

Learning to Name Classes for Vision and Language Models

Sarah Parisot Yongxin Yang Steven McDonagh
 Huawei Noah’s Ark Lab

{sarah.parisot, yongxin.yang, steven.mcdonagh}@huawei.com

Abstract

Large scale vision and language models can achieve impressive zero-shot recognition performance by mapping class specific text queries to image content. Two distinct challenges that remain however, are high sensitivity to the choice of hand-crafted class names that define queries, and the difficulty of adaptation to new, smaller datasets. Towards addressing these problems, we propose to leverage available data to learn, for each class, an optimal word embedding as a function of the visual content. By learning new word embeddings on an otherwise frozen model, we are able to retain zero-shot capabilities for new classes, easily adapt models to new datasets, and adjust potentially erroneous, non-descriptive or ambiguous class names. We show that our solution can easily be integrated in image classification and object detection pipelines, yields significant performance gains in multiple scenarios and provides insights into model biases and labelling errors.

1. Introduction

The introduction of large-scale vision and language pre-training techniques has led to significant breakthroughs in object recognition, notably for tasks where data efficiency forms a crucial component (e.g. zero-shot scenarios). The key idea is to learn a mapping between image and text spaces such that a latent image representation will be close to sentences that describe the image content, in representation space [33, 36]. Learning from image-specific descriptions, rather than a fixed and common class label, with limited semantic meaning, yields models with very strong representational power [36]. Such multi-modal approaches address one key shortcoming of visual-only recognition methods; they provide an ability to seamlessly expand the model’s output space without the requirement of additional training (i.e. open-set recognition). In contrast, visual-only classification and object detection methods have historically relied on training a linear classification layer in order to identify objects of interest, resulting in the undesirable consequence that any introduction of new classes (with or without available data) is very challenging. A plethora of complex mechanisms have been proposed, towards addressing this problem [12, 19]. Vision



Figure 1. Vision and language recognition models are highly sensitive to class prompts. Ambiguous or highly technical class names can lead to poor classification performance. Rather than adjusting class names directly, current solutions rely on prompt context to reduce ambiguity, e.g.: ‘a photo of a bow, a type of weapon’.

and language multi-modal approaches alternatively perform classification by computing a similarity between image features and a set of textual features, corresponding to text descriptions from the classes of interest. As the text component constitutes an input, such models naturally enable a route to identify new categories, simply by modifying the queried text, yielding very strong zero-shot and open-set performance [22, 36].

We draw attention to the fact that such performant multi-modal models suffer from two remaining key limitations. Firstly, performance is highly sensitive to class-prompt textual content; both in terms of (i) the selected class name and (ii) the surrounding sentence context. Class name sensitivity can be illustrated by considering class homonyms; words that have the same spelling but different meanings and origins (Figure 1a). These lead to ambiguous queries that negatively effect performance and mitigation necessitates time consuming prompt

engineering. A further example of class name sensitivity can be demonstrated by classes exhibiting highly technical characterisations (Figure 1b). Poor performance is typically observed for datasets exhibiting such complex naming strategies [36]. We hypothesise that lack of clear semantic meaning leave it difficult for multi-modal models to obtain valuable embeddings.

Secondly; model adaptation to new data, that becomes available, is not straightforward – especially when datasets are small. A natural solution is to use the existing representation layers and learn a new classification layer (*i.e.* linear probing). However this strategy must sacrifice the flexibility of text input, for potential subsequent adaptations [36]. Alternatives consider replacing text queries with image queries [33], or learning optimal prompt contexts from the new data [6, 50, 51]. The former can struggle to leverage larger amounts of available data successfully as image queries are averaged over all available class specific data. The latter alternatively updates a fixed set of parameters, which can lead to forgetting issues under sequential adaptation. Furthermore, the approach fails to learn a new class representation or mapping from the new data, as the focus lies solely with the prompt context. Finally; while context learning is highly successful in simple classification tasks, it becomes more challenging for object detection purposes and requires more complex training mechanisms [6]. We conjecture that this may be attributed to the fact that context then depends on the spatial region of the image, in addition to object type (*i.e.* foreground or background).

These considerations lead us to propose a complementary solution to prompt learning, which can be successfully leveraged for both classification and detection problems, without any task-specific tailoring or requirements. We build upon ideas from the field of text to image generation, and propose to leverage the concept of textual inversion [10] to adapt models to new data and thus address suboptimal, and potentially wrong, labels (*e.g.* incorrect class names, mislabelled data). The crux of our idea is to supersede manual word engineering, describing categories of interest, by leveraging available data to learn optimal word embeddings, as a function of the visual content. This simple solution affords several advantages: 1) in contrast to linear probing, open-set capabilities are maintained; learned class representations remain in text space, 2) our approach learns class specific parameters, therefore preventing forgetting under sequential training, 3) direct applicability to any classification or object detection technique that utilises textual input, 4) interpretability properties, with regards to learned class names.

Experimentally, we realise our ideas in conjunction with two contemporary models; for image classification [36] and object detection [33]. Given a pre-trained model and (potentially previously unseen, additional) training dataset, we introduce a new set of word embeddings in the model text branch, each corresponding to a class of interest. Following context tuning methods [50, 51], embeddings are learned using standard losses on an otherwise frozen model. Comprehensive experiments across thirteen classification and detection datasets help to evidence that

learning optimal class names can significantly enhance model performance, maintain open-vocabulary properties, successfully enable continual model adaptation, and improve performance for rare and longtail classes. Last but not least, we demonstrate the method’s potential for interpretability, notably with regards to mislabelling and inadequately chosen original class names.

