

---

# Enhancing CLIP with CLIP: Exploring Pseudolabeling for Limited-Label Prompt Tuning

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Fine-tuning vision-language models (VLMs) like CLIP to downstream tasks is  
2 often necessary to optimize their performance. However, a major obstacle is the  
3 limited availability of labeled data. We study the use of pseudolabels, i.e., heuristic  
4 labels for unlabeled data, to enhance CLIP via prompt tuning. Conventional pseu-  
5 dolabeling trains a model on labeled data and then generates labels for unlabeled  
6 data. VLMs’ zero-shot capabilities enable a “second generation” of pseudolabeling  
7 approaches that do not require task-specific training on labeled data. By using zero-  
8 shot pseudolabels as a source of supervision, we observe that learning paradigms  
9 such as semi-supervised, transductive zero-shot, and unsupervised learning can  
10 all be seen as optimizing the same loss function. This unified view enables the  
11 development of versatile training strategies that are applicable across learning  
12 paradigms. We investigate them on image classification tasks where CLIP exhibits  
13 limitations, by varying prompt modalities, e.g., textual or visual prompts, and  
14 learning paradigms. We find that (1) unexplored prompt tuning strategies that iter-  
15 atively refine pseudolabels consistently improve CLIP accuracy, by 19.5 points in  
16 semi-supervised learning, by 28.4 points in transductive zero-shot learning, and by  
17 15.2 points in unsupervised learning, and (2) unlike conventional semi-supervised  
18 pseudolabeling, which exacerbates model biases toward classes with higher-quality  
19 pseudolabels, prompt tuning leads to a more equitable distribution of per-class  
20 accuracy.

## 21 1 Introduction

22 Large pre-trained vision-language models (VLMs) [31, 43, 17] achieve remarkable accuracy without  
23 task-specific training but still require adaptation for optimal performance. Prompt-tuning [13, 18] is  
24 an approach to efficiently enhance VLMs performance on downstream tasks by learning inputs to the  
25 model. While learning prompts with a few labeled data can yield significant improvements [48, 2], a  
26 broader range of learning settings such as semi-supervised, transductive zero-shot, and unsupervised  
27 learning are still underexplored. All of these settings share access to unlabeled data, and the versatile  
28 zero-shot classification abilities of VLMs make pseudolabeling a natural approach to leveraging it.  
29 This paper investigates how the use of out-of-the-box pseudolabels assigned by CLIP can contribute  
30 to improving CLIP’s own performance. To this end, we conduct an extensive exploration of learning  
31 scenarios by varying prompt modalities, learning paradigms, and training strategies. We present  
32 empirical evidence showcasing the effectiveness of iterative prompt-training strategies that leverage  
33 CLIP-based pseudolabels, regardless of learning paradigms and prompt modalities, resulting in  
34 significant improvements in CLIP’s image classification performance across different settings.

35 Pseudolabels are heuristic labels assigned by a model to unlabeled data, which are leveraged to  
36 further train the model [20]. Successful training with pseudolabels relies on two factors: the quality

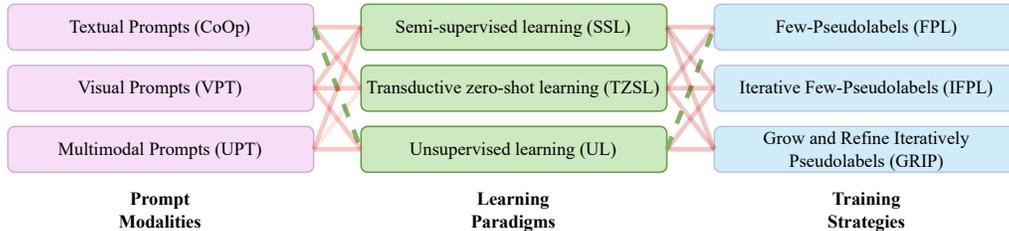


Figure 1: Our design space to explore the effect of leveraging pseudolabels in a unified way across prompt modalities, learning paradigms, and training strategies. The green (dashed) path has already been explored [15], while the red (solid) lines are the unexplored combinations for prompt tuning.

37 of the labels and how they are used during training. To address the first, conventional methods assign  
 38 labels to instances with high-confidence predictions [36]. For pseudolabeling using CLIP, Huang et  
 39 al. propose to select the most confident samples for each class [15], mitigating CLIP’s bias [38] and  
 40 miscalibration [22] (see Section 3). To assign pseudolabels, we rely on this approach and address the  
 41 second point by exploring how to make the best use of them. We design a broad space of analysis  
 42 considering three dimensions: prompt modalities, which are the model inputs we learn; learning  
 43 paradigms, which define the data we have available; and training strategies, which describe the  
 44 process used to optimize performance (Figure 1).

45 Research on prompt tuning has demonstrated that training strategies used for learning prompts in  
 46 one modality can be transferred to learning prompts in a different modality. For instance, Visual  
 47 Prompt Tuning [18] was originally designed to effectively fine-tune large vision models but can be  
 48 adapted to efficiently fine-tune CLIP using the same training strategy as standard textual prompt  
 49 tuning [48, 34, 44]. On the contrary, different learning paradigms with limited labeled data typically  
 50 require distinct approaches specifically tailored to extract information from the available data [27, 12].  
 51 However, we observe that this changes by using VLM’s generated pseudolabels. Unlike conventional  
 52 pseudolabeling approaches that bootstrap off labeled data and are used as semi-supervised learning  
 53 techniques [36, 3, 40], VLMs can generate pseudolabels in any learning setting. This offers a  
 54 significant advantage, expanding the scope of pseudolabeling beyond semi-supervised learning, and  
 55 making it a promising approach for other settings, such as transductive zero-shot and unsupervised  
 56 learning. By using CLIP-based pseudolabels as a source of supervision, we can view these settings as  
 57 optimizing the same loss function, which is simply a weighted sum of the errors on labeled data, if  
 58 available, and pseudolabeled data. Given that we can express different settings as the same problem,  
 59 we can propose training strategies, i.e., the way of using pseudolabels, that suit them all.

60 By standardizing the training strategies across various prompt modalities and learning settings, we  
 61 can conduct experiments on different applications of pseudolabels for various combinations of prompt  
 62 modalities, learning paradigms, and training strategies, as illustrated in Figure 1. To the best of  
 63 our knowledge, only one potential path has been explored thus far; specifically, fine-tuning textual  
 64 prompts in an unsupervised learning context using a few pseudolabels [15]. Rather than relying  
 65 on a fixed set of pseudolabels, we propose iterative training techniques that allow for the ongoing  
 66 refinement and expansion of the pool of pseudolabeled data used during training. In this way, with  
 67 each iteration, we progressively enhance CLIP’s pseudolabeling ability, allowing us to extend the set  
 68 of pseudolabeled data while maintaining the high quality of the initial pseudolabels, which tend to be  
 69 reliable.

70 We conduct experiments on six tasks where CLIP has been observed to underperform [31], such as  
 71 satellite-image classification, flower-species identification, and texture-image recognition, among  
 72 others. Our findings reveal that iterative approaches effectively fine-tune prompts irrespective of  
 73 their modality and learning paradigms. Recent studies have identified the “Matthew effect” as a  
 74 potential issue for semi-supervised models that use pseudolabels [49, 38]. This phenomenon causes  
 75 models to perform well on classes with accurate pseudolabels but poorly on those with inaccurate  
 76 ones, thereby reinforcing the model’s original bias towards certain classes. Our analysis reveals that  
 77 using pseudolabels generated by CLIP for prompt-tuning with iterative strategies not only improves  
 78 CLIP’s overall performance but also corrects its natural bias towards certain classes.

79 We summarize the main takeaways of our work:

- 80 • General purpose zero-shot learners used as general purpose pseudolabelers open the opportunity to develop training strategies that leverage pseudolabeled data beyond semi-supervised learning. We point out that different learning paradigms, such as semi-supervised, transductive zero-shot, and unsupervised learning, can be all considered as special cases of a single objective function, by using pseudolabels as a source of supervision.
- 81
- 82
- 83
- 84
- 85 • We demonstrate that simple iterative training strategies for refining pseudolabels are highly effective approaches for limited-label prompt tuning. In fact, regardless of the prompt modality and learning setting, these strategies improve CLIP, by on average 19.5 points in semi-supervised learning, 28.4 in transductive zero-shot learning, and 15.2 in unsupervised learning.
- 86
- 87
- 88
- 89
- 90 • We show that prompts learned with iterative strategies help mitigate the "rich get richer, poor get poorer" effect observed in semi-supervised approaches leveraging pseudolabels. By redistributing the quality of pseudolabels across different classes, we observe a "Robin Hood effect" where the extremely rich classes' accuracy stays the same or decreases, while poorer classes get richer, leading to a more equitable distribution of per-class accuracy.
- 91
- 92
- 93
- 94

## 95 2 Background and related work

96 **Vision-language models** Vision-language models such as CLIP [31], ALIGN [17], and Florence [43] are models that align images and text. We focus on CLIP, which is composed of two components: a text encoder,  $\psi$ , and an image encoder,  $\phi$ , which are jointly trained using a contrastive loss to learn a multi-modal embedding space which aligns the representations of similar text and image inputs. This pre-training enables CLIP to perform zero-shot image classification. Given an image  $x$  and a set of classes  $\mathcal{Y} = \{y_1, \dots, y_C\}$ , CLIP classifies  $x$  by measuring the similarity between the image representation  $z = \phi(x)$  and each class representation  $w_i = \psi(\pi_i)$ , based on their cosine distance in the shared embedding space. Here,  $\pi_i$  is a natural language prompt such as "a photo of a [CLASS<sub>*i*</sub>]", where CLASS<sub>*i*</sub> is the specific class name, such as "orange dahlia," "forest" or "Boeing 737". The image  $x$  gets assigned to the class with the highest similarity score. In this work, we study how to learn better prompts that enhance CLIP by leveraging pseudolabels.

