. One may wonder whether the CM loss is still needed when using a vision-and-language model such as CLIP, given that such models have text and visual spaces which are inherently aligned. We analyse this in Table 6, where we show results on the base model that were obtained with the CLIP (ViT-B/32_{CLIP}) image encoder together with the word embeddings from the CLIP text encoder (Word_{CLIP}). As can be seen, even in this case, using the CM loss leads to substantial performance gains.

5.4.2 Different Feature Extractors. For the main experiments, we have primarily used ResNet-50 [26] as the feature extractor (backbone) for convenient comparison with other methods. Tab. 7 shows the impact of using different feature extractors. In addition to the standard ResNet-50, we also explore backbones from foundation models. Specifically, we leverage ResNet-50 and ViT-B/32 from the CLIP[47] image encoder (denoted as ResNet-50_{CLIP} and ViT-B/32_{CLIP} respectively). To adapt them to our approach, we remove the final attention layer and keep all grid features for ResNet-50_{CLIP} and ViT-B/32_{CLIP} respectively to receive 3D local features. The results presented

Table 7. Results (in %) on COCO for different feature extractors. The best results are highlighted in bold.

Method	Feature Extractor	Results		Margin of errors	
		1-shot	5-shot	1-shot	5-shot
Base model	ResNet-50	47.02	48.85	0.031	0.039
LCM model		48.92	52.79	0.038	0.041
Base model	ResNet-50 _{CLIP}	53.83	56.31	0.054	0.038
LCM model		54.88	60.54	0.057	0.046
Base model	ViT-B/32 _{CLIP}	58.70	61.53	0.055	0.048
LCM model		60.35	65.01	0.042	0.037

Table 8. Results (in %) on COCO for different word embeddings. The best results are highlighted in bold.

Method	Word Emb.	1-shot	5-shot
Base model	FastText	39.84	41.19
LCM model		40.67	41.55
Base model	Skip-Gram	42.68	43.11
LCM model		43.56	44.87
Base model	GloVe	47.02	48.85
LCM model		48.92	52.79
Base model	Word _{CLIP}	45.34	48.08
LCM model		46.67	49.32
Base model	BERT _{avg}	40.47	43.41
LCM model		42.86	44.60
Base model	mirrorWIC	42.54	46.16
LCM model		45.32	48.07
Base model	BERT _{biencoder}	47.79	49.05
LCM model		49.30	52.61
Base model	BERT _{concn}	49.29	51.25
LCM model		51.27	54.17
Base model	mirrorBERT	49.49	51.49
LCM model		51.55	54.98

in Tab. 7 demonstrate remarkable improvements by utilizing feature extractors from CLIP. This observation can be attributed to the extensive training of the CLIP model on a large volume of text-image data, resulting in its image encoder possessing a rich visual knowledge that aligns closely with textual features.

5.4.3 Robustness of Our Approach. To verify the robustness of our approach, we train our model 10 times with different seeds, and the margin of error is presented by the standard deviation across these diverse runs. The results in Tab. 7 show the robustness of our approach.

5.4.4 Word Embeddings. Tab. 8 compares the results we obtained with different word embeddings: FastText [5], Word2Vec [44], GloVe [46], BERT [16], and word vectors from the CLIP text encoder. To obtain pre-trained label embeddings from BERT, we considered a number of different strategies, as explained in Sec. 5.1.3. In comparison to the image encoder of CLIP, as demonstrated in Tab. 7,

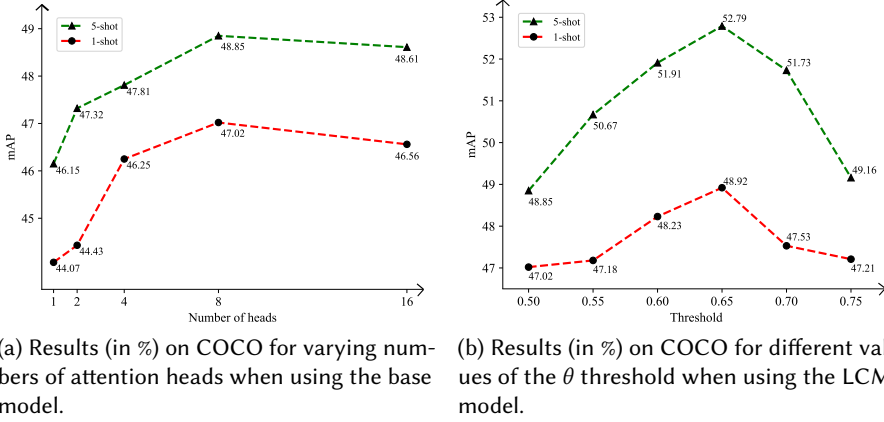


Fig. 3. Ablation study for the number of attention heads and threshold θ .

the word vectors produced by the text encoder of CLIP (obtained as explained in Section 5.1.3) exhibit weaker performance than GloVe and some of the BERT strategies.