Our main contributions can be summarised as:

- A simple proposition for data-efficient adaptation of large vision-language models: optimisation of class names from visual content that retains attractive open-vocabulary properties.
- Performance improvements *complementary* to those of prompt tuning; consistent experimental gains across 13 classification and detection datasets. Strong sequential adaptation, open-vocabulary and long tail performance.
- Model interpretability insights: we visualise class name changes and related semantic meanings. We highlight how model biases and labelling errors can be identified, and how vision-language models suffer from long-tail issues.

2. Related work

Vision and language models. Early vision and language multi-modal alignment work tended to enable learning by encoding distinct modalities individually and independently. Model branch outputs would then be connected by losses [41]. Cross-modality alignment has been performed using techniques including metric learning [8], multi-label classification [20] and transformers [48]. A contemporary alternative approach pertains to so called “vision-language” models [9, 18, 29, 36] that recently gain popularity and instead are designed to jointly learn modality encoders; sharing information between modalities to the benefit of the global model. A representative approach is CLIP [36], which trains two neural network-based encoders on a very large image-caption dataset using a contrastive cross-entropy loss, to match pairs of images and texts, and shows impressive zero-shot image recognition performance. This seminal work has been extensively investigated in recent literature, and several improvements have been proposed [9, 18, 46].

Open vocabulary object detection was originally explored in a two-step manner [45]: firstly learning a mapping between the two modalities using a large-scale dataset, then leveraging the obtained encoders and mapping layers to learn an object detector that can perform classification by computing a distance to class names’ text embeddings. More recent methods fully integrate vision-language mappings into the object detection training process, providing more flexible and adaptable multi-modal detection solutions [2, 13, 22, 26, 31, 33]. Alternative methods generate a multi-modal joint representation via a fusion block [22, 31], while contemporary solutions adopt a CLIP-like parallel training strategy [2, 13, 33], facilitating querying of large numbers of

classes. Similarly to the classification setting, multi-modal detection models achieve impressive zero-shot performance.

Prompting methods. Prompt learning aims to exploit the rich knowledge contained in large pre-trained language models (e.g. GPT [37]), to predict new downstream tasks. The idea prescribes that we can reformulate these new tasks to “look more like” those solved during the original pre-training stage, through the help of an appropriate textual prompt, such that the model can be used to predict the desired output without any additional training [11,30,37,39]. The disadvantage with such methods, is the necessity for *prompt engineering*; searching for the most suitable prompt to allow the language model to solve the desired task. Recent work has looked to learn optimal prompts *cf.* manually designing them. Work on continuous prompt learning methods [24,28,49] can be considered closely related to our work, where the main idea is to convert the prompt into a set of continuous vectors that can be optimised, with respect to an objective function, in end-to-end manner.

The benefits of the prompting strategy have recently started to be explored in computer vision [21,38,44,47,51]. CoOp [51] can be considered the first work to bring continuous prompt learning to the vision domain, towards offering efficient solutions for adaptation of pre-trained vision-language models to downstream vision tasks. CoOp was recently extended by [50] where they propose to learn an additional light-weight neural network to generate, for each image, an input-conditional token that yields an image-specific dynamic prompt, outperforming CoOp’s static prompts. This comes at the price of efficiency however, as the proposed extension is substantially both more memory and computationally expensive. We note that [50,51] facilitate model adaptation to new domains in a highly parameter efficient manner, yet may still struggle in settings where new classes (and their names) widely differ from that which the model has initially learned. A more complex context prompting strategy has been recently proposed, specifically for detection tasks [6]; integrating a background specific loss function, as well as a context dependent scheme in the prompt learning process.

Crucially, for all of these methods, classification is still performed with respect to a handcrafted class name, making accuracy highly dependent on the quality of chosen names. In certain settings, finding an accurate semantic description of a class can be challenging, and a poor choice can impair the model’s performance. To the best of our knowledge, this limitation has not been addressed yet in classification and object detection settings.

Textual Inversion. Textual inversion was recently introduced in [10] for text to image generative diffusion models, as a means to condition image generation towards incorporating specific objects or visual components. Given a set of images, all displaying the same object in diverse settings, a specific word embedding is learned so as to ensure the model is able to reconstruct images, sampled from this small set. Simple prompt contexts are used (e.g. ‘a photo of’) to ensure a simple image of the object is generated. After training, the model displays impressive ability




Figure 2. Overview of our method. Classes of interest are mapped to a set of learnable word embeddings, and integrated within the input sentence representation. These embeddings can be learned for classification and object detection tasks using standard losses, and used in conjunction with pre-trained word embeddings in open-vocabulary settings.

to integrate specific objects in generated images, via inputting learned words at specific locations in the input sentence. In the solution proposed by [10], each word is learned individually.