107 **Prompt tuning** Prompt tuning is a technique that enhances the practical application of large pre-trained models like CLIP [31] and GPT [32, 7]. It involves providing task-specific information to the model during inference through textual or visual inputs, leading to improved performance on downstream tasks [1, 33, 6, 7]. While discrete prompts are manually crafted natural language descriptions of classes that guide the model, they may not yield optimal results [47]. Soft prompting [21, 24], on the other hand, optimizes prompts as continuous vectors. These can be optimized by backpropagating through the frozen pre-trained model, resulting in better performance. Soft prompts can be learned for various modalities, e.g., text or image, [48, 13, 17, 2, 44, 19] and applications [34, 27, 12, 28] by training on a small number of labeled examples per class. If only unlabeled data is accessible, it is possible to learn textual soft prompts by leveraging CLIP-based pseudolabels [15]. Expanding on this concept, we further investigate the use of pseudolabels across a broader range of prompt modalities and learning approaches, and we introduce novel training strategies to leverage pseudolabels more effectively.

120 **Learning from pseudolabels** Pseudolabeling is the practice of assigning labels to unlabeled data based on the prediction of a model [20]. Then, pseudolabels are used to improve the performance of the model itself. There are different ways to obtain and use pseudolabels and each impacts the final predictions of the model [41, 45, 16, 35]. Some approaches use confidence thresholds [36, 3, 40] and others average predictions from multiple augmentations [4]. Pseudolabeling is a semi-supervised learning technique, and it is rarely used in transductive zero-shot learning [42, 5, 25]. Applying such techniques requires a few labeled examples related to the target task to learn a baseline model capable of pseudolabeling. However, this limitation has been overcome by VLMs, which are capable of pseudolabeling examples without task-specific training. The conventional pseudolabeling scheme based on confidence threshold is not effective if we assign pseudolabels based on CLIP. In fact CLIP is miscalibrated [22] and has imbalanced predictions [38] which may induce noise in the pseudolabels. An alternative approach selects the top- $K$  most confident examples per class to improve performance [15]. In our analysis, we rely on this scheme (Section 3).

### 133 3 Design space

134 Our analysis encompasses the design space consisting of various combinations of prompt modalities,  
135 learning paradigms, and training strategies (Figure 1). Within this space, two key components remain  
136 constant: the pseudolabeling scheme and a unified loss function. This section begins by introducing  
137 these components and subsequently delves into a comprehensive discussion of each dimension within  
138 the design space to be explored.

139 **Pseudolabeling scheme** The use of CLIP to generate pseudolabels has been investigated in [15].  
140 Given unlabeled data  $X_u$  with target classes  $\{y_1, \dots, y_C\}$ , the goal is to assign labels to data points in  
141 which the model is most confident. Typically, pseudo labeling schemes use a confidence threshold  
142 ( $P(y|x) > \tau$ ) to select instances to pseudolabel. However, this approach does not work well for CLIP  
143 due to its miscalibration [22] and imbalanced predictions [38]. Instead, one can use a top- $K$  pseudo  
144 labeling approach, where the top- $K$  most confident examples per class are used as pseudolabeled  
145 data [15]. This approach guarantees an equal distribution of pseudolabeled training samples for  
146 each class, effectively addressing the natural bias in CLIP’s pseudolabels [38]. Moreover, it enables  
147 leveraging the similarity signals from CLIP’s probabilities, irrespective of the distribution of its  
148 predictions.

149 This top- $K$  pseudolabeling scheme is applicable to unlabeled data, regardless of the availability of  
150 labeled data. As a result, we can extend the use of pseudolabels to any learning setting that involves  
151 unlabeled data. We observe that by treating pseudolabeled examples as true labeled data, we can view  
152 all learning settings as optimizing the same objective function.

153 **Unified objective function** Consider a  $C$ -class image classification task, where  $X_L$  and  $Y_L$  rep-  
154 resent the image representations and labels of the labeled data, and  $X_U$  and  $\tilde{Y}_U$  denote the image  
155 representations and pseudolabels for the unlabeled data. We define a loss function that combines two  
156 cross-entropy losses, one accounting for the error on the labeled data points and the other accounting  
157 for the error on pseudolabeled data:

$$\mathcal{L} = \gamma \mathcal{L}_{CE}(X_L, Y_L) + \lambda \mathcal{L}_{CE}(X_U, \tilde{Y}_U)$$

158 where  $\gamma$  and  $\lambda$  define the training balance between the errors on labeled and pseudolabeled data.

#### 159 3.1 Prompt modalities

160 Learning prompts is the process of training a set of vectors  $P = [p]_1 \dots [p]_K$  that are prepended  
161 to the textual or visual inputs of the encoders within the CLIP architecture. By prepending these  
162 vectors to specific inputs, we can learn *textual* prompts, *visual* prompts, or *multimodal* prompts  
163 when applying a set of vectors to both inputs simultaneously. We provide a technical and detailed  
164 explanation in Appendix A.1.

165 In our exploration, we consider all three types of prompts. The efficacy of prompts can vary depending  
166 on the task. Text prompt tuning may be most beneficial when image features are well-separated  
167 by class but may not be aligned with the corresponding textual prompt. Visual prompts rearrange  
168 the image features within the projection space, and it has the potential to improve CLIP when the  
169 pre-trained image features are not well separated by class. Finally, multimodal prompts allows for  
170 beneficial interaction between the two separate modalities, which might lead to both separable visual  
171 features, and text classifiers that are well-aligned with the corresponding visual features.

#### 172 3.2 Learning paradigms

173 By adjusting the values of parameters  $\gamma$  and  $\lambda$  and using the appropriate sets of labeled and pseudola-  
174 beled data, the unified objective loss can be customized for each learning paradigm.

175 **Semi-supervised learning** In the semi-supervised learning (SSL) scenario we have access to a  
176 limited number of labeled data for all the target classes  $D_L = \{(x, y)\}$  where  $x$  is an input feature and  
177  $y \in \mathcal{Y} = [C]$  is the corresponding label. In addition, we have access to unlabeled data  $X_U = \{x\}$ ,

178 where  $x$  is an image in the target domain  $\mathcal{Y}$ . From  $X_U$ , we get  $\mathcal{D}_{PL} = \{(x, \tilde{y})\}$ , where  $\tilde{y} \in [C]$   
179 is  $x$ 's pseudolabel. When using the unified loss in this setting, we set  $\gamma$  to  $|\mathcal{D}_{PL}|/|\mathcal{D}_L|$ . As  $|\mathcal{D}_L|$   
180 is much smaller than  $|\mathcal{D}_{PL}|$ ,  $\gamma$  acts as an upweighting factor for the few-labeled instances, thus  
181 counterbalancing the learning effect of pseudolabels ( $\lambda=1$ ).

182 **Transductive zero-shot learning** In transductive zero-shot learning (TZSL), we are provided with  
183 labeled data  $D_L = \{(x, y)\}$  for some target classes  $S$  (referred to as *seen* classes), where  $x$  represents  
184 input features, and  $y \in [S]$  is the corresponding label. Additionally, we have access to unlabeled  
185 data  $X_U = \{x\}$  for a disjoint set of classes  $U$  (referred to as *unseen* classes). Using  $X_U$ , we obtain  
186  $\mathcal{D}_{PL} = (x, \tilde{y})$ , where  $\tilde{y} \in [U]$  denotes the pseudolabels for  $x$ . The value of  $\lambda$  in the unified loss is  
187 set to  $|\mathcal{D}_L|/|\mathcal{D}_{PL}|$ , which makes the weight of the pseudolabel loss equivalent to that of the labeled  
188 data ( $\gamma = 1$ ). This is necessary because an imbalance in the number of labeled and pseudolabeled  
189 samples can result in a skewed training distribution, leading to better performance on seen classes  
190 while the performance on unseen classes may either remain stagnant or degrade. Studying this setting  
191 is interesting beyond transductive zero-shot learning. In fact, it has the potential to generalize to  
192 scenarios where the target task involves unseen classes, while the seen classes consist of auxiliary  
193 labeled data from the same domain but different task [30].

194 **Unsupervised learning** In the unsupervised learning (UL) setting, we have access only to unlabeled  
195 data  $X_U = \{x\}$ , from which we obtain  $\mathcal{D}_{PL} = (x, \tilde{y})$ , where  $\tilde{y} \in [C]$  denotes the pseudolabel for  $x$ .  
196 In this case,  $\gamma$  is set to 0, as there is no labeled data, and  $\lambda = 1$ . The use of this setting was initially  
197 explored in [15], who leveraged a few pseudolabels per class to learn textual prompts. In this paper,  
198 we build on their work by investigating a variety of training strategies and prompt modalities.

199 **Supervised learning** In supervised learning (SL), we are only provided with labeled data  $D_L =$   
200  $(x, y)$ , where  $x$  represents an input feature, and  $y \in [C]$  is the corresponding label. If we set  $\lambda$  to 0,  
201 the unified loss function is equivalent to the objective functions of default prompt-tuning approaches  
202 that optimize the prompts using a few labeled instances per target class. This setting is not strictly  
203 part of our design space. However, we will refer to it to define baselines in Section 4.

### 204 3.3 Training strategies

205 The unified objective function enables the development of training strategies broadly applicable  
206 across various learning paradigms. We explore three distinct learning strategies to effectively utilize  
207 pseudolabels in this context. The first strategy utilizes pseudolabels in a static manner. The other  
208 two strategies, which are novel and proposed here for the first time for prompt tuning, involve the  
209 dynamic use of pseudolabeled data.

210 **Few-pseudolabels (FPL)** We select  $K$  pseudolabels per target class, resulting in a pseudolabeled  
211 dataset of size  $K \cdot C$ . We learn the prompts by minimizing the objective function via backpropagation  
212 through CLIP's encoders. This strategy aligns with Unsupervised Prompt Learning (UPL) in [15]. We  
213 refer to it as few-pseudolabels (FPL) to encompass its applicability for learning prompts of diverse  
214 modalities across learning paradigms.

215 **Iterative Refinement of FPL (IFPL)** Similar to FPL, we assign  $K$  pseudolabels per target class.  
216 After the first iteration, we recompute the set of  $C \cdot K$  pseudolabels. Then, we conduct  $I$  iterations of  
217 training-pseudolabeling reinitializing the prompt every time. With this iterative approach, if training  
218 with the initial pseudolabel set leads to an improvement in the model's performance, the model itself  
219 can become a more effective pseudolabeler, refining the pseudolabels in each subsequent iteration.