5.4.5 Number of Attention Heads. In Fig. 3a, we analyze the importance of using multiple attention heads. It is important to note that adjustments to the number of attention heads necessitate corresponding changes in the division of visual feature channels. However, despite these changes, the dimensionality of the final prototypes remains constant. The results in Fig. 3a show that with fewer than eight attention heads, the performance is notably lower, while there is no obvious benefit to increasing the number of attention heads beyond eight.

5.4.6 LCM Threshold. The hyper-parameter θ controls the threshold for selecting local features based on the learned importance weights. Fig. 3b shows the impact of θ for our LCM model on the COCO dataset. Note that $\theta = 0.5$ corresponds to the base model since 0.5 is the minimum value that can be obtained after applying the sigmoid activation function to a non-negative input. As shown, the LCM model achieves the best performance in both 1-shot and 5-shot settings when $\theta = 0.65$.

Table 9. Zero-shot results (in %) on COCO of various different word embeddings. The best result is highlighted in bold.

Word Embedding	mAP
GloVe	18.79
Word _{CLIP}	18.02
mirrorBERT	23.65

5.5 Discussion

5.5.1 Potential for Zero-shot Classification. While our method was not designed for the zero-shot setting, the fact that word embeddings are aligned with visual features in the joint space means that they can be used as noisy prototypes, which could in principle be used for zero-shot classification. We test the effectiveness of this approach for three types of word embeddings: the conventional GloVe embeddings [46], Word_{CLIP} from the vision-language model CLIP [47], and mirrorBERT [41] from the BERT language model. As shown in Tab. 9, GloVe is slightly better than Word_{CLIP}. Notably,




Fig. 4. Visualization of attention weights among local features used to construct label prototypes for large objects. Darker areas indicate higher importance of the corresponding local features.

mirrorBERT significantly outperforms both, with improvements of 4.86 and 5.63 mAP scores over GloVe and Word_{CLIP}, respectively. These results suggest that embeddings from language models, when aligned with visual features, may be more effective for multi-label zero-shot learning (ML-ZSL) than other label embedding strategies. While these findings are clearly preliminary, we hope these insights can guide the selection of word embeddings in future ML-ZSL methods.

5.5.2 Potential Future Directions. There are two key limitations in our method. First, using off-the-shelf word embeddings for representing labels may not be the optimal strategy. For instance, some approaches for inferring word embeddings from language models are based on contextualised representations, meaning that suitable mentions of these words in a corpus have to be selected. Previous work has shown that selecting the most informative mentions can have an important impact on the quality of the resulting embeddings [63]. One promising avenue might thus be to study strategies for selecting those mentions that are most relevant for vision tasks. Exploring the use of Large Language Models for obtaining more informative embeddings also remains as a promising avenue. Second, our multi-modal cross-interaction module includes two operations between word embeddings and visual features: channel-wised cross-attention guided by word embedding and word embedding-based dynamic convolution. There may be more effective multi-modal interaction operations to explore.

5.5.3 More Than Local Objects. As illustrated in Fig. 4, our model can effectively identify relevant image regions, even when the label pertains to the content of the whole image. This highlights that our method is not constrained to local objects within an image.

6 CONCLUSION

We introduced a novel metric-based method for multi-label few-shot image classification. The main challenge, in this setting, comes from the fact that only a subset of the local features from support images correspond to the given label. To construct informative prototypes, it is essential to estimate which local features are most relevant. With limited support images, we argue that incorporating prior knowledge can effectively address this challenge. In particular, we have proposed to use word embeddings to obtain noisy prototypes, which are then processed by a multi-modal cross-interaction mechanism. This mechanism aims to construct the final prototypes by leveraging multiple interactions between visual features and word embeddings. To further alleviate the influence of irrelevant regions in the image, we have also proposed a loss change measurement module, which analyzes the gradient of the loss function to identify the most representative local features. Our model achieved substantially better results than existing models, both on COCO and newly proposed splits of the PASCAL VOC, NUS-WIDE, and iMaterialist datasets.

ACKNOWLEDGMENTS

This research was supported in part by the China Postdoctoral Science Foundation (2023M740079, GZC20230058). Zied Bouraoui is supported by ANR-22-CE23-0002 ERIANA. Shoaib Jameel is supported by the NVIDIA Academic Hardware Grant. Steven Schockaert was supported by EPSRC grant EP/V025961/1.