3. Methodology

We consider a vision and language object recognition model pre-trained on a large-scale dataset. Following the multi-modal architecture, the model takes as input an image-text pair $\mathbf{x} = \{I, T\}$, where the text provides a list of N candidate classes t_i to recognise; $T = [t_1, t_2, \dots, t_N]$. We assume the i -th class specific text t_i to be in the format: $t_i = [\text{prompt prefix}] + [\text{CLASS}] + [\text{prompt suffix}]$, where prompt prefixes and suffixes are shared across all classes. To obtain text feature representations, a tokenizer converts each sentence into a series of tokens such that each token maps to a pre-trained word embedding \mathbf{e}_k . Embeddings are then concatenated and used as input to a text transformer. Our final text representation is the transformer output for the End Of Sequence (EOS) token $f(T) = f_i(t_i)$, where $i \in \{1, \dots, N\}$.

3.1. Baseline multi-modal models

Multi-modal Image classification. In the context of image classification [36], our objective is to assign a specific category to the image as a whole. Image inputs are processed in parallel to the text input through an image encoder $g(\cdot)$, yielding repre-

sentation $g(I)$. The classification task is then simply performed by evaluating the cosine similarity between image and text embeddings, and selecting the class with the highest similarity.

Multi-modal Object detection. There are two main contemporary strategies for vision and language based object detection. The first (simpler) strategy is similar to multi-modal classification, wherein image and text information are only combined latterly at the bounding box classification stage [2, 13, 33]. In this setting, the image input is processed through an object detector entirely disjointly from the text modality. The detector outputs a set of coordinates and a class embedding for each bounding box candidate. We then compute class embedding similarity with text representations $f(T)$ following the multi-modal classification approach to assign a classification label to each bounding box.

The second class of strategies propose to fuse image and text inputs; either when generating text and image embeddings [27] or via a third encoder $h(g(I), f(T))$ that outputs bounding box coordinates and class embeddings [13, 22, 31]. For this latter fusion strategy, classification is performed over sentence tokens rather than across sentences, with an aim of associating bounding boxes with a keyword. This technique is typically more computationally expensive than the first method class as it requires computation of new image features, for each queried class.

3.2. Learning optimal class names

We conjecture that the performance of pre-trained recognition models is intrinsically linked to the suitability of the class names, chosen to represent certain visual categories. Selecting class names that do not reflect an object’s visual appearance can substantially alter recognition performance. We thus propose to remove the model’s sensitivity to hand crafted class names, by *learning class specific word embeddings from image content*. An overview of the method is provided in Figure 2.

We consider the matrix of pre-trained word embeddings $E \in \mathbb{R}^{V \times F}$, where V is the vocabulary size, and F the feature dimension. We propose to extend this matrix with a set of new, learnable word embeddings $E^l \in \mathbb{R}^{N \times F}$, where each row of E^l corresponds to a class specific learnable embedding and N pertains to the total number of dataset classes. In this setting, we replace the class name in each text query with a class specific placeholder $t_i = [\text{prompt prefix}] + [\text{pl}^i] + [\text{prompt suffix}]$. While the tokenizer maps context prompts to matrix E , our placeholder maps to its corresponding class embedding e_i^l . This new class query is used as input to the text transformer encoder, and the remainder of the recognition model is unchanged. We note that it is possible to map one category to a series of m word embeddings, using queries of the type $t_i = [\text{prompt prefix}] + [\text{pl}_1^i] + \dots + [\text{pl}_m^i] + [\text{prompt suffix}]$.

Training. During training, we freeze the entire pre-trained model (including pre-trained word embeddings) and learn our new set of word embeddings E^l with the standard losses used to train the base model (e.g. cross entropy loss for classification,

and bounding box regression losses for detection). We identify three key scenarios where our approach can be beneficial: 1) adapting the pre-trained model on a new dataset with a new set of classes; 2) adjusting suboptimal class names on the same dataset used to train the model; 3) continual adaptation of a pre-trained model, where new classes are available after a first training stage.

Inference. At inference time, we follow the process described in Section 3.1. Categories whose names were learned are mapped to their new embeddings $e^l \in E^l$. For unseen object categories (i.e. not in our new dataset) we use their original embeddings, $e \in E$, from the pre-trained model. As a result, our solution maintains the ability to perform open-vocabulary recognition.

4. Experiments

We evaluate the impact of learning class names for the tasks of image classification and object detection. We consider four settings where our strategy can be beneficial:

Task 1: Model adaptation to new datasets: considering a new dataset of interest, we adapt our pre-trained model to this new set of potentially unseen classes by learning optimal class names from the new image content.

Task 2: Post-training class name adjustment: the model is first fine-tuned on our target dataset, then class names are optimised in a subsequent learning stage. This setting is more valuable for object detection where fine-tuning can have important impact on bounding box localisation.

Task 3: Open-vocabulary recognition: open-vocabulary is one of the key advantages of vision-language models, and maintaining strong performance is crucial. Evaluation of the compatibility between our learned word embeddings, and standard pre-trained word embeddings can provide a measure of strength on this axis. We adapt the model to only a subset of classes from our new dataset and measure overall performance on all classes, such that inference combines learned and text based word embeddings.

Task 4: Continual model adaptation: to measure the model ability for sequential learning of new word embeddings, we first adapt the model to a subset of classes, then learn the remaining class names in a second, separate, training stage.

We highlight that we do not seek state of the art performance, as this is intrinsically linked to the pre-trained model’s base performance, pre-training data, and architecture. Rather, our objective is to demonstrate the potential of learning class names to improve model performance, while maintaining zero-shot properties. We therefore perform experiments on models with simple multi-modal mechanisms, use a single word embedding per class, and a standard prompt: (prompt prefix = ‘a photo of a’; prompt suffix = ‘.’), in all experimental settings unless specified otherwise.