220 **Grow and Refine Iteratively Pseudolabels (GRIP)** Although the iterative refinement of FPL can  
221 improve the quality of the  $K$  pseudolabels used for training, it still limits learning to a few examples  
222 per target class. To overcome this constraint, we expand the set of pseudolabels at each iteration such  
223 that, at the  $i$ th iteration, we use  $i/I$ -th of the unlabeled data. The rationale behind this strategy is that  
224 as the model's accuracy in generating pseudolabels improves, we can increase the total number of  
225 pseudolabels without introducing excessive noise.

## 226 4 Experiments

227 We explore the design space outlined in Section 3 to understand the effectiveness of leveraging  
228 pseudolabels for limited-label prompt tuning. We show that (1) iterative strategies significantly  
229 improve CLIP’s performance across prompt modalities and learning settings, (2) using CLIP-based  
230 pseudolabels with iterative strategies induces a more equitable distribution of per-class accuracy.

231 **Datasets** We conduct the analysis on six tasks, covering specialized and fine-grained domains,  
232 where CLIP shows deficiencies [31]. We call this set of tasks FRAMED, and it includes  
233 Flowers102 [29], RESICS45 [9], FGVC-Aircraft [26], MNIST [11], EuroSAT [14], and DTD [10].  
234 For each dataset we use the training and test splits provided in [23]. For the transductive zero-shot  
235 learning setting we randomly generate three splits of seen and unseen classes with a 62-38 ratio.  
236 Further details are in Appendix A.2.

237 **Baselines** To evaluate the effectiveness of the training strategies described in Section 3.3, we  
238 compare the performance of CLIP when queried with the learned soft prompts to CLIP zero-shot  
239 with default prompts such as “a photo of a [CLASS].” In addition, we compare with default  
240 supervised prompt-tuning baselines, for which we only use the available labeled data: CoOp [48]  
241 for textual prompts, VPT [18] for visual prompts, and UPT [44] for multimodal prompts. We defer  
242 to Appendix A.1 the technical details of these methods.

243 **Evaluation metrics** We assess the performance of each method by measuring the accuracy of the  
244 test set, averaging the results over five runs. In the case of TRZSL, we report the harmonic mean to  
245 account for potentially imbalanced performance between seen and unseen classes [39].

246 **Training settings** For all experiments, datasets, and learning strategies, we use ViT-B/32 as the  
247 vision backbone. For both visual and textual prompt learning, we set the prefix size to 16 [48, 18].  
248 Multimodal prompts have length 8 [44]. We use SGD as the optimizer and train for 150 epochs.  
249 We utilize 5 warmup epochs at a learning rate of 0.0001, and then set the learning rate to  $l$ , which  
250 is decayed by the cosine annealing rule. For textual and visual prompt learning,  $l = .1$ , while for  
251 multimodal prompt learning,  $l = .01$ . In SSL, we use 2 labeled samples per class to assess the impact  
252 of pseudolabels in the scenario of very few labeled data and abundant unlabeled data. FPL and IFPL  
253 have the number of pseudolabels per class fixed to 16 [15]. The number of iterations  $I = 10$ .

### 254 4.1 Exploring the design space

255 **GRIP consistently enhances CLIP across prompt modalities and learning settings** Table 1  
256 reports the performance of GRIP, the best performing among the training strategies in Section 3.3,  
257 compared to CLIP and prompt-tuning baselines. Overall, GRIP consistently improves the performance  
258 of CLIP and the baselines across prompt modalities and learning settings. By tuning textual prompts,  
259 the average improvement over CLIP is 20.7 points in SSL, 14.9 in UL, and 32.4 in TRZSL, while  
260 the improvement on CoOp is 9.6 points in SSL, and 26.6 in TRZSL. Similar results for the visual  
261 prompts show that GRIP improves CLIP by 18.2 points in SSL, 15.7 in UL, and 30.8 in TRZSL, and  
262 VPT by 12.9 points in SSL, and 20.8 in TRZSL. We note that CoOp and VPT applied to the SSL  
263 setting correspond to learning only on the labeled data, and we do not run them in the UL setting as  
264 there is no labeled data. Results are similar for multimodal prompts. We defer them to Appendix A.3,  
265 due to space constraints.

266 **Unsupervised learning is equivalent or more robust than learning with very few shots** The  
267 accuracy of GRIP when applied to the fully unsupervised setting is either higher or equivalent to  
268 the accuracy of VPT, which is trained using two labeled instances per class (Table 1). This shows  
269 that pseudolabeled data can substitute very few labeled examples for prompt tuning. However, the  
270 significant improvement of GRIP over CoOp and VPT in the semi-supervised setting (see Table 1)  
271 suggests that leveraging unlabeled data through pseudolabeling is advantageous in scenarios where  
272 labeled data is scarce but there is an abundance of unlabeled data.

273 **Transductive zero-shot learning effectively transfers knowledge** In the TRZSL setting,  
274 GRIP improves over CLIP and the baselines by a large margin (Table 1). Figure 2  
275 displays the balance of seen and unseen classes of each method alongside its accuracy.

Textual prompts									
Method	Flowers102			RESICS45			FGVCAircraft		
	SSL	UL	TRZSL	SSL	UL	TRZSL	SSL	UL	TRZSL
CLIP	63.67 <sub>0.00</sub>	-	63.40 <sub>0.00</sub>	54.48 <sub>0.00</sub>	-	54.46 <sub>0.00</sub>	17.58 <sub>0.00</sub>	-	17.86 <sub>0.00</sub>
CoOp	76.76 <sub>1.11</sub>	-	63.22 <sub>0.02</sub>	58.53 <sub>10.81</sub>	-	63.37 <sub>0.02</sub>	14.91 <sub>3.22</sub>	-	21.70 <sub>0.03</sub>
GRIP	<b>83.60</b> <sub>0.68</sub>	<b>69.84</b> <sub>1.06</sub>	<b>86.26</b> <sub>0.00</sub>	<b>74.11</b> <sub>0.68</sub>	<b>70.55</b> <sub>0.88</sub>	<b>81.07</b> <sub>0.00</sub>	16.98 <sub>0.82</sub>	15.22 <sub>0.71</sub>	<b>26.08</b> <sub>0.00</sub>
$\Delta$ CLIP	$\uparrow$ 19.93	$\uparrow$ 6.17	$\uparrow$ 22.86	$\uparrow$ 19.63	$\uparrow$ 16.07	$\uparrow$ 26.61	$\downarrow$ 0.6	$\downarrow$ 2.36	$\uparrow$ 8.22
$\Delta$ CoOp	$\uparrow$ 6.84	-	$\uparrow$ 23.04	$\uparrow$ 15.58	-	$\uparrow$ 17.70	$\uparrow$ 2.07	-	$\uparrow$ 4.38
Visual prompts									
Method	Flowers102			RESICS45			FGVCAircraft		
	SSL	UL	TRZSL	SSL	UL	TRZSL	SSL	UL	TRZSL
CLIP	63.67 <sub>0.00</sub>	-	63.40 <sub>0.00</sub>	54.48 <sub>0.00</sub>	-	54.46 <sub>0.00</sub>	17.58 <sub>0.00</sub>	-	17.86 <sub>0.00</sub>
VPT	63.73 <sub>1.52</sub>	-	64.71 <sub>0.00</sub>	60.80 <sub>1.65</sub>	-	67.06 <sub>0.00</sub>	17.76 <sub>0.68</sub>	-	<b>26.69</b> <sub>0.00</sub>
GRIP	<b>67.95</b> <sub>1.2</sub>	63.09 <sub>0.55</sub>	<b>77.18</b> <sub>0.00</sub>	<b>71.22</b> <sub>0.77</sub>	<b>68.43</b> <sub>0.61</sub>	<b>82.19</b> <sub>0.00</sub>	<b>19.43</b> <sub>0.5</sub>	17.51 <sub>0.61</sub>	26.42 <sub>0.00</sub>
$\Delta$ CLIP	$\uparrow$ 4.28	$\downarrow$ 0.58	$\uparrow$ 13.78	$\uparrow$ 16.74	$\uparrow$ 13.95	$\uparrow$ 27.73	$\uparrow$ 1.85	$\downarrow$ 0.07	$\uparrow$ 8.56
$\Delta$ VPT	$\uparrow$ 4.22	-	$\uparrow$ 12.47	$\uparrow$ 10.42	-	$\uparrow$ 15.13	$\uparrow$ 1.67	-	$\downarrow$ 0.27
Visual prompts									
Method	Flowers102			RESICS45			FGVCAircraft		
	SSL	UL	TRZSL	SSL	UL	TRZSL	SSL	UL	TRZSL
CLIP	63.67 <sub>0.00</sub>	-	63.40 <sub>0.00</sub>	54.48 <sub>0.00</sub>	-	54.46 <sub>0.00</sub>	17.58 <sub>0.00</sub>	-	17.86 <sub>0.00</sub>
VPT	63.73 <sub>1.52</sub>	-	64.71 <sub>0.00</sub>	60.80 <sub>1.65</sub>	-	67.06 <sub>0.00</sub>	17.76 <sub>0.68</sub>	-	<b>26.69</b> <sub>0.00</sub>
GRIP	<b>67.95</b> <sub>1.2</sub>	63.09 <sub>0.55</sub>	<b>77.18</b> <sub>0.00</sub>	<b>71.22</b> <sub>0.77</sub>	<b>68.43</b> <sub>0.61</sub>	<b>82.19</b> <sub>0.00</sub>	<b>19.43</b> <sub>0.5</sub>	17.51 <sub>0.61</sub>	26.42 <sub>0.00</sub>
$\Delta$ CLIP	$\uparrow$ 4.28	$\downarrow$ 0.58	$\uparrow$ 13.78	$\uparrow$ 16.74	$\uparrow$ 13.95	$\uparrow$ 27.73	$\uparrow$ 1.85	$\downarrow$ 0.07	$\uparrow$ 8.56
$\Delta$ VPT	$\uparrow$ 4.22	-	$\uparrow$ 12.47	$\uparrow$ 10.42	-	$\uparrow$ 15.13	$\uparrow$ 1.67	-	$\downarrow$ 0.27
Visual prompts									
Method	Flowers102			RESICS45			FGVCAircraft		
	SSL	UL	TRZSL	SSL	UL	TRZSL	SSL	UL	TRZSL
CLIP	63.67 <sub>0.00</sub>	-	63.40 <sub>0.00</sub>	54.48 <sub>0.00</sub>	-	54.46 <sub>0.00</sub>	17.58 <sub>0.00</sub>	-	17.86 <sub>0.00</sub>
VPT	63.73 <sub>1.52</sub>	-	64.71 <sub>0.00</sub>	60.80 <sub>1.65</sub>	-	67.06 <sub>0.00</sub>	17.76 <sub>0.68</sub>	-	<b>26.69</b> <sub>0.00</sub>
GRIP	<b>67.95</b> <sub>1.2</sub>	63.09 <sub>0.55</sub>	<b>77.18</b> <sub>0.00</sub>	<b>71.22</b> <sub>0.77</sub>	<b>68.43</b> <sub>0.61</sub>	<b>82.19</b> <sub>0.00</sub>	<b>19.43</b> <sub>0.5</sub>	17.51 <sub>0.61</sub>	26.42 <sub>0.00</sub>
$\Delta$ CLIP	$\uparrow$ 4.28	$\downarrow$ 0.58	$\uparrow$ 13.78	$\uparrow$ 16.74	$\uparrow$ 13.95	$\uparrow$ 27.73	$\uparrow$ 1.85	$\downarrow$ 0.07	$\uparrow$ 8.56
$\Delta$ VPT	$\uparrow$ 4.22	-	$\uparrow$ 12.47	$\uparrow$ 10.42	-	$\uparrow$ 15.13	$\uparrow$ 1.67	-	$\downarrow$ 0.27
Visual prompts									
Method	Flowers102			RESICS45			FGVCAircraft		
	SSL	UL	TRZSL	SSL	UL	TRZSL	SSL	UL	TRZSL
CLIP	63.67 <sub>0.00</sub>	-	63.40 <sub>0.00</sub>	54.48 <sub>0.00</sub>	-	54.46 <sub>0.00</sub>	17.58 <sub>0.00</sub>	-	17.86 <sub>0.00</sub>
VPT	63.73 <sub>1.52</sub>	-	64.71 <sub>0.00</sub>	60.80 <sub>1.65</sub>	-	67.06 <sub>0.00</sub>	17.76 <sub>0.68</sub>	-	<b>26.69</b> <sub>0.00</sub>
GRIP	<b>67.95</b> <sub>1.2</sub>	63.09 <sub>0.55</sub>	<b>77.18</b> <sub>0.00</sub>	<b>71.22</b> <sub>0.77</sub>	<b>68.43</b> <sub>0.61</sub>	<b>82.19</b> <sub>0.00</sub>	<b>19.43</b> <sub>0.5</sub>	17.51 <sub>0.61</sub>	26.42 <sub>0.00</sub>
$\Delta$ CLIP	$\uparrow$ 4.28	$\downarrow$ 0.58	$\uparrow$ 13.78	$\uparrow$ 16.74	$\uparrow$ 13.95	$\uparrow$ 27.73	$\uparrow$ 1.85	$\downarrow$ 0.07	$\uparrow$ 8.56
$\Delta$ VPT	$\uparrow$ 4.22	-	$\uparrow$ 12.47	$\uparrow$ 10.42	-	$\uparrow$ 15.13	$\uparrow$ 1.67	-	$\downarrow$ 0.27
Visual prompts									
Method	Flowers102			RESICS45			FGVCAircraft		
	SSL	UL	TRZSL	SSL	UL	TRZSL	SSL	UL	TRZSL
CLIP	63.67 <sub>0.00</sub>	-	63.40 <sub>0.00</sub>	54.48 <sub>0.00</sub>	-	54.46 <sub>0.00</sub>	17.58 <sub>0.00</sub>	-	17.86 <sub>0.00</sub>
VPT	63.73 <sub>1.52</sub>	-	64.71 <sub>0.00</sub>	60.80 <sub>1.65</sub>	-	67.06 <sub>0.00</sub>	17.76 <sub>0.68</sub>	-	<b>26.69</b> <sub>0.00</sub>
GRIP	<b>67.95</b> <sub>1.2</sub>	63.09 <sub>0.55</sub>	<b>77.18</b> <sub>0.00</sub>	<b>71.22</b> <sub>0.77</sub>	<b>68.43</b> <sub>0.61</sub>	<b>82.19</b> <sub>0.00</sub>	<b>19.43</b> <sub>0.5</sub>	17.51 <sub>0.61</sub>	26.42 <sub>0.00</sub>
$\Delta$ CLIP	$\uparrow$ 4.28	$\downarrow$ 0.58	$\uparrow$ 13.78	$\uparrow$ 16.74	$\uparrow$ 13.95	$\uparrow$ 27.73	$\uparrow$ 1.85	$\downarrow$ 0.07	$\uparrow$ 8.56
$\Delta$ VPT	$\uparrow$ 4.22	-	$\uparrow$ 12.47	$\uparrow$ 10.42	-	$\uparrow$ 15.13	$\uparrow$ 1.67	-	$\downarrow$ 0.27