REFERENCES

- [1] Rabab Abdelfattah, Qing Guo, Xiaoguang Li, Xiaofeng Wang, and Song Wang. 2023. Cdul: Clip-driven unsupervised learning for multi-label image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 1348–1357.
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems* 35 (2022), 23716–23736.
- [3] Amit Alfassy, Leonid Karlinsky, Amit Aides, Joseph Shtok, Sivan Harary, Rogerio Feris, Raja Giryes, and Alex M Bronstein. 2019. Laso: Label-set operations networks for multi-label few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6548–6557.
- [4] Yuexuan An, Hui Xue, Xingyu Zhao, Ning Xu, Pengfei Fang, and Xin Geng. 2024. Leveraging Bilateral Correlations for Multi-Label Few-Shot Learning. *IEEE Transactions on Neural Networks and Learning Systems* (2024).
- [5] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics* 5 (2017), 135–146. https://doi.org/10.1162/tacl_a_00051
- [6] Shang-Fu Chen, Yi-Chen Chen, Chih-Kuan Yeh, and Yu-Chiang Wang. 2018. Order-free rnn with visual attention for multi-label classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [7] Tianshui Chen, Liang Lin, Xiaolu Hui, Riquan Chen, and Hefeng Wu. 2020. Knowledge-guided multi-label few-shot learning for general image recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020).
- [8] Zhengyu Chen, Jixie Ge, Heshen Zhan, Siteng Huang, and Donglin Wang. 2021. Pareto Self-Supervised Training for Few-Shot Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13663–13672.
- [9] Zhao-Min Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo. 2019. Multi-label image recognition with graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5177–5186.
- [10] Hao Cheng, Yufei Wang, Haoliang Li, Alex C Kot, and Bihan Wen. 2023. Disentangled feature representation for few-shot image classification. *IEEE Transactions on Neural Networks and Learning Systems* (2023).
- [11] Xiao Chu, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. 2016. Structured feature learning for pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4715–4723.
- [12] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yantao Zheng. 2009. NUS-WIDE: a real-world web image database from National University of Singapore. In *Proceedings of the 8th ACM International Conference on Image and Video Retrieval, CIVR 2009, Santorini Island, Greece, July 8-10, 2009*. ACM.
- [13] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yantao Zheng. 2009. Nus-wide: a real-world web image database from national university of singapore. In *Proceedings of the ACM international conference on image and video retrieval*. 1–9.
- [14] Elijah Cole, Oisín Mac Aodha, Titouan Llorient, Pietro Perona, Dan Morris, and Nebojsa Jojic. 2021. Multi-label learning from single positive labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 933–942.
- [15] Leichao Dai, Lin Feng, Xinglin Shang, and Han Su. 2023. Cross modal adaptive few-shot learning based on task dependence. *Chinese Journal of Electronics* 32, 1 (2023), 85–96.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings NAACL-HLT*.
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [18] Thibaut Durand, Nazanin Mehrasa, and Greg Mori. 2019. Learning a deep convnet for multi-label classification with partial labels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 647–657.
- [19] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. 2015. The pascal visual object classes challenge: A retrospective. *International journal of computer vision* 111, 1 (2015), 98–136.

- [20] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*. 1126–1135.
- [21] Amit Gajbhiye, Luis Espinosa-Anke, and Steven Schockaert. 2022. Modelling commonsense properties using pre-trained bi-encoders. *arXiv preprint arXiv:2210.02771* (2022).
- [22] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. 2024. Clip-adapter: Better vision-language models with feature adapters. *International Journal of Computer Vision* 132, 2 (2024), 581–595.
- [23] Shijie Geng, Jianbo Yuan, Yu Tian, Yuxiao Chen, and Yongfeng Zhang. 2023. HiCLIP: Contrastive language-image pretraining with hierarchy-aware attention. *arXiv preprint arXiv:2303.02995* (2023).
- [24] Sheng Guo, Weilin Huang, Xiao Zhang, Prasanna Srikhanta, Yin Cui, Yuan Li, Hartwig Adam, Matthew R Scott, and Serge Belongie. 2019. The materialist fashion attribute dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*. 0–0.
- [25] Zixian Guo, Bowen Dong, Zhilong Ji, Jinfeng Bai, Yiwen Guo, and Wangmeng Zuo. 2023. Texts as images in prompt tuning for multi-label image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2808–2817.
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proc. CVPR*. 770–778.
- [27] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*. PMLR, 4904–4916.
- [28] Jongmin Kim, Taesup Kim, Sungwoong Kim, and Chang D Yoo. 2019. Edge-labeling graph neural network for few-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 11–20.
- [29] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*.
- [30] Wenlan Kuang and Zhixin Li. 2024. Multi-label image classification with multi-layered multi-perspective dynamic semantic representation. *Machine Learning* 113, 6 (2024), 3443–3461.
- [31] Kibok Lee, Hao Yang, Satyaki Chakraborty, Zhaowei Cai, Gurumurthy Swaminathan, Avinash Ravichandran, and Onkar Dabeer. 2022. Rethinking few-shot object detection on a multi-domain benchmark. In *European Conference on Computer Vision*. 366–382.
- [32] Aoxue Li, Weiran Huang, Xu Lan, Jiashi Feng, Zhenguo Li, and Liwei Wang. 2020. Boosting few-shot learning with adaptive margin loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 12576–12584.
- [33] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*. 12888–12900.
- [34] Na Li, Zied Bouraoui, José Camacho-Collados, Luis Espinosa Anke, Qing Gu, and Steven Schockaert. 2021. Modelling General Properties of Nouns by Selectively Averaging Contextualised Embeddings. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021*. 3850–3856.
- [35] Na Li, Hanane Kteich, Zied Bouraoui, and Steven Schockaert. 2023. Distilling semantic concept embeddings from contrastively fine-tuned language models. *arXiv preprint arXiv:2305.09785* (2023).
- [36] Wenbin Li, Lei Wang, Jinglin Xu, Jing Huo, Yang Gao, and Jiebo Luo. 2019. Revisiting local descriptor based image-to-class measure for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7260–7268.
- [37] Zhixin Li, Lan Lin, Canlong Zhang, Huifang Ma, Weizhong Zhao, and Zhiping Shi. 2021. A semi-supervised learning approach based on adaptive weighted fusion for automatic image annotation. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 17, 1 (2021), 1–23.
- [38] Zeqian Li, Michael Mozer, and Jacob Whitehill. 2021. Compositional embeddings for multi-label one-shot learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 296–304.
- [39] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. 2017. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835* (2017).
- [40] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*. Springer, 740–755.
- [41] Fanguy Liu, Ivan Vulic, Anna Korhonen, and Nigel Collier. [n. d.]. Fast, Effective, and Self-Supervised: Transforming Masked Language Models into Universal Lexical and Sentence Encoders. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021*. 1442–1459.
- [42] Qianchu Liu, Fanguy Liu, Nigel Collier, Anna Korhonen, and Ivan Vulic. 2021. MirrorWiC: On Eliciting Word-in-Context Representations from Pretrained Language Models. In *Proceedings of the 25th Conference on Computational*