4.1. Image classification: CLIP model

To evaluate our method on classification tasks, we employ CLIP [36] as our base pre-trained model with a vision

Class group	CLIP*	CoOp	Ours	Ours*
Base classes	69.4	82.6	82.7	82.8
New classes	74.2	64.1	73.1	74.2
All classes	65.2	65.7	66.5	68.7

Table 1. CLIP model: Base to new classification accuracy. Results on Base classes show method ability to adapt to new data (Task 1). Results on New and All classes show open vocabulary ability (Task 3). * Manually engineered prompt templates.

transformer ViT-B/16 [5] fulfilling the image encoder $g(\cdot)$ role. Here, we seek to evaluate method performance in previously described **Tasks 1, 3 and 4**; *i.e.* model adaptation, open-vocabulary performance, and continual adaptation.

Datasets. We follow the setting proposed in [50, 51]. We adapt the CLIP model to a selection of eleven diverse datasets, focusing on different types of recognition tasks: standard image classification (ImageNet [4], Caltech 101 [7]), fine-grained classification (Stanford cars [23], Oxford Pets [35], Flowers 102 [34], Food 101 [1] and FGVC Aircraft [32]), scene recognition (SUN 397 [43]), action recognition (UCF 101 [42]), texture classification (DTD [3]) and satellite image recognition (Eurosat [17]).

Training details. We adopt a set of parameters similar to [51]. For each dataset, we learn class names with a batch size of 32, over 200 epochs (50 for ImageNet), SGD optimiser with a learning rate of $2e^{-4}$ and cosine scheduling, with one warm up epoch. We initialise our learnable word embeddings using the average embedding of all word tokens associated with hand-crafted class names.

For each dataset, 16 training images per category are sampled, and we report mean statistics over three random seeds. **Baselines.** The method we directly compare to is the zero-shot CLIP model, as we aim to evaluate the benefits of learning class names. We additionally report classification accuracy for the CoOp model [51], which learns optimal prompt context with fixed handcrafted class names. We retrain CoOp in our settings of interest and use four context words, initialised with ‘a photo of a’, as in [50] (this differs from the 16 words used in [51]). This allows to compare advantages and disadvantages of learning class names over prompt context, and evaluate any complementary nature between the strategies. We report the best possible CLIP performance with engineered prompts, and report our results with a standard prompt, as well as using the same engineered prompts as CLIP. We note that four of the datasets (Stanford cars, ImageNet, Caltech 101 and Sun 397) do not have engineered prompt, and we provide all dataset-specific prompts in the supplementary materials.

Model adaptation (Task 1) and Open vocabulary (Task 3). We adopt the base-to-new generalisation setting proposed in [50]. For each dataset, we equally split classes into a set of ‘base’ and ‘new’ classes. We use base classes for training and evaluate performance on new classes in a zero-shot manner. Additionally, to measure open-vocabulary performance, we report performance on ‘all’ classes, *i.e.* a mixture of learned and handcrafted class

Class group	CLIP*	CoOp	Ours	Ours*
New classes	74.0	84.4	85.9	85.9
All classes	65.2	71.7	75.7	76.3

Table 2. CLIP model: Classification accuracy for continual model adaptation (Task 4) experiments. * manually engineered prompt templates.

names. Results are reported in Table 1. Results on base classes highlight our ability to perform model adaptation (Task 1) by learning class names only. For this task, we observe that our strategy can achieve equivalent gains as context tuning on base classes (82.6% (CoOp) vs. 82.8% (Ours) average accuracy). However, our approach guarantees that zero-shot CLIP performance is maintained on new classes, as our learned parameters are class specific. Our open-vocabulary experiments (Task 3) show that we are able to maintain the CLIP model’s ability to expand the output set, and to introduce new classes from text alone (CLIP and Ours 74.2%, new classes). Experiments on all classes confirm that we are able to use a mixture of learned and new classes successfully, outperforming CoOp (65.7% CoOp, 68.7% Ours, all classes). We further note sensitivity to context prompts, as substantial gains are obtained using engineered prompts (+2.2% Ours, all classes). This is to be expected, as new class performance can be strongly impacted by prompts, and highlights the importance of prompt context tuning in this setting.

Sequential adaptation (Task 4). We extend the generalisation setting to evaluate our ability to perform sequential adaptation; *i.e.* after training on base classes, we subsequently train on new classes with no access to base classes. We measure performance on new classes, as well as all classes after this second training stage. For CoOp [51] experiments, we fine-tune prompt parameters, and adjust learning rate to $2e^{-4}$ which provided best empirical performance in this setting. Results are reported in Table 2. As in the previous setting, performance on new classes is highly similar when performing context or class name learning (84.4% CoOp, 85.9% Ours, new classes), achieving strong gains over zero-shot CLIP (10+%). On all classes, we can see that we achieve significant gains (+4%, Ours vs. CoOp), which can be attributed to the fact that class names are (all) learned independently. Last but not least, we note a much higher robustness to change in prompt context when all classes are learned; we achieve very similar results when using a standard or engineered template in this setting (+0.6% vs. +2.2% in Table 1, all classes).