Table 1: For each learning paradigm, we compare the accuracy of GRIP with CLIP zero-shot (ViT-B/32), CoOp, and VPT. Results are for SSL, UL, and TRZSL on FRAMED. We average the accuracy on 5 seeds and report the standard deviation.  $\Delta$  METHOD is the difference between the accuracy of GRIP and METHOD. We note that for UL we can not apply CoOp and VPT since no labeled data is available.

276 The *class balance* is  $(acc_{unseen} - acc_{seen})/acc_{seen}$ ,  
277 where values close to zero indicate a good balance, neg-  
278 ative values indicate better accuracies for seen classes,  
279 and positive values indicate better accuracies for unseen  
280 classes. Methods employing an iterative usage of pseudola-  
281 bels maintain a good balance, as opposed to CoOp/VPT  
282 and FPL. This balance in accuracy is likely a combined  
283 effect of the quality of the pseudolabels and the transfer of  
284 knowledge from the seen to the unseen classes. The latter  
285 point holds significant relevance as it can imply that even  
286 if we only possess unlabeled data for a specific target task,  
287 we can still utilize labeled data from related classes [30]  
288 within the same domain to enhance CLIP’s performance.

289 **There is a trade-off between quality and quantity of**  
290 **pseudolabels**

291 Table 2 shows the performance of CLIP  
292 employing prompts learned with different training strategies, all leveraging pseudolabels (Section 3.3).  
293 Iterative strategies are more effective than FPL which, similarly to [15], uses a static set of a few  
294 pseudolabels for one iteration. On Flowers102, RESICS45, and DTD, IFPL improves on average  
295 FPL by 5.6 points in SSL, 1.7 in UL, and 5.6 in TRZSL. GRIP boosts the performance even more by  
296 on average 7.8 points in SSL, 3.1 in UL, and 9.7 in TRZSL. Results on the other tasks are comparable  
297 or larger and we report them in Appendix A.3 due to space constraints.

297 Figure 3 shows the progression of pseudolabels quality for the iterative learning of textual prompts.  
298 IFPL maintains a fixed set of 16 pseudolabels, improving their quality with each iteration. On  
299 the other hand, GRIP expands pseudolabels by incorporating an additional decile of unlabeled

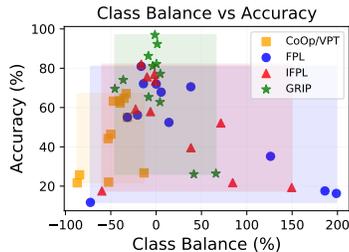


Figure 2: Balance of seen and unseen accuracies vs. model’s overall accuracy. Points close to 0 indicate a good balance. Negatives represent better accuracy for the seen classes.

Textual prompts									
Method	Flowers102			RESICS45			DTD		
	SSL	UL	TRZSL	SSL	UL	TRZSL	SSL	UL	TRZSL
FPL	75.96 <sub>0.74</sub>	65.67 <sub>0.23</sub>	80.97 <sub>0.00</sub>	68.13 <sub>0.55</sub>	63.07 <sub>0.38</sub>	72.11 <sub>0.00</sub>	37.10 <sub>5.45</sub>	44.96 <sub>0.55</sub>	46.3 <sub>0.03</sub>
IFPL	78.68 <sub>0.75</sub>	<b>69.56</b> <sub>1.05</sub>	82.08 <sub>0.00</sub>	70.52 <sub>1.12</sub>	64.11 <sub>0.98</sub>	75.51 <sub>0.00</sub>	<b>55.24</b> <sub>0.97</sub>	<b>47.77</b> <sub>1.15</sub>	59.14 <sub>0.02</sub>
GRIP	<b>83.60</b> <sub>0.48</sub>	<b>69.84</b> <sub>1.06</sub>	<b>86.26</b> <sub>0.00</sub>	<b>74.11</b> <sub>0.68</sub>	<b>70.55</b> <sub>0.88</sub>	<b>81.07</b> <sub>0.00</sub>	<b>56.07</b> <sub>0.85</sub>	<b>46.09</b> <sub>1.06</sub>	<b>65.30</b> <sub>0.01</sub>
Δ IFPL	↑ 2.72	↑ 3.89	↑ 1.11	↑ 2.39	↑ 1.07	↑ 3.4	↑ 18.14	↑ 2.81	↑ 12.84
Δ GRIP	↑ 7.64	↑ 4.17	↑ 5.29	↑ 5.89	↑ 7.48	↑ 8.96	↑ 18.97	↑ 1.13	↑ 19.00
Visual prompts									
FPL	67.03 <sub>0.65</sub>	<b>65.50</b> <sub>0.41</sub>	71.94 <sub>0.00</sub>	65.14 <sub>0.25</sub>	62.24 <sub>0.22</sub>	67.85 <sub>0.00</sub>	47.60 <sub>1.09</sub>	47.69 <sub>0.48</sub>	52.43 <sub>0.00</sub>
IFPL	<b>68.69</b> <sub>0.45</sub>	<b>66.12</b> <sub>0.46</sub>	76.91 <sub>0.00</sub>	67.11 <sub>1.19</sub>	62.93 <sub>1.23</sub>	73.53 <sub>0.00</sub>	<b>51.65</b> <sub>0.70</sub>	<b>50.34</b> <sub>0.65</sub>	57.86 <sub>0.01</sub>
GRIP	<b>67.95</b> <sub>1.2</sub>	63.09 <sub>0.56</sub>	<b>77.18</b> <sub>0.00</sub>	<b>71.22</b> <sub>0.77</sub>	<b>68.43</b> <sub>0.61</sub>	<b>82.19</b> <sub>0.00</sub>	<b>54.57</b> <sub>4.86</sub>	<b>50.51</b> <sub>0.99</sub>	<b>63.42</b> <sub>0.00</sub>
Δ IFPL	↑ 1.66	↓ 0.38	↑ 4.97	↑ 1.97	↑ 0.69	↑ 5.68	↑ 4.05	↑ 2.65	↑ 5.43
Δ GRIP	↑ 0.92	↓ 3.41	↑ 5.24	↑ 6.08	↑ 6.19	↑ 14.34	↑ 6.97	↑ 2.82	↑ 10.35
Multimodal prompts									
FPL	72.54 <sub>0.36</sub>	<b>65.26</b> <sub>0.38</sub>	77.47 <sub>0.00</sub>	62.84 <sub>1.05</sub>	62.32 <sub>0.65</sub>	71.43 <sub>0.00</sub>	43.71 <sub>2.19</sub>	44.85 <sub>0.31</sub>	54.86 <sub>0.00</sub>
IFPL	<b>73.14</b> <sub>0.87</sub>	<b>65.39</b> <sub>1.33</sub>	81.47 <sub>0.00</sub>	70.60 <sub>1.04</sub>	63.69 <sub>0.53</sub>	46.04 <sub>0.36</sub>	<b>53.21</b> <sub>1.24</sub>	<b>47.59</b> <sub>1.04</sub>	43.17 <sub>0.25</sub>
GRIP	<b>74.56</b> <sub>2.02</sub>	<b>64.82</b> <sub>1.63</sub>	<b>82.01</b> <sub>0.00</sub>	<b>73.78</b> <sub>0.91</sub>	<b>69.37</b> <sub>0.61</sub>	<b>82.17</b> <sub>0.00</sub>	<b>54.07</b> <sub>2.25</sub>	<b>47.37</b> <sub>0.70</sub>	<b>63.42</b> <sub>0.00</sub>
Δ IFPL	↑ 1.91	↑ 0.13	↑ 4.00	↑ 7.76	↑ 1.37	↓ 25.39	↑ 9.5	↑ 2.74	↓ 11.69
Δ GRIP	↑ 2.02	↓ 0.44	↑ 4.54	↑ 10.84	↑ 7.05	↑ 10.74	↑ 10.36	↑ 2.52	↑ 8.56