Natural Language Learning, CoNLL 2021. 562–574.

- [43] Tomáš Mikolov, Quoc V. Le, and Ilya Sutskever. 2013. Exploiting Similarities among Languages for Machine Translation. *CoRR abs/1309.4168* (2013). arXiv:1309.4168 <http://arxiv.org/abs/1309.4168>
- [44] Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013. Linguistic Regularities in Continuous Space Word Representations. In *Proceedings NAACL-HLT*. 746–751.
- [45] Alex Nichol, Joshua Achiam, and John Schulman. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999* (2018).
- [46] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [47] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. 8748–8763.
- [48] Sachin Ravi and Hugo Larochelle. 2017. Optimization as a model for few-shot learning. In *International conference on learning representations*.
- [49] Andrei A. Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. 2019. Meta-Learning with Latent Embedding Optimization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- [50] Victor Garcia Satorras and Joan Bruna Estrach. 2018. Few-shot learning with graph neural networks. In *International conference on learning representations*.
- [51] Julio Silva-Rodriguez, Sina Hajimiri, Ismail Ben Ayed, and Jose Dolz. 2023. A Closer Look at the Few-Shot Adaptation of Large Vision-Language Models. *arXiv preprint arXiv:2312.12730* (2023).
- [52] Christian Simon, Piotr Koniusz, and Mehrtash Harandi. 2022. Meta-learning for multi-label few-shot classification. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 3951–3960.
- [53] Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems* 30 (2017).
- [54] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, et al. 2021. Sparse r-cnn: End-to-end object detection with learnable proposals. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 14454–14463.
- [55] Ximeng Sun, Ping Hu, and Kate Saenko. 2022. Dualcoop: Fast adaptation to multi-label recognition with limited annotations. *Advances in Neural Information Processing Systems* 35 (2022), 30569–30582.
- [56] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. 2018. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1199–1208.
- [57] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2818–2826.
- [58] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems* 30 (2017).
- [59] Grigorios Tsoumakos and Ioannis Katakis. 2007. Multi-label classification: An overview. *International Journal of Data Warehousing and Mining (IJDW)* 3, 3 (2007), 1–13.
- [60] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [61] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. Matching Networks for One Shot Learning. In *Proc. NIPS*. 3630–3638.
- [62] Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu. 2016. Cnn-rnn: A unified framework for multi-label image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2285–2294.
- [63] Yixiao Wang, Zied Bouraoui, Luis Espinosa Anke, and Steven Schockaert. 2022. Sentence Selection Strategies for Distilling Word Embeddings from BERT. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference, LREC 2022, Marseille, France, 20-25 June 2022*. European Language Resources Association, 2591–2600.
- [64] Ya Wang, Dongliang He, Fu Li, Xiang Long, Zhichao Zhou, Jinwen Ma, and Shilei Wen. 2020. Multi-label classification with label graph superimposing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 12265–12272.
- [65] Zhouxia Wang, Tianshui Chen, Guanbin Li, Ruijia Xu, and Liang Lin. 2017. Multi-label image recognition by recurrently discovering attentional regions. In *Proceedings of the IEEE international conference on computer vision*. 464–472.
- [66] Zihao Wei, Zixuan Pan, and Andrew Owens. 2024. Efficient Vision-Language Pre-training by Cluster Masking. *arXiv preprint arXiv:2405.08815* (2024).