4.2. Object detection: OWL-vit model

For object detection experiments, we evaluate all 4 Tasks. We select OWL-vit [33] as our base multi-modal detection model due to noted strong zero-shot performance and clear separation between vision and language branches. This allows us to evaluate the impact of class learning more directly, in contrast to methods using complex text and image fusion mechanisms [26]. OWL-vit leverages a pre-trained multi-modal classification

Method	backbone	pre-training (image)	pre-training (detection)	AP	AP/f	AP/c	AP/r
Zero-shot models							
MDETR [22]	DETR	N/A	MDETR data	6.4	9.8	3.6	1.9
XDETR [2]	DETR	N/A	MDETR data, LocNar	16.4	18.8	15.2	9.6
GLIP-T [26]	SWIN-T	Cap4M	O365,GOLDG	26.0	31.0	21.4	20.8
GLIP-L [26]	SWIN-L	Cap24M	O365,FourODs	37.3	41.5	43.3	28.2
Zero-shot OWL-vit base model							
OWL-vit zero-shot [33]	VIT B/16	CLIP	O365, VG	25.3	27.0	25.0	17.6
Prompt learning all [51]	VIT B/16	CLIP	O365, VG	24.4	25.9	24.0	17.8
Ours base (Task 3)	VIT B/16	CLIP	O365, VG	25.7	26.8	25.9	19.0
Ours sequential (Task 4)	VIT B/16	CLIP	O365, VG	26.0	26.6	25.9	23.1
Ours all (Task 1)	VIT B/16	CLIP	O365, VG	26.6	27.0	27.5	19.8
OWL-vit base model fine-tuned on LVIS base classes							
OWL-vit [33]	VIT B/16	CLIP	O365, VG	28.8	34.4	24.7	17.7
Prompt learning all [51]	VIT B/16	CLIP	O365, VG	31.3	34.7	29.2	22.4
Ours base (Task 3)	VIT B/16	CLIP	O365, VG	33.3	35.2	33.9	20.3
Ours rare (Task 3)	VIT B/16	CLIP	O365, VG	30.4	34.5	25.4	34.5
Ours sequential (Task 4)	VIT B/16	CLIP	O365, VG	34.1	35.0	33.8	30.6
Ours all (Task 2)	VIT B/16	CLIP	O365, VG	34.5	35.6	34.4	28.5
OWL-vit base model fine-tuned on all LVIS classes							
OWL-vit [33]	VIT B/16	CLIP	O365, VG	34.5	38.5	33.2	19.1

Table 3. Average precision results on LVIS mini-validation set for all (AP), frequent (AP/f), common (AP/c) and rare (AP/r) classes. Our method uses only 10% of the LVIS training dataset. State of the art methods (in grey) use different training datasets and architectures of different capacities. They are added for context. {base, rare, all} describes which class group is used for training. Sequential means base and rare classes are trained sequentially. Task 1: model adaptation to new data, Task 2: class name learning after fine-tuning, Task 3: Open vocabulary ability (mixture of seen and unseen classes), Task 4: continual adaptation (learning class groups in a sequential manner).

model (e.g. CLIP) with a vision transformer backbone. To adapt this model to the object detection task, the pooling layer that combines vision transformer outputs is removed, and two new modules are learned: a class prediction module and a box prediction module that output class embeddings and bounding-box coordinates, respectively, for all transformer outputs. The OWL-vit model is fine-tuned from the CLIP weights using a mixture of standard object detection datasets. Additionally, a set of engineered CLIP prompt templates are used as class prompt contexts.

Datasets. We evaluate the impact of learning class names on two object detection datasets. Firstly, we consider LVIS [14], a large scale natural image object detection dataset comprising over 100k training images and 1203 classes, characterised by a long tail distribution with ‘frequent’ ($x > 100$), ‘common’ ($100 \geq x > 10$), and ‘rare’ ($10 \geq x \geq 1$) class groups, for x images per class. Following prior work [2], we use pre-defined category names, but remove the text in the parentheses, e.g., “flip-flop (sandal)” is replaced by “flip-flop”. We note that this leads to a few instances of classes sharing the same name (e.g. “bow”).

Secondly, we run experiments on CODA 2.0 [25], a self-driving dataset focusing on detection of rare, corner case instances. It comprises validation and test sets containing 5k images each, and 43 classes, of which only 29 appear in the validation set. In addition to the base OWL-vit model, we consider a model fine-tuned on the SODA dataset [15] (fully labelled instance), closer to our test domain. Designed for semi-supervised

training of self-driving models, SODA comprises six classes that overlap with CODA classes. Following [25], we refer to the {*pedestrian, cyclist, car, bus, tram, truck, tricycle*} classes as common classes, and remaining classes as corner-cases.

Training details. We fine-tune the pre-trained OWL-vit model on LVIS, SODA datasets for 10 epochs, using 7 GPUs (batch size 2/GPU), with learning rates $5e^{-6}$ (main model) and $2e^{-7}$ (text branch) using cosine scheduling and an Adam optimiser. Weight decay is unused and gradient clipping is set to 0.1. For LVIS fine-tuning, we train two models: one on all classes, and one on base (frequent + common) classes only, thus following the evaluation set-up of [33]. For CODA 2.0 experiments, models are fine-tuned on the SODA and LVIS datasets, to simulate the setting where corner cases are discovered after training.

For our class learning process, we freeze the entire model aside from new class embeddings. Embeddings are learned for 20 epochs, using a learning rate of $1e^{-2}$ (cosine schedule), 7 GPUS (batch size 10/GPU). They are initialised with the average token embedding of the corresponding class name. For the LVIS dataset, we learn class names with a balanced subsampling of 10% of the training set.