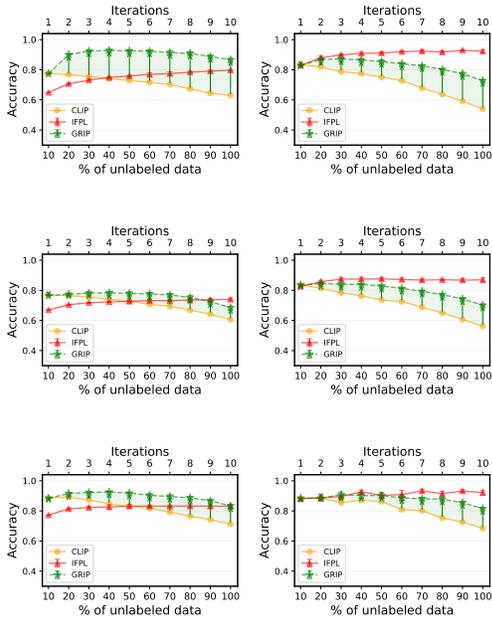
Table 2: For each learning paradigm, we compare FPL, IFPL, and GRIP on Flowers102, RESICS45, and DTD, for all the learning settings SSL, UL, TRZSL. We average across 5 runs and report the standard deviation. Δ METHOD is the difference between the accuracy of FPL and METHOD.

300 data in each iteration. Initially, GRIP maintains accuracy, but as it nears completion the qual-  
301 ity tends to decrease, while a larger dataset with good-quality pseudolabels becomes available.  
302 Comparing GRIP and CLIP, GRIP’s expanded  
303 pseudolabels exhibit superior quality and performs  
304 better (Table 1). Even though IFPL’s pseudolabel  
305 accuracy surpasses GRIP in the final iteration,  
306 GRIP’s overall performance remains better due  
307 to training on a larger number of pseudolabels  
308 (Table 2). This suggests that numerous, slightly  
309 noisier pseudolabels can yield better results, high-  
310 lighting a trade-off and offering insights for future  
311 approaches.

312 **GRIP benefits adaptation even for larger image**  
313 **encoders** We measure how much the effect of  
314 the iterative strategies changes if we consider a  
315 larger pre-trained image encoder. In Table 3, we  
316 report the average improvements of GRIP on CLIP  
317 for Flowers102, RESICS45, and DTD. The mag-  
318 nitude of improvements slightly decreases when  
319 using a larger image encoder. However, we still  
320 see significant benefits for both modalities.

## 321 4.2 The Robin Hood effect

322 Although training models with pseudolabels can  
323 lead to good performance, it can also result in bi-  
324 ased predictions and generate disparate impacts on  
325 sub-populations, i.e., the “Matthew effect” [49, 8].  
326 Particularly, the use of pseudolabels can lead to  
327 improved performance in *well-behaved* (high accuracy) classes but can cause stagnation or decreased  
328 performance in *poorly behaved* (low accuracy) classes. As we explore the use of pseudolabels, we  
329 investigate how the accuracy of the analyzed approaches distributes across classes. Figure 4 shows  
330 an opposite scenario from typical SSL. For all learning paradigms, the iterative training strategies  
331 increase the accuracy of classes where CLIP is not proficient, while maintaining or decreasing the accu-



(a) Flowers102 (b) RESICS45

Figure 3: Evolution of pseudolabels accuracy during training. The rows refer to SSL, UL, and TRZSL, in order. IFPL refers to the top x-axis, while CLIP and GRIP to the bottom.

332 racy of initially well-behaved classes. This effect, we call the “Robin Hood effect,” is very interesting  
 333 because it shows how CLIP can mitigate its own bias toward certain classes by learning from itself.

Textual prompts			
	SSL	UL	TRZSL
Avg. $\Delta$ CLIP (ViT-B/32)	17.46 <sub>12.83</sub>	8.36 <sub>2.85</sub>	23.77 <sub>2.51</sub>
Avg. $\Delta$ CLIP (ViT-L/14)	15.85 <sub>6.44</sub>	8.16 <sub>6.12</sub>	19.96 <sub>6.19</sub>
Visual prompts			
	SSL	UL	TRZSL
Avg. $\Delta$ CLIP (ViT-B/32)	10.78 <sub>4.28</sub>	6.88 <sub>-0.58</sub>	20.28 <sub>13.78</sub>
Avg. $\Delta$ CLIP (ViT-L/14)	7.61 <sub>3.27</sub>	4.89 <sub>-0.48</sub>	16.14 <sub>11.13</sub>

341 Table 3: Average improvement of GRIP with different backbones on Flowers102, RESICS45, and  
 342 DTD.  $\Delta$  CLIP is the difference between the accuracy of GRIP and CLIP. Alongside the average, we  
 343 provide the minimum improvement across tasks.  
 344

345 We find that both approaches yield similar overall accuracies. However, we observe the  
 346 "Matthew effect" when using the first approach. In contrast, when using CLIP-based pseudolabels,  
 347 the class disparity of the regressor trained solely on seen classes is reduced. Particularly, we see a  
 348 significant improvement on initially poor classes, together with a significant diminish of the accuracy  
 349 of well-behaved classes. We observe a clear manifestation of the “Robin Hood effect.” We present  
 350 plots illustrating this effect in Appendix A.4.  
 351

352 **Prompt tuning retains the accuracy of already rich classes better than linear probing** To evaluate the role of prompt  
 353 tuning in the “Robin Hood effect,” we train a linear classifier and textual prompts in the UL setting using GRIP’s training  
 354 strategy. Comparing the per-class accuracies of the two approaches, GRIP on prompts shows an average improvement  
 355 of 22.85 points for the poor classes across tasks, along with a slight average decrease of 0.3 points for the rich classes. On the  
 356 other hand, linear probing determines a 14.42-point improvement for the poor classes, but it results in an average decrease  
 357 of 9.39 points in accuracy for the rich classes (Appendix A.4).  
 358

## 359 5 Conclusions

360 We show that prompt tuning using pseudolabels generated by CLIP itself is a successful approach to enhance CLIP across  
 361 various learning settings. Training strategies that iteratively refine pseudolabels turn out to be effective ways of leveraging  
 362 pseudolabeled data. These approaches not only enhance CLIP’s accuracy but also mitigate model biases toward certain  
 363 classes. We hope this work lays a solid groundwork for reducing reliance on labeled data when adapting pre-trained  
 364 vision-language models like CLIP to new tasks.

365 **Limitations** The effectiveness of the training strategies examined in this paper depends on both the strategies themselves  
 366 and the quality of pseudolabels. The latter is particularly crucial. If CLIP performs poorly on a task, we may struggle to obtain a reliable set of pseudolabels to begin with, potentially diminishing CLIP’s  
 367 performance. Despite this potential risk, we have not observed any relevant failure of GRIP, even in tasks where CLIP’s initial accuracy is extremely low (such as FGVCAircraft). The pseudolabeling  
 368 strategy we adopt involves selecting  $K$  pseudolabels per class, which can create a strong assumption about the distribution of the training data if we attempt to cover all unlabeled data. In fact, during the  
 369 final iteration, it is as if we assume a uniform data distribution.  
 370

To understand the roots of the Robin Hood effect, we examine two factors: (1) the role of pseudolabels generated by CLIP, and (2) the role of prompt tuning. To disentangle these factors, we explore the variation in per-class accuracy of a basic linear classifier trained on CLIP’s ViT-B/32 image representation.

### “Second generation” pseudolabels are a good treatment for class disparity

We train the linear classifier in the SSL setting on 2 labeled examples per class and pseudolabels. The pseudolabels are obtained through conventional methods, where a threshold of .95 is applied, or by using CLIP to generate 16 pseudolabels per

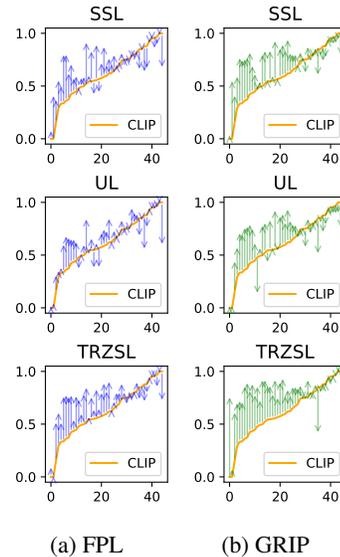


Figure 4: Improvements of FPL and GRIP on CLIP’s per-class accuracies (RESICS45). X-axis is the ranked class index, while y-axis is the accuracy.