- [67] Chen Xing, Negar Rostamzadeh, Boris Oreshkin, and Pedro O O Pinheiro. 2019. Adaptive cross-modal few-shot learning. *Advances in Neural Information Processing Systems* 32 (2019).
- [68] Kun Yan, Zied Bouraoui, Ping Wang, Shoaib Jameel, and Steven Schockaert. 2021. Aligning visual prototypes with BERT embeddings for few-shot learning. In *Proceedings of the 2021 International Conference on Multimedia Retrieval*. 367–375.
- [69] Kun Yan, Zied Bouraoui, Ping Wang, Shoaib Jameel, and Steven Schockaert. 2021. Few-shot image classification with multi-facet prototypes. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1740–1744.
- [70] Kun Yan, Lingbo Liu, Jun Hou, and Ping Wang. 2021. Representative Local Feature Mining for Few-Shot Learning. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1730–1734.
- [71] Kun Yan, Chenbin Zhang, Jun Hou, Ping Wang, Zied Bouraoui, Shoaib Jameel, and Steven Schockaert. 2022. Inferring prototypes for multi-label few-shot image classification with word vector guided attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 2991–2999.
- [72] Kaicheng Yang, Jiankang Deng, Xiang An, Jiawei Li, Ziyong Feng, Jia Guo, Jing Yang, and Tongliang Liu. 2023. Alip: Adaptive language-image pre-training with synthetic caption. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2922–2931.
- [73] Vacit Oguz Yazici, Abel Gonzalez-Garcia, Arnau Ramisa, Bartlomiej Twardowski, and Joost van de Weijer. 2020. Orderless recurrent models for multi-label classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13440–13449.
- [74] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. 2020. Few-shot learning via embedding adaptation with set-to-set functions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 8808–8817.
- [75] Renchun You, Zhiyao Guo, Lei Cui, Xiang Long, Yingze Bao, and Shilei Wen. 2020. Cross-modality attention with semantic graph embedding for multi-label classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 12709–12716.
- [76] Zhonghui You, Kun Yan, Jinmian Ye, Meng Ma, and Ping Wang. 2019. Gate decorator: Global filter pruning method for accelerating deep convolutional neural networks. *Advances in Neural Information Processing Systems* 32 (2019).
- [77] Jin Yuan, Shikai Chen, Yao Zhang, Zhongchao Shi, Xin Geng, Jianping Fan, and Yong Rui. 2023. Graph attention transformer network for multi-label image classification. *ACM Transactions on Multimedia Computing, Communications and Applications* 19, 4 (2023), 1–16.
- [78] Chi Zhang, Yujun Cai, Guosheng Lin, and Chunhua Shen. 2022. Deepemd: Differentiable earth mover’s distance for few-shot learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).
- [79] Ruru Zhang, E Haihong, Meina Song, and Xun Cao. 2024. FSCIL-EACA: Few-Shot Class-Incremental Learning Network Based on Embedding Augmentation and Classifier Adaptation for Image Classification. *Chinese Journal of Electronics* 33, 1 (2024), 139–152.
- [80] Yunfeng Zhao, Guoxian Yu, Lei Liu, Zhongmin Yan, Carlotta Domeniconi, and Lizhen Cui. 2021. Few-shot partial multi-label learning. In *2021 IEEE International Conference on Data Mining (ICDM)*. IEEE, 926–935.
- [81] Yunfeng Zhao, Guoxian Yu, Lei Liu, Zhongmin Yan, Carlotta Domeniconi, Xiayan Zhang, and Lizhen Cui. 2023. Few-shot partial multi-label learning via prototype rectification. *Knowledge and Information Systems* 65, 4 (2023), 1851–1880.
- [82] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 16816–16825.
- [83] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022. Learning to prompt for vision-language models. *International Journal of Computer Vision* 130, 9 (2022), 2337–2348.
- [84] Wei Zhou, Peng Dou, Tao Su, Haifeng Hu, and Zhijie Zheng. 2023. Feature learning network with transformer for multi-label image classification. *Pattern Recognition* 136 (2023), 109203.
- [85] Wei Zhou, Yanke Hou, Dihou Chen, Haifeng Hu, and Tao Su. 2023. Attention-augmented memory network for image multi-label classification. *ACM Transactions on Multimedia Computing, Communications and Applications* 19, 3 (2023), 1–24.
- [86] Wei Zhou, Zhiwu Xia, Peng Dou, Tao Su, and Haifeng Hu. 2023. Double attention based on graph attention network for image multi-label classification. *ACM Transactions on Multimedia Computing, Communications and Applications* 19, 1 (2023), 1–23.
- [87] Feng Zhu, Hongsheng Li, Wanli Ouyang, Nenghai Yu, and Xiaogang Wang. 2017. Learning spatial regularization with image-level supervisions for multi-label image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5513–5522.
- [88] Ke Zhu and Jianxin Wu. 2021. Residual attention: A simple but effective method for multi-label recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 184–193.

A EXPERIMENTS

Table 10. Labels covered by the training, validation, and novel sets for the four considered multi-label few-shot image classification datasets.