LVIS experiments. We provide average precision results for all 4 Tasks on the LVIS mini-validation set [22] in Table 3. For context, we report methods competitive with OWL-vit, but highlight the key point of comparison is with respect to the base model that exhibits pre-defined class names. We also provide comparison

Method	AP-Com	AP-A	AR-A	AR-Cor	Sum
RetinaNet	0.25	0.38	0.48	0.25	1.36
Faster RCNN	0.28	0.37	0.43	0.17	1.27
Cascade RCNN	0.30	0.42	0.48	0.21	1.41
Zero-shot OWL-vit base model					
OWL-vit zero-shot	0.10	0.20	0.44	0.45	1.19
Ours (Task 3)	0.17	0.31	0.47	0.50	1.45
OWL-vit model fine-tuned on SODA					
OWL-vit	0.21	0.35	0.50	0.44	1.44
Ours (Task 3)	0.23	0.37	0.51	0.47	1.58
OWL-vit model fine-tuned on SODA + LVIS					
OWL-vit	0.21	0.35	0.52	0.51	1.59
Ours (Task 3)	0.23	0.39	0.53	0.53	1.67

Table 4. Detection results on the CODA test set. AP-Com (AP-Common): average precision (AP) on common classes, AP-A (AP-agnostic): AP in the class agnostic setting. AR-A and AR-Cor: average recall for class agnostic and corner case detection respectively. Task 3: Open vocabulary ability (mixture of seen and unseen classes)

to prompt learning, using four learnable context words initialised as ‘a photo of a’, and identical training parameters to ours.

When learning names for all classes, we achieve consistent performance gains for all class groups in both zero-shot (Task 1: adapting to a new dataset, +1.3 AP) and fine-tuned settings (Task 2: post training class adjustment, +5.7 AP). Larger gains are observed for the fine-tuned model, which can be attributed to the improved bounding box regression component, facilitating the classification task. Gains are more substantial for rare classes (+2.2 AP (zero-shot), +10.8 AP (fine-tuned)), which can be attributed to the removal of name ambiguity and the fact that these names are less common in pre-training datasets, leading to potentially poorer mapping with image content (e.g. ‘arctic’ shoe class). In contrast, frequent classes remain mostly stable (+0 AP (zero-shot), +1.2 (fine-tuned)), highlighting the quality of image to name mapping for this group. We provide more evidence of this in Section 4.3.

Open vocabulary experiments (Task 3), where class names are only learned on base classes, show modest gains in the zero-shot setting (+ 0.5 AP), but large improvements on a fine-tuned model (+ 4.5 AP), notably highlighting the importance of the localisation task. Due to the multi-label nature of LVIS, base and rare classes are not competing, facilitating the open vocabulary task. Learning base class names additionally reduces ambiguity with rare class names, improving rare class performance.

Sequential learning experiments (Task 4; training on base then rare classes) achieve comparable performance to directly training on all classes (35.0 AP (seq.) vs. 35.6 AP (all classes)), further highlighting how well suited our approach is for sequential adaptation. LVIS is a highly imbalanced dataset. While our 10% subsampling strategy reduces class imbalance, it is not entirely eliminated. This is evidenced by the improved rare class performance in the sequential setting (+3.3




Figure 3. Correlation between class rarity (log scale) and embedding similarity between learned and pre-defined class names. LVIS dataset.

(zero-shot), +2.1 (fine-tuned)). Fine-tuning the entire OWL-vit model on all classes (100% of the data) further evidences the impact of this imbalance and our model’s potential to address long tail issues. By reducing data requirements (only 10% achieves equivalent overall performance; 34.5 AP), we are able to train on a balanced data subset. In both fine-tuned and zero-shot settings, learning class names on 10% of the data achieves better or equivalent rare class performance (AP/r) than the fine-tuned model (-0.1 to +15.4 AP/r gains).

Finally, prompt tuning methods struggle to achieve good Task 1 performance in the zero-shot setting (-0.9 AP), but obtain better gains after fine-tuning (+2.5 AP). We note that prompt tuning performance is overall poorer than our strategy (-2.2 zero shot, -3.2 fine-tuned vs. ours all), which we hypothesise is partially linked to OWL-vit pre-training using fixed templates, and the diversity of object detection contexts.

CODA experiments, Task 3. Experimental results on the CODA 2.0 dataset are reported in Table 4. Following the setup proposed in [25], we report average precision (AP) on common classes (AP-Com), class agnostic AP on all classes (AP-A), class agnostic average recall on all classes (AR-A) and on corner classes (AR-Cor). We compare to state of the art self-driving models, trained in a semi-supervised setting on the SODA dataset [25]. The zero-shot OWL-vit model is not trained on any self-driving data, therefore achieves an overall poorer common class performance (0.10 AP-Com). However, the open-vocabulary setting allows this model to achieve very strong corner classes detection (0.45 AR-Cor). Our method largely improves performance for all metrics, with an overall performance (1.45 Sum) exceeding that of self-driving models (1.41 Sum), including our OWL-vit model fine-tuned on SODA (1.44 Sum). Fine-tuning OWL-vit on SODA, and a mixture of SODA and LVIS training data further improves results (1.44 and 1.59 Sum, respectively), and our strategy consistently increases all metrics (1.58 and 1.67 Sum, respectively). Common class performance remains below the reference models (-0.07 AP-Com), which we conjecture can be attributed to our relatively small self driving training data (only 10k images).