384 **References**

- 385 [1] Stephen Bach, Victor Sanh, Zheng Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak,  
386 Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, Zaid Alyafeai, Manan Dey,  
387 Andrea Santilli, Zhiqing Sun, Srulik Ben-david, Canwen Xu, Gunjan Chhablani, Han Wang,  
388 Jason Fries, Maged Al-shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru  
389 Tang, Dragomir Radev, Mike Tian-jian Jiang, and Alexander Rush. PromptSource: An integrated  
390 development environment and repository for natural language prompts. In *Proceedings of the*  
391 *60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.  
392 Association for Computational Linguistics, 2022.
- 393 [2] Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Exploring visual  
394 prompts for adapting large-scale models. *arXiv preprint arXiv:2203.17274*, 2022.
- 395 [3] David Berthelot, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang,  
396 and Colin Raffel. Remixmatch: Semi-supervised learning with distribution matching and  
397 augmentation anchoring. In *International Conference on Learning Representations*, 2020.
- 398 [4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A  
399 Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural*  
400 *Information Processing Systems*, 2019.
- 401 [5] Liu Bo, Qiulei Dong, and Zhanyi Hu. Hardness sampling for self-training based transductive  
402 zero-shot learning. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition*  
403 *(CVPR)*, pages 16494–16503, 2021.
- 404 [6] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx,  
405 Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson,  
406 S. Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen A.  
407 Creel, Jared Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano  
408 Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren E.  
409 Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto,  
410 Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas F.  
411 Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling,  
412 Fereshte Khani, O. Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kuditipudi,  
413 Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa  
414 Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric  
415 Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Benjamin  
416 Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, J. F. Nyarko, Giray Ogut,  
417 Laurel J. Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher  
418 Potts, Aditi Raghunathan, Robert Reich, Hongyu Ren, Frieda Rong, Yusuf H. Roohani, Camilo  
419 Ruiz, Jack Ryan, Christopher R’e, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy  
420 Shih, Krishna Parasuram Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian  
421 Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie,  
422 Michihiro Yasunaga, Jiaxuan You, Matei A. Zaharia, Michael Zhang, Tianyi Zhang, Xikun  
423 Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and  
424 risks of foundation models. *ArXiv*, abs/2108.07258, 2021.
- 425 [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
426 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel  
427 Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler,  
428 Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott  
429 Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya  
430 Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural*  
431 *Information Processing Systems*, 2020.
- 432 [8] Baixu Chen, Junguang Jiang, Ximei Wang, Pengfei Wan, Jianmin Wang, and Mingsheng  
433 Long. Debaised self-training for semi-supervised learning. In *Advances in Neural Information*  
434 *Processing Systems*, 2022.
- 435 [9] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification:  
436 Benchmark and state of the art. *Proceedings of the IEEE*, 2017.

- 437 [10] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi.  
438 Describing textures in the wild. *2014 IEEE Conference on Computer Vision and Pattern  
439 Recognition*, 2013.
- 440 [11] Li Deng. The mnist database of handwritten digit images for machine learning research [best of  
441 the web]. *IEEE Signal Processing Magazine*, 2012.
- 442 [12] Yunhe Gao, Xingjian Shi, Yi Zhu, Hongya Wang, Zhiqiang Tang, Xiong Zhou, Mu Li, and Dim-  
443 itris N. Metaxas. Visual prompt tuning for test-time domain adaptation. *ArXiv*, abs/2210.04831,  
444 2022.
- 445 [13] Chunjiang Ge, Rui Huang, Mixue Xie, Zihang Lai, Shiji Song, Shuang Li, and Gao Huang.  
446 Domain adaptation via prompt learning. *ArXiv*, abs/2202.06687, 2022.
- 447 [14] Patrick Helber, Benjamin Bischke, Andreas R. Dengel, and Damian Borth. Eurosat: A novel  
448 dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal  
449 of Selected Topics in Applied Earth Observations and Remote Sensing*, 2017.
- 450 [15] Hao Huang, Jack Chu, and Fangyun Wei. Unsupervised prompt learning for vision-language  
451 models. *ArXiv*, abs/2204.03649, 2022.
- 452 [16] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Label propagation for  
453 deep semi-supervised learning. *2019 IEEE/CVF Conference on Computer Vision and Pattern  
454 Recognition (CVPR)*, pages 5065–5074, 2019.
- 455 [17] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-  
456 Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation  
457 learning with noisy text supervision. In *International Conference on Machine Learning*, 2021.
- 458 [18] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan,  
459 and Ser-Nam Lim. Visual prompt tuning. In *ECCV 2022: 17th European Conference on  
460 Computer Vision*, 2022.
- 461 [19] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fa-  
462 had Shahbaz Khan. Maple: Multi-modal prompt learning. *ArXiv*, abs/2210.03117, 2022.
- 463 [20] Dong-Hyun Lee. Pseudo-label : The simple and efficient semi-supervised learning method  
464 for deep neural networks. *ICML 2013 Workshop : Challenges in Representation Learning  
465 (WREPL)*, 2013.
- 466 [21] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient  
467 prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural  
468 Language Processing*, 2021.
- 469 [22] Will LeVine, Benjamin Pikus, Pranav Vishnu Raja, and Fernando Amat. Enabling calibration  
470 in the zero-shot inference of large vision-language models. In *ICLR 2023 Workshop on Pitfalls  
471 of limited data and computation for Trustworthy ML*, 2023.
- 472 [23] Chunyuan Li, Haotian Liu, Liunian Li, Pengchuan Zhang, Jyoti Aneja, Jianwei Yang, Ping  
473 Jin, Houdong Hu, Zicheng Liu, Yong Jae Lee, and Jianfeng Gao. Elevator: A benchmark and  
474 toolkit for evaluating language-augmented visual models. In *Advances in Neural Information  
475 Processing Systems*, 2022.
- 476 [24] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation.  
477 In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics  
478 and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long  
479 Papers)*, 2021.
- 480 [25] Bo Liu, Lihua Hu, Qiulei Dong, and Zhanyi Hu. An iterative co-training transductive framework  
481 for zero shot learning. *IEEE Transactions on Image Processing*, 30:6943–6956, 2021.
- 482 [26] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew B. Blaschko, and Andrea Vedaldi. Fine-  
483 grained visual classification of aircraft. *ArXiv*, abs/1306.5151, 2013.

- 484 [27] Shu Manli, Nie Weili, Huang De-An, Yu Zhiding, Goldstein Tom, Anandkumar Anima, and  
485 Xiao Chaowei. Test-time prompt tuning for zero-shot generalization in vision-language models.  
486 In *NeurIPS*, 2022.
- 487 [28] Nihal V. Nayak, Peilin Yu, and Stephen Bach. Learning to compose soft prompts for composi-  
488 tional zero-shot learning. In *The Eleventh International Conference on Learning Representa-*  
489 *tions*, 2023.
- 490 [29] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large  
491 number of classes. *2008 Sixth Indian Conference on Computer Vision, Graphics & Image*  
492 *Processing*, 2008.
- 493 [30] Wasu Piriyaakulkij, Cristina Menghini, Ross Briden, Nihal V. Nayak, Jeffrey Zhu, Elaheh Raisi,  
494 and Stephen H. Bach. TAGLETS: A system for automatic semi-supervised learning with  
495 auxiliary data. In *Conference on Machine Learning and Systems (MLSys)*, 2022.
- 496 [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agar-  
497 wal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya  
498 Sutskever. Learning transferable visual models from natural language supervision. In *Interna-*  
499 *tional Conference on Machine Learning*, 2021.
- 500 [32] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language  
501 models are unsupervised multitask learners. 2019.
- 502 [33] Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai,  
503 Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish  
504 Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal  
505 Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica,  
506 Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala  
507 Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan  
508 Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M Rush. Multitask prompted training enables zero-shot task generalization. In *International Conference on Learning Representations*, 2022.
- 511 [34] Sheng Shen, Shijia Yang, Tianjun Zhang, Bohan Zhai, Joseph E. Gonzalez, Kurt Keutzer, and  
512 Trevor Darrell. Multitask vision-language prompt tuning. *arXiv preprint arXiv:2211.11720*,  
513 2022.
- 514 [35] Weiwei Shi, Yihong Gong, C. Ding, Zhiheng Ma, Xiaoyu Tao, and Nanning Zheng. Transductive  
515 semi-supervised deep learning using min-max features. In *European Conference on Computer*  
516 *Vision*, 2018.
- 517 [36] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raf-  
518 fel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-  
519 supervised learning with consistency and confidence. In *Advances in Neural Information*  
520 *Processing Systems*, 2020.
- 521 [37] Ximeng Sun, Ping Hu, and Kate Saenko. Dualcoop: Fast adaptation to multi-label recognition  
522 with limited annotations. *ArXiv*, 2022.
- 523 [38] Xudong Wang, Zhi-Li Wu, Long Lian, and Stella X. Yu. Debaised learning from naturally  
524 imbalanced pseudo-labels. *2022 IEEE/CVF Conference on Computer Vision and Pattern*  
525 *Recognition (CVPR)*, 2022.
- 526 [39] Yongqin Xian, Bernt Schiele, and Zeynep Akata. Zero-shot learning — the good, the bad and  
527 the ugly. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- 528 [40] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data aug-  
529 mentation for consistency training. In *Advances in Neural Information Processing Systems*,  
530 2020.
- 531 [41] Yi Xu, Lei Shang, Jinxing Ye, Qi Qian, Yu-Feng Li, Baigui Sun, Hao Li, and Rong Jin. Dash:  
532 Semi-supervised learning with dynamic thresholding. In *International Conference on Machine*  
533 *Learning*, 2021.