ML-FSIC Datasets				
	COCO	PASCAL VOC	NUS-WIDE	iMaterialist
Training Set:	{ hair drier, clock, wine glass, book, cake, tie, motorcycle, sheep, bottle, giraffe, cell phone, suitcase, remote, bench, mouse, carrot, banana, train, sports ball, toothbrush, fire hydrant, airplane, tv, bus, refrigerator, couch, knife, toilet, elephant, truck, parking meter, car, potted plant, kite, skateboard, orange, horse, cat, tennis racket, bowl, scissors, baseball glove, apple, traffic light, handbag, donut, dog, hot dog, oven, umbrella, sink, pizza }	{ person, car, chair, bottle, horse, bicycle, motorbike, diningtable }	{ animal, whales, railroad, clouds, person, leaf, elk, frost, map, snow, waterfall, toy, wedding, protest, book, running, cat, airport, harbor, valley, soccer, train, sky, moon, swimmers, window, dog, boats, buildings, plants, road, tattoo, nighttime, castle, dancing, plane, zebra, tiger, earthquake, sports, grass, water, glacier, rainbow, fox, rocks, cow, ocean, surf, mountain, computer, flags }	{ epaulette, fringe, cape, top_t-shirt_sweatshirt, sequin, bag_wallet, bow, leg warmer, shoe, sleeve, pants, dress, headband_head-covering_hair-accessory, tassel, ribbon, pocket, umbrella, neckline }
Validation Set:	{ cow, dining table, zebra, sandwich, bear, toaster, person, laptop, bed, teddy bear, baseball bat, skis }	{ boat, cow, train, airplane, bus, bird }	{ tower, tree, food, sun, town, beach, military, birds, flowers, house, lake, garden }	{ watch, lapel, shorts, bead, glasses, hat, sweater, tights_stockings, buckle, skirt, applique, rivet }
Novel Set:	{ bicycle, boat, stop sign, bird, backpack, frisbee, snowboard, surfboard, cup, fork, spoon, broccoli, chair, keyboard, microwave, vase }	{ dog, sofa, cat, potted plant, tv monitor, sheep }	{ bear, statue, bridge, sunset, sign, coral, fish, fire, cars, sand, vehicle, cityscape, police, temple, horses, reflection, street }	{ sock, jumpsuit, glove, scarf, coat, belt, zipper, cardigan, jacket, hood, vest, flower, tie, collar, ruffle, shirt_blouse }

A.1 Datasets

An overview of the labels contained in four datasets, including COCO [40], PASCAL VOC [19], NUS-WIDE [12], and iMaterialist [24], can be found in Tab. 10.

A.2 Experimental Results

The experimental results for PASCAL VOC, NUS-WIDE, and iMaterialist are shown in Tab. 11, Tab. 12, and Tab. 13, respectively. Our method consistently achieves the best performance.

A.3 Ablation Study

A.3.1 Importance of the Attention Mechanism. To evaluate the effect of the proposed attention mechanism, in Tab. 14, we report results for a variant that relies on a simpler mechanism for generating prototypes (shown as *Simple attention*). In particular, in this variant, we generate the p_c^{att} of a label c by taking a weighted average of the global features of the support images which

Table 11. Overview of the main results on PASCAL VOC, in terms of micro/macro AP and F1 scores (in %). The best results are highlighted in bold font. “-” indicates that no word embeddings are used.

	Word-Emb.	1-shot				5-shot			
		Mi-AP	Mi-F1	Ma-AP	Ma-F1	Mi-AP	Mi-F1	Ma-AP	Ma-F1
ResNet-50	-	16.57	20.77	27.24	21.92	24.47	28.53	35.06	29.43
ResNet-101	-	19.23	23.86	30.11	22.17	27.07	29.37	39.01	31.26
ViT	-	16.69	17.76	27.72	17.15	31.70	37.35	40.13	36.29
ResNet-50 + ViT	-	18.71	13.70	28.76	12.86	34.69	35.83	42.73	33.99
PLA	Self-learned	41.98	36.81	47.29	37.83	48.09	42.01	55.54	45.40
PLA (GloVe)	GloVe	41.31	35.84	46.54	36.55	47.50	40.60	55.26	45.33
LaSO	-	22.12	27.02	32.50	25.20	30.61	33.64	41.89	34.96
MAML	-	40.21	35.11	44.15	36.08	45.95	39.26	50.03	42.74
WGA	GloVe	46.28	40.19	53.26	43.35	51.10	45.03	58.27	48.20
Base model	GloVe	49.92	44.00	56.59	47.05	54.35	48.21	61.48	51.21
LCM model	GloVe	50.98	44.18	58.03	47.31	56.74	49.08	65.41	52.00
Base model	mirrorBERT	51.19	44.37	58.43	47.96	55.36	49.37	63.33	51.55
LCM model	mirrorBERT	53.30	46.05	59.13	49.65	58.30	51.19	66.78	54.01

Table 12. Overview of the main results on NUS-WIDE, in terms of micro/macro AP and F1 scores (in %). The best results are highlighted in bold font. “-” indicates that no word embeddings are used.