Dataset	Original name	Closest to new name (embedding similarity)	Example image
Appearance based Correction			
LVIS	arctic	boot (0.46), ski boot (0.30)	
CODA	tricycle	cart (0.62), rickshaw (0.67)	
CODA	misc	waterwheel (0.47), wheel (0.45)	
British to American English Correction			
LVIS	clothes hamper wall socket garbage postbox trousers	laundry basket (0.52), shopping basket (0.39) outlet (0.62), power outlet (0.60) trash can (0.47), trashcan (0.45) mailbox (0.47), telephone booth (0.46) clothes (0.42), pants (0.41)	
Object type first			
LVIS	bell pepper bridal gown horse buggy horse carriage celery cymbal	vegetable (0.55), pepper (0.51), bell pepper (0.51) gown (0.57), dress (0.57), bridal gown (0.55) horse-drawn vehicle (0.66), horse buggy (0.64) horse-drawn vehicle (0.68), horse carriage (0.65) green vegetables (0.50), celery (0.47) musical instrument (0.55), cymbal (0.46)	

Table 5. Interpretability results: changes to original class names.

4.3. Interpretability

One key advantage of the proposed approach is a beneficial effect on model interpretability; the method allows evaluation of the manner in which class names are modified, based on image content. Here, we aim to evaluate how object classes were renamed in our object detection experiments. This experiment can usefully inform us about model biases, sources of classification errors, potential mislabelling errors and the suitability of handcrafted names. For this purpose, we build a reference vocabulary of 3031 object names, using the following dataset class names (filtered for duplicates): LVIS (1203 classes), CODA (43 classes), Objects365 [40] (365 classes), ImageNet (1000 classes), and the THINGS [16] database (1854 concepts).

Individual word embeddings may not have very valuable meanings in terms of interpretability, especially considering that certain class names comprise multiple tokens. As a result, we compare learned representations to our reference vocabulary at the output of the text transformer encoder, using the shared prompt template context: ‘a photo of a’. More formally, for each word \mathbf{e}^l , in our list of learned embeddings, we compute the sentence representation $f(t(\mathbf{e}^l))$, where $t(\mathbf{e}^l)$ is the sentence constructed using our prompt prefix and word embedding \mathbf{e}^l . We repeat the procedure for each instance in our vocabulary, using the model’s pre-trained word embeddings \mathbf{e}^v , and obtain $f(t(\mathbf{e}^v))$. We then simply compute the cosine distance between all pairs of learned and reference word representations.

We carry out this process for models trained on the LVIS dataset (base model fine-tuned on base classes), and on the CODA 2.0 dataset (zero-shot model). In Table 5, we show examples of classes that were renamed towards a semantic meaning

closer to image content, and highlight some labelling errors. “Arctic” and “misc” are evidence of non representative class names. The “tricycle” class appears to have been mislabelled, with labelled vehicles looking closer to “rickshaws” and “carts”. We additionally report two interesting trends: 1) the model tends to adjust word embeddings from British English representations to American English, where the latter is expected to be the dominant language, on which large scale CLIP models were trained; 2) fine-grained category representations tend to be closer to supercategory word types (e.g. vegetable, musical instrument), suggesting that the model also leverages cross-class similarities towards the learning of representations for rare classes. The first observation notably highlights our method’s potential for cross-lingual model adaptation: given a dataset in one language, and a model trained in another language, we can learn optimal class names adapted to the pre-trained language model without any additional language translation efforts.

In Figure 3 we highlight the relationship connecting class rarity and embedding similarity between learned and original class names, on the LVIS dataset. As previously conjectured in Section 4.2, there is a clear correlation between these variables, highlighting how multi-modal models also suffer from common long tail related issues (i.e. poorer semantic mapping between rare words and their corresponding visual appearance). Additional results are available in our supplementary materials.

5. Conclusion

This work focuses on the concept of class name and resulting impact on performance of vision-language recognition models. Complementary to prompt tuning, which seeks to optimise the context of text queries, we propose to automatically learn optimal class names from visual content. Our solution can easily be integrated in classification and object detection models, achieving significant gains, with respect to the base model, in all settings. We additionally demonstrate promising interpretability properties by showing how class semantic meanings were adjusted, providing data and model analysis properties.

Learning class names can be viewed as an alternative to linear probing. The former approach, in contrast to probing, maintains key properties inherent to vision-language models (e.g. open-vocabulary). One observed limitation of our approach, in common with linear probing, is relatively weaker performance, in comparison with prompt tuning and image queries, under very low data regimes (i.e. 1-2 shots). Exploring a combination of learned text and image queries makes for a promising future avenue. While learning class names reduces dependency on prompt context, we still observe improvements using engineering prompts, in particular in open-vocabulary settings. As such, our method has the potential to further benefit from combination with prompt tuning techniques [50, 51], where prompts can provide key domain adaptation properties. Last but not least, learning class names has the potential to benefit class agnostic detection, via clustering and learning cluster semantics.