- 534 [42] Yunlong Yu, Zhong Ji, Xi Li, Jichang Guo, Zhongfei Zhang, Haibin Ling, and Fei Wu. Trans-  
535 ductive zero-shot learning with a self-training dictionary approach. *IEEE Transactions on*  
536 *Cybernetics*, 48(10):2908–2919, 2018.
- 537 [43] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel C. F. Codella, Xiyang Dai, Jianfeng Gao,  
538 Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, Ce Liu, Mengchen Liu, Zicheng Liu,  
539 Yumao Lu, Yu Shi, Lijuan Wang, Jianfeng Wang, Bin Xiao, Zhen Xiao, Jianwei Yang, Michael  
540 Zeng, Luowei Zhou, and Pengchuan Zhang. Florence: A new foundation model for computer  
541 vision. *ArXiv*, abs/2111.11432, 2021.
- 542 [44] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Unified vision and  
543 language prompt learning. *ArXiv*, abs/2210.07225, 2022.
- 544 [45] Bowen Zhang, Yidong Wang, Wenxin Hou, HAO WU, Jindong Wang, Manabu Okumura, and  
545 Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo  
546 labeling. In *Advances in Neural Information Processing Systems*, 2021.
- 547 [46] X. Zhang, Yusuke Iwasawa, Yutaka Matsuo, and Shixiang Shane Gu. Domain prompt learning  
548 for efficiently adapting clip to unseen domains. 2021.
- 549 [47] Tony Z. Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. Calibrate before use:  
550 Improving few-shot performance of language models. 2021.
- 551 [48] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for  
552 vision-language models. *International Journal of Computer Vision*, 2021.
- 553 [49] Zhaowei Zhu, Tianyi Luo, and Yang Liu. The rich get richer: Disparate impact of semi-  
554 supervised learning. In *International Conference on Learning Representations*, 2022.

## 555 A Appendix

556 We include here extra information that supports the results presented in the main body of the paper.

557 **Reproducibility** We have provided the code to run the experiments as supplementary material for  
558 the submission. However, we plan to release it as an open repository upon acceptance.

### 559 A.1 Trainable Prompts

560 **Text Prompt Tuning** The primary objective of text prompt tuning is to improve the alignment  
561 between the class token and the image features extracted by the image encoder. This is achieved by  
562 adding learnable vectors, i.e., *prefix*, before the CLASS token to create a contextualized representation.  
563 Specifically, the sequence

$$\mathbf{t} = [\mathbf{V}]_1[\mathbf{V}]_2 \dots [\mathbf{V}]_M[\text{CLASS}]$$

564 is fed into the textual encoder, where each vector  $[\mathbf{V}]_m$  ( $m \in 1, \dots, M$ ) has the same dimension as  
565 word embeddings, and  $M$  is a hyperparameter that determines the length of the prefix.

566 Context Optimization (CoOp) [48] was the first work to explore continuous prompts for VLMs.  
567 Follow-up works have experimented with different training strategies to enhance the generalizability  
568 of the learned prompts while preserving the core concept of continuous vector tuning [34, 12, 27, 46,  
569 13, 37].

570 Tuning the text prefix vector changes the resulting  $n$  linear weight vectors  $w_i = \psi(p_i)$ , while leaving  
571 the image features unchanged. Therefore, text prompt tuning may be most beneficial when image  
572 features are well-separated by class but may not be aligned with the corresponding textual prompt.  
573 Conversely, text prompt tuning may not be as effective when the image features are poorly separated,  
574 as in specialized or novel domains where CLIP may lack sufficient training data.

	Num. classes ( $ \mathcal{Y} $ )	Num. seen classes ( $ \mathcal{S} $ )	Num. unseen classes ( $ \mathcal{U} $ )	Size training data	Avg. labeled data per class	Size test
Flowers102	102	63	39	2040	16	6149
RESICS45	45	27	18	6300	110	25200
FGVC-Aircraft	100	62	38	6667	53	3333
MNIST	10	6	4	60000	4696	10000
EuroSAT	10	6	4	27000	2200	5000
DTD	47	29	18	3760	64	1880

Table 4: For each dataset we report the number of classes, the number of seen and unseen classes in the TRZSL setting, the size of training data (including both labeled and unlabeled data), the average number of labeled examples per class, and the size of the test set which is the same across learning paradigms. We recall that we use the datasets gathered by the recent ELEVATER [23] benchmark for vision-language models.

575 **Visual Prompt Tuning** Instead of tuning the text prompts, one can also tune the inputs of the vision  
576 encoder. In this case, a learnable visual prefix is prepended to the image tokens as input to the image  
577 transformer as follows:

$$\hat{\mathbf{I}} = [p]_1 \dots [p]_K [I]_1 \dots [I]_P$$

578 where  $p$  represents a sequence of  $K$  learnable prefix vectors, and  $[I]_1 \dots [I]_P$  are the image tokens  
579 from the corresponding  $P$  patches of the input images. The new sequence  $\hat{\mathbf{I}}$  is the input to the image  
580 encoder  $\phi$ .

581 Visual Prompt Tuning (VPT) was introduced in the context of efficiently adapting pre-trained vision  
582 transformers to downstream tasks [18]. However, the approach has since been applied in the context  
583 of VLM [34].

584 Whereas text prompt tuning does not alter the image features, visual prompt tuning does. By  
585 rearranging the image features within the projection space, VPT has the potential to improve CLIP  
586 when the image features are not well separated by class, such as in specialized domains.

587 **Multimodal Prompt Tuning** The previous approaches are unimodal, as they either involve modi-  
588 fying the text or visual input, but never both. This choice may be suboptimal as it does not allow the  
589 flexibility to dynamically adjust both representations on a downstream task. Recently, multimodal  
590 prompt tuning has been introduced [44, 19]. We focus on Unified Prompt Tuning (UPT) [44] which  
591 essentially learns a tiny neural network to jointly optimize prompts across different modalities. UPT  
592 learns a set of prompts  $\mathbf{U} = [\mathbf{U}_T, \mathbf{U}_V] \in \mathbb{R}^{d \times n}$  with length  $n$ , where  $\mathbf{U}_T \in \mathbb{R}^{d \times n_T}$ ,  $\mathbf{U}_V \in \mathbb{R}^{d \times n_V}$ .  
593  $\mathbf{U}$  is transformed as follows:

$$\begin{aligned} \mathbf{U}' &= \text{SA}(\mathbf{U}) + \text{LN}(\mathbf{U}) \\ \hat{\mathbf{U}} &= \text{FFN}(\text{LN}(\mathbf{U}')) + \text{LN}(\mathbf{U}') \end{aligned}$$

594 where SA is the self-attention operator, LN is the layer normalization operator, and FFN is a feed  
595 forward network. After transformation, we obtain  $\hat{\mathbf{U}} = [\hat{\mathbf{U}}_T, \hat{\mathbf{U}}_V] \in \mathbb{R}^{d \times n}$ , such that  $\hat{\mathbf{U}}_T$  is to be  
596 used as a text prompt, and  $\hat{\mathbf{U}}_V$  is to be used as a visual prompt.

597 The author of UPL argue that self-attention allows for beneficial interaction between the two separate  
598 modalities, which leads to both separable visual features, and text classifiers that are well-aligned  
599 with the corresponding visual features [44].

600 **Prompts initialization** We initialize textual and visual prompts from a normal distribution of mean  
601 0 and variance 0.02. We note that we learn shallow visual prompts by modifying only the input to the  
602 image encoder. Multimodal prompts are initialized from a uniform distribution. We found that the  
603 latter was not working properly for textual and visual prompts.

604 **Additional training settings** For training, the batch size is 64.

## 605 A.2 Datasets details

606 We use six datasets from specialized or fine-grained domains. Here we provide a description of each  
607 of them. In Table 4, we report the details about the number of classes and data available for each  
608 dataset. For each dataset, we also show CLIP’s prediction distribution over classes Figure 5.

609 **Flowers102 [29]** It is a dataset collecting images for 102  
 610 flower categories commonly occurring in the United King-  
 611 dom. For each class we have between 40 and 258 images.  
 612 Figure 5a shows that CLIP’s predictions are skewed toward  
 613 certain classes, which are predicted more often than what we  
 614 would expect according to the real class distribution on the test  
 615 set.

616 **RESICS45 [9]** This is a publicly available benchmark for  
 617 Remote Sensing Image Scene Classification. It collects 45 kind  
 618 of scenes. Figure 5b shows that CLIP predicts more often a  
 619 subset of classes.

620 **FGVC-Aircraft [26]** It describes the fine-grained task of cat-  
 621 egorizing aircraft. We consider the task of classifying aircrafts  
 622 into 100 variants. Also for this task, CLIP assigns images to a  
 623 reduced set of classes (Figure 5c).

624 **MNIST [11]** MNIST is a database of handwritten digits. The  
 625 digits are size-normalized and centered in a fixed-size image.  
 626 We observe that CLIP never predicts 6 out of 10 classes (Fig-  
 627 ure 5c).

628 **EuroSAT [14]** EuroSAT represents the task of categorizing  
 629 satellite images of scenes. It consists of 10 classes. In Figure 5e,  
 630 we show CLIP’s predictions distribution over the classes.

631 **DTD [10]** DTD stands for Describable Textures Dataset. It is  
 632 an evolving collection of textural images in the wild, and it is  
 633 annotated relying on human-centric attributes, inspired by the  
 634 perceptual properties of textures. The zero-shot CLIP predic-  
 635 tions show the model’s bias toward certain classes (Figure 5f).

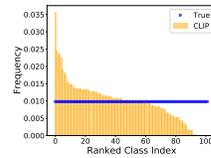
### 636 A.3 Experiments

637 In this section, we report tables and plots that complement the  
 638 results presented in Section 4.

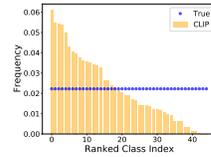
639 **The effect of GRIP on multimodal prompts** Table 5 shows  
 640 the improvements of GRIP on CLIP and Unified Prompt Tuning  
 641 (UPL) [44]. Similar to the results in Table 1, GRIP consistently  
 642 improves CLIP with respect to the baselines. The improve-  
 643 ments on CLIP are by 18.2 in semi-supervised learning, 14.8 in  
 644 unsupervised learning, and 30.7 in transductive zero-shot learn-  
 645 ing. While GRIP outperforms UPL by 4.7 in semi-supervised  
 646 learning, and 19.5 in transductive zero-shot learning.

647 **Comparison across iterative strategies** In Table 6, we report  
 648 a comparison between FPL and the iterative strategies (IFPL  
 649 and GRIP) on MNIST, EuroSAT, and FGVC-Aircraft. Results on the other tasks can be found in the  
 650 main body of the paper Section 4.1. While GRIP largely and consistently outperforms FPL by on  
 651 average 16.7 points in accuracy, IFPL is not robust and it leads to performances that are inferior to  
 652 FPL by on average 4.4 points in accuracy.