	Word-Emb.	1-shot				5-shot			
		Mi-AP	Mi-F1	Ma-AP	Ma-F1	Mi-AP	Mi-F1	Ma-AP	Ma-F1
ResNet-50	-	16.46	21.49	20.78	21.77	21.71	25.33	24.89	26.02
ResNet-101	-	15.10	20.27	17.90	20.18	22.60	25.82	24.81	26.25
ViT	-	17.08	20.76	19.94	21.97	24.85	28.58	27.04	26.57
ResNet-50 + ViT	-	16.05	18.92	18.10	15.13	26.32	27.84	29.09	25.15
PLA	Self-learned	22.31	23.92	27.85	24.63	30.32	31.59	35.24	30.33
PLA (GloVe)	GloVe	22.47	23.86	27.69	24.29	29.78	30.47	36.01	31.04
LaSO	-	20.27	23.17	25.95	22.33	25.35	28.91	29.63	27.77
MAML	-	25.32	26.94	30.09	24.82	32.95	32.52	36.17	30.50
WGA	GloVe	32.91	31.53	38.12	30.71	36.81	35.16	41.25	35.63
Base model	GloVe	37.88	35.31	42.39	35.53	42.17	40.36	47.02	40.80
LCM model	GloVe	38.77	36.09	44.96	36.08	45.65	42.31	50.91	43.65
Base model	mirrorBERT	38.85	35.32	43.53	35.89	44.16	41.85	48.54	42.54
LCM model	mirrorBERT	39.64	36.61	45.00	37.30	46.19	43.81	52.58	44.43

have that label. The weights are obtained by computing the cosine similarity between the vectors $\widehat{\mathbf{w}}_c$ and the feature vectors $\widehat{\mathbf{f}}_{gl}^l$, multiplying these cosine similarities with the scalar λ , and feeding the resulting values to a softmax layer. As shown in Tab. 14, the performance degradation in this simplified setting is clearly evident. We also present results for a variant employing the standard attention mechanism, denoted as *Standard Attention*. In this variant, the inputs of *value* and *key* are generated by processing visual features through fully connected layers, rather than splitting them evenly along the channel dimension. Additionally, we report results when local features are used instead of global features within our cross-attention mechanism (*Attention with global features*). Finally, we experimented with a variant that learns visual prototypes in the same way as in KGGR, where the weights of local features are generated by a low-rank bilinear pooling method

Table 13. Overview of the main results on iMaterialist, in terms of micro/macro AP and F1 scores (in %). The best results are highlighted in bold font. “-” indicates that no word embeddings are used.

	Word-Emb.	1-shot				5-shot			
		Mi-AP	Mi-F1	Ma-AP	Ma-F1	Mi-AP	Mi-F1	Ma-AP	Ma-F1
ResNet-50	-	23.58	27.02	22.16	19.33	34.90	31.84	30.06	25.10
ResNet-101	-	16.71	26.35	21.76	21.45	38.79	32.89	34.39	28.32
ViT	-	23.39	29.45	24.20	15.54	38.98	37.18	35.44	28.20
ResNet-50 + ViT	-	24.32	30.50	25.27	17.13	40.46	38.96	35.73	28.44
PLA	Self-learned	31.61	27.82	36.41	27.65	41.69	30.78	37.08	30.65
PLA (GloVe)	GloVe	30.43	25.94	34.97	27.46	40.46	29.91	36.49	30.09
LaSO	-	29.69	25.31	30.51	24.82	38.78	29.25	35.87	27.82
MAML	-	31.42	29.40	37.96	27.32	41.70	31.10	39.89	30.17
WGA	GloVe	41.01	38.64	46.37	34.74	52.13	41.48	47.26	38.32
Base model	GloVe	45.90	43.45	51.23	40.06	57.96	46.60	52.57	43.49
LCM model	GloVe	47.93	44.13	53.09	41.07	61.50	48.97	56.74	45.15
Base model	mirrorBERT	48.26	43.80	54.25	42.13	59.83	47.07	54.96	44.89
LCM model	mirrorBERT	50.09	44.71	55.27	42.25	62.43	48.45	57.62	45.18

Table 14. Results (in %) on COCO for different attention mechanism. The best results are highlighted in bold.

Method	Difference	1-shot	5-shot
Base model	Simple attention	41.42	43.25
LCM model		41.83	44.17
Base model	Standard attention	45.84	46.91
LCM model		47.00	49.45
Base model	Attention with global features	42.68	45.97
LCM model		43.11	47.60
Base model	Low-rank bilinear pooling	38.82	41.00
LCM model		39.29	42.71
Base model	-	47.02	48.85
LCM model		48.92	52.79

Table 15. Results (in %) on COCO for different ways to select representative local features for the LCM model based on the 5-shot setting. The best results are highlighted in bold.