References

- [1] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *European conference on computer vision*, pages 446–461. Springer, 2014. [5](#)
- [2] Zhaowei Cai, Gukyeong Kwon, Avinash Ravichandran, Erhan Bas, Zhuowen Tu, Rahul Bhotika, and Stefano Soatto. X-detr: A versatile architecture for instance-wise vision-language tasks. *arXiv preprint arXiv:2204.05626*, 2022. [2](#), [4](#), [6](#)
- [3] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3606–3613, 2014. [5](#)
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [5](#)
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [5](#)
- [6] Yu Du, Fangyun Wei, Ziheng Zhang, Miaojing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14084–14093, 2022. [2](#), [3](#)
- [7] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pages 178–178. IEEE, 2004. [5](#)
- [8] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems*, 26, 2013. [2](#)
- [9] Andreas Furst, Elisabeth Rumetshofer, Viet Tran, Hubert Ramsauer, Fei Tang, Johannes Lehner, David Kreil, Michael Kopp, Günter Klambauer, Angela Bitto-Nemling, et al. Cloob: Modern hopfield networks with infoloob outperform clip. *arXiv preprint arXiv:2110.11316*, 2021. [2](#)
- [10] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermanto, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. [2](#), [3](#)
- [11] Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*, 2020. [3](#)
- [12] Chuanxing Geng, Sheng-jun Huang, and Songcan Chen. Recent advances in open set recognition: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3614–3631, 2020. [1](#)
- [13] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Zero-shot detection via vision and language knowledge distillation. *CoRR*, abs/2104.13921, 2021. [2](#), [4](#)
- [14] Agrim Gupta, Piotr Dollar, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. [6](#)
- [15] Jianhua Han, Xiwen Liang, Hang Xu, Kai Chen, Lanqing Hong, Jiageng Mao, Chaoqiang Ye, Wei Zhang, Zhenguo Li, Xiaodan Liang, et al. Soda10m: A large-scale 2d self/semi-supervised object detection dataset for autonomous driving. *arXiv preprint arXiv:2106.11118*, 2021. [6](#)
- [16] Martin N. Hebart, Adam H. Dickter, Alexis Kidder, Wan Y. Kwok, Anna Corriveau, Caitlin Van Wicklin, and Chris I. Baker. Things: A database of 1,854 object concepts and more than 26,000 naturalistic object images. *PLOS ONE*, 14:1–24, 10 2019. [8](#)
- [17] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Introducing eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. In *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 204–207, 2018. [5](#)
- [18] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. [2](#)
- [19] KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5830–5840, 2021. [1](#)
- [20] Armand Joulin, Laurens van der Maaten, Allan Jabri, and Nicolas Vasilache. Learning visual features from large weakly supervised data. In *European Conference on Computer Vision*, pages 67–84. Springer, 2016. [2](#)
- [21] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *European Conference on Computer Vision*, pages 105–124. Springer, 2022. [3](#)
- [22] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1780–1790, 2021. [1](#), [2](#), [4](#), [6](#)
- [23] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561, 2013. [5](#)
- [24] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. [3](#)
- [25] Kaican Li, Kai Chen, Haoyu Wang, Lanqing Hong, Chaoqiang Ye, Jianhua Han, Yukuai Chen, Wei Zhang, Chunjing Xu, Dit-Yan Yeung, et al. Coda: A real-world road corner case dataset for object detection in autonomous driving. *arXiv preprint arXiv:2203.07724*, 2022. [6](#), [7](#)
- [26] Liunian Harold Li*, Pengchuan Zhang*, Haotian Zhang*, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training. In *CVPR*, 2022. [2](#), [5](#), [6](#)
- [27] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image

- pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10965–10975, 2022. 4
- [28] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. 3
- [29] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208*, 2021. 2
- [30] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*, 2021. 3
- [31] Muhammad Maaz, Hanoona Bangalath Rasheed, Salman Hameed Khan, Fahad Shahbaz Khan, Rao Muhammad Anwer, and Ming-Hsuan Yang. Multi-modal transformers excel at class-agnostic object detection. *arXiv*, 2021. 2, 4
- [32] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. 06 2013. 5
- [33] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection with vision transformers. *arXiv preprint arXiv:2205.06230*, 2022. 1, 2, 4, 5, 6
- [34] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008. 5
- [35] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012. 5
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 1, 2, 3, 4
- [37] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 3
- [38] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18082–18091, 2022. 3
- [39] Timo Schick and Hinrich Schütze. It’s not just size that matters: Small language models are also few-shot learners. *arXiv preprint arXiv:2009.07118*, 2020. 3
- [40] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8429–8438, 2019. 8
- [41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2
- [42] Khurram Soomro, Amir Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, 12 2012. 5
- [43] Jianxiong Xiao, James Hays, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3485–3492, 2010. 5
- [44] Yuan Yao, Ao Zhang, Zhengyan Zhang, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. Cpt: Colorful prompt tuning for pre-trained vision-language models. *arXiv preprint arXiv:2109.11797*, 2021. 3
- [45] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14393–14402, 2021. 2
- [46] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. *CoRR*, abs/2111.07991, 2021. 2
- [47] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8552–8562, 2022. 3
- [48] Chen Zheng, Quan Guo, and Parisa Kordjamshidi. Cross-modality relevance for reasoning on language and vision. *arXiv preprint arXiv:2005.06035*, 2020. 2
- [49] Zexuan Zhong, Dan Friedman, and Danqi Chen. Factual probing is [mask]: Learning vs. learning to recall. *arXiv preprint arXiv:2104.05240*, 2021. 3
- [50] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16816–16825, 2022. 2, 3, 5, 8
- [51] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. 2, 3, 5, 6, 8