Method	Ma-AP	Mi-AP
Base model	48.85	42.84
Using $\rho_{j,k}$	51.46	45.18
Directly using $g(\cdot)$	50.16	43.55
Taylor approximation	51.73	45.63
LCM model	52.79	46.35

followed by a fully connected layer with softmax activation (*Low-rank bilinear pooling*). The results in Tab. 14 show that all of these variants result in lower performance.

A.3.2 Selection of Local Features. There are a number of ways in which the LCM module can be simplified. First, we could use the importance weights $\rho_{j,k}$ directly, i.e. we could select those local

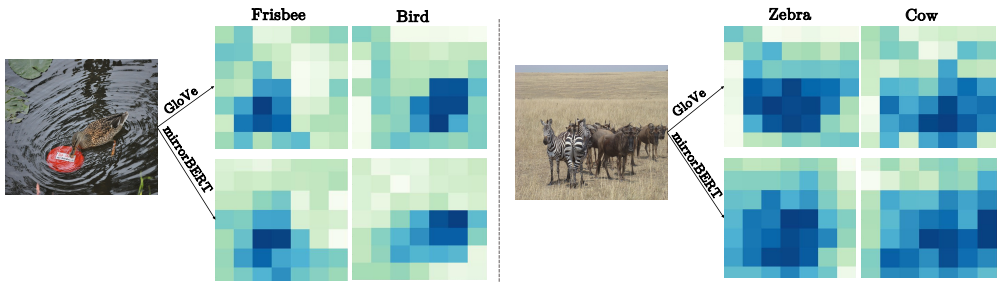


Fig. 5. Visualization of the attention weights among local features for constructing label prototypes, using GloVe and mirrorBERT word vectors, across various categories. The color intensity indicates the relative importance of the corresponding local feature.

features whose $\rho_{j,k}$ -value is above a certain threshold. Next, the implementation of loss change measurement itself could also be simplified. For instance, we could evaluate the loss change $g(\cdot)$ directly using Eq. 8. A final possibility we consider is to use the Taylor approximation, i.e. to use Eq. 10, but without the momentum-based strategy. These different variants are compared with the full LCM model in Tab. 15. In addition, we include results of the base model in the table for comparison. Notably, all variants exhibit superior performance compared to the base model, with the proposed LCM model attaining the best overall results.

A.3.3 Local Features for Query Images. In our method, we initially compare the global features of query images with prototypes. Additionally, we experimented with a variant of the base model where the local features of the query image are matched against those of the prototypes. In this variant, the similarity score is derived from the sum of the top- k values, with $k = 12$, chosen based on the validation split. This variant achieves mAP scores of 47.54 and 49.11 in the 1-shot and 5-shot settings, respectively. While this represents a slight improvement over our base model (47.02/48.85), it incurs a significantly higher computational cost.

A.3.4 Visualization of the Attention Mechanism. Fig. 5 illustrates which regions the proposed attention mechanism focuses on according to different categories for two choices of the word vectors. We take the sum of attention weights across all attention heads as the visualization scores of the local features. As illustrated in the first example of Fig. 5, the model is often successful in identifying the most relevant image region for different categories and the attention weights are indeed label-specific. In this example, the model correctly selects the frisbee or the duck depending on the selected label. This is even though no images with these labels were present in the training data. On the other hand, as the last example shows, for labels that are semantically closely related, such as *zebra* and *cow* in this case, word vectors are not sufficiently informative. For both labels, the model correctly selects the group of animals, but it fails to make a finer selection.

A.3.5 Visualization of LCM Importance Weights. Fig. 6 visualizes the importance of the different local features, as predicted by the LCM module. The darker the color, the more important the local features are considered to be. Most background local features are given low scores. In fact, many of these features will have a score below the threshold θ and will thus be completely ignored when aggregating prototypes.

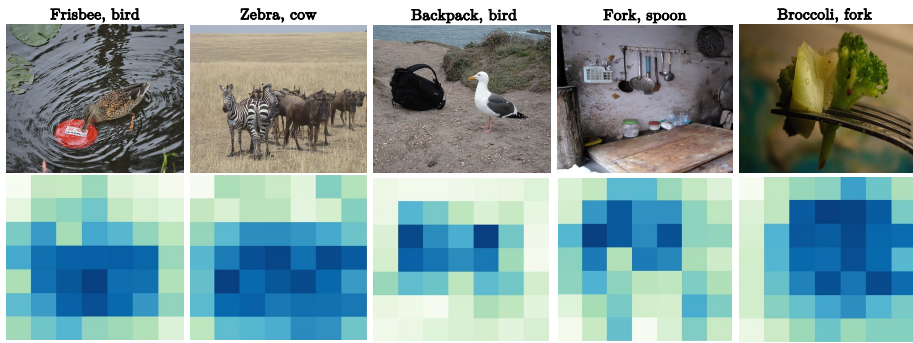


Fig. 6. Visualization of the distribution of importance scores among local features in the LCM module. The more important the local feature, the darker the corresponding visualization.