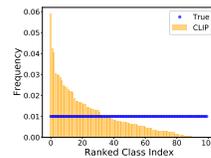
653 **The evolving accuracy of dynamic pseudolabels** Figure 6 represents the evolution of pseudolabels  
 654 accuracy during training for all datasets, but Flowers102 and RESICS45 presented in Figure 3. We  
 655 observe that the accuracy of the pseudolabels characterizes the overall performance of the models  
 656 reported in Table 6. For instance, IFPL for EuroSAT in the TRZSL setting is highly variable,



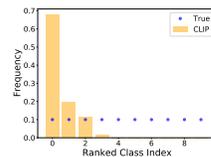
(a) Flowers102



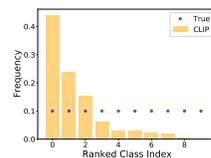
(b) RESICS45



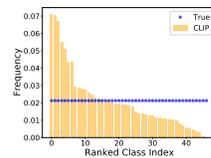
(c) FGVC-Aircraft



(d) MNIST



(e) EuroSAT



(f) DTD

Figure 5: For each dataset we show the distribution of CLIP’s predictions over classes on the test set. The blue dots represent the true class distribution.

Multimodal prompts									
Method	Flowers102			RESICS45			FGVCAircraft		
	SSL	UL	TRZSL	SSL	UL	TRZSL	SSL	UL	TRZSL
CLIP	63.67 <sub>0.00</sub>	-	63.40 <sub>0.00</sub>	54.48 <sub>0.00</sub>	-	54.46 <sub>0.00</sub>	<b>17.58</b> <sub>0.00</sub>	-	<b>17.86</b> <sub>0.00</sub>
UPT	68.03 <sub>1.29</sub>	-	61.05 <sub>0.04</sub>	62.84 <sub>1.05</sub>	-	58.79 <sub>0.04</sub>	11.13 <sub>4.98</sub>	-	15.89 <sub>0.07</sub>
GRIP	<b>74.56</b> <sub>2.02</sub>	64.82 <sub>1.63</sub>	<b>82.01</b> <sub>0.01</sub>	<b>73.68</b> <sub>0.91</sub>	69.37 <sub>0.61</sub>	<b>82.17</b> <sub>0.00</sub>	17.36 <sub>0.43</sub>	14.73 <sub>0.08</sub>	<b>17.85</b> <sub>10.30</sub>
$\Delta$ CLIP	$\uparrow$ 10.89	$\uparrow$ 1.15	$\uparrow$ 18.61	$\uparrow$ 19.2	$\uparrow$ 14.89	$\uparrow$ 27.71	$\downarrow$ 0.22	$\downarrow$ 2.85	$\downarrow$ 0.01
$\Delta$ UPT	$\uparrow$ 6.53	-	$\uparrow$ 20.96	$\uparrow$ 10.84	-	$\uparrow$ 22.38	$\uparrow$ 6.23	-	$\uparrow$ 1.96
MNIST									
EuroSAT									
DTD									
CLIP	25.10 <sub>0.00</sub>	-	20.77 <sub>0.00</sub>	32.88 <sub>0.00</sub>	-	30.54 <sub>0.00</sub>	43.24 <sub>0.00</sub>	-	43.45 <sub>0.00</sub>
UPT	<b>64.44</b> <sub>3.66</sub>	-	63.59 <sub>0.11</sub>	<b>68.85</b> <sub>9.92</sub>	-	60.43 <sub>0.04</sub>	43.71 <sub>2.18</sub>	-	36.91 <sub>0.04</sub>
GRIP	<b>65.94</b> <sub>2.23</sub>	<b>68.18</b> <sub>run</sub>	<b>73.75</b> <sub>2.93</sub>	<b>60.38</b> <sub>1.77</sub>	<b>61.52</b> <sub>3.04</sub>	<b>95.52</b> <sub>0.40</sub>	<b>54.07</b> <sub>2.25</sub>	<b>47.37</b> <sub>0.7</sub>	<b>63.42</b> <sub>0.00</sub>
$\Delta$ CLIP	$\uparrow$ 40.84	$\uparrow$ 43.08	$\uparrow$ 52.98	$\uparrow$ 27.5	$\uparrow$ 28.64	$\uparrow$ 64.98	$\uparrow$ 10.83	$\uparrow$ 4.13	$\uparrow$ 19.97
$\Delta$ UPT	$\uparrow$ 2.35	-	$\uparrow$ 10.16	$\downarrow$ 8.47	-	$\uparrow$ 35.09	$\uparrow$ 10.36	-	$\uparrow$ 26.51

Table 5: For each learning paradigm, we compare the accuracy of GRIP with CLIP zero-shot (ViT-B/32), and UPL. Results are for SSL, UL, and TRZSL on FRAMED. We average the accuracy on 5 seeds and report the standard deviation.  $\Delta$  METHOD is the difference between the accuracy of GRIP and METHOD. We note that for UL we can not apply UPL since no labeled data is available.

Textual prompts									
Method	MNIST			EuroSAT			FGVCAircraft		
	SSL	UL	TRZSL	SSL	UL	TRZSL	SSL	UL	TRZSL
FPL	66.06 <sub>1.10</sub>	40.03 <sub>2.63</sub>	9.73 <sub>19.45</sub>	<b>62.05</b> <sub>1.64</sub>	48.96 <sub>1.49</sub>	53.70 <sub>26.87</sub>	<b>20.02</b> <sub>0.77</sub>	<b>16.62</b> <sub>0.67</sub>	17.55 <sub>0.37</sub>
IFPL	59.14 <sub>3.43</sub>	28.94 <sub>2.05</sub>	0.00 <sub>0.00</sub>	<b>61.28</b> <sub>1.59</sub>	<b>56.46</b> <sub>3.26</sub>	14.36 <sub>28.71</sub>	18.00 <sub>0.35</sub>	13.80 <sub>0.67</sub>	21.72 <sub>0.77</sub>
GRIP	<b>71.78</b> <sub>3.59</sub>	<b>67.88</b> <sub>2.76</sub>	<b>74.06</b> <sub>0.29</sub>	58.66 <sub>2.64</sub>	<b>57.21</b> <sub>1.77</sub>	<b>92.33</b> <sub>0.69</sub>	16.98 <sub>0.82</sub>	15.22 <sub>0.71</sub>	<b>26.08</b> <sub>0.25</sub>
$\Delta$ IFPL	$\downarrow$ 6.92	$\downarrow$ 11.09	$\downarrow$ 9.73	$\downarrow$ 0.77	$\uparrow$ 7.50	$\downarrow$ 39.34	$\downarrow$ 2.02	$\downarrow$ 2.82	$\uparrow$ 4.17
$\Delta$ GRIP	$\uparrow$ 5.72	$\uparrow$ 27.85	$\uparrow$ 64.33	$\downarrow$ 3.39	$\uparrow$ 8.25	$\uparrow$ 38.63	$\downarrow$ 3.04	$\downarrow$ 1.40	$\uparrow$ 8.53
Visual prompts									
FPL	42.84 <sub>16.80</sub>	39.62 <sub>6.53</sub>	31.82 <sub>17.53</sub>	52.47 <sub>2.53</sub>	48.79 <sub>3.69</sub>	68.68 <sub>14.74</sub>	<b>20.14</b> <sub>0.26</sub>	<b>18.28</b> <sub>0.33</sub>	16.28 <sub>0.45</sub>
IFPL	52.91 <sub>8.99</sub>	37.17 <sub>6.27</sub>	38.38 <sub>4.21</sub>	<b>57.85</b> <sub>6.52</sub>	32.52 <sub>10.00</sub>	48.13 <sub>11.13</sub>	18.77 <sub>0.48</sub>	16.36 <sub>0.37</sub>	19.29 <sub>0.36</sub>
GRIP	<b>69.66</b> <sub>5.51</sub>	<b>68.04</b> <sub>1.11</sub>	<b>69.54</b> <sub>1.31</sub>	<b>63.48</b> <sub>3.09</sub>	<b>63.68</b> <sub>3.42</sub>	<b>96.97</b> <sub>0.77</sub>	<b>19.43</b> <sub>0.50</sub>	<b>17.51</b> <sub>0.61</sub>	<b>26.42</b> <sub>0.30</sub>
$\Delta$ IFPL	$\uparrow$ 10.07	$\downarrow$ 2.45	$\uparrow$ 6.56	$\uparrow$ 5.38	$\downarrow$ 16.27	$\downarrow$ 20.55	$\downarrow$ 1.37	$\downarrow$ 1.92	$\uparrow$ 3.01
$\Delta$ GRIP	$\uparrow$ 26.82	$\uparrow$ 28.42	$\uparrow$ 37.72	$\uparrow$ 11.01	$\uparrow$ 14.89	$\uparrow$ 28.29	$\downarrow$ 0.71	$\downarrow$ 0.77	$\uparrow$ 10.14

Table 6: For each learning paradigm, we compare FPL, IFPL, and GRIP on MNIST, EuroSAT, and FGVCAircraft. We average across 5 runs and report the standard deviation.  $\Delta$  METHOD is the difference between the accuracy of FPL and METHOD.

657 explaining the low average accuracy of the model on the test set (Table 6). Similarly, for MNIST in  
658 the TRZSL we observe that after the first iteration, the pseudolabels get very noisy.

659 **GRIP performance on transductive zero-shot learning** We show how the effectiveness of GRIP  
660 is consistent over the three random splits of seen and unseen classes which we randomly generated.  
661 The splits are reported in Table 9. Table 8 gathers the accuracy of seen and unseen classes, along  
662 with the harmonic mean for all three splits using textual prompts. Beyond the consistent improvement  
663 induced by GRIP training strategy, we observe that the accuracy of GRIP on the seen classes is often  
664 lower than the accuracy of CoOp on the same set of classes. We speculate this can result from two  
665 factors: (1) we learn to distinguish between seen and unseen losing knowledge specialized on the  
666 seen classes, and (2) the parameter  $\lambda$  that upweights the error on the pseudolabeled data is too large  
667 (Section 3.3) and further training might be needed.

#### 668 A.4 The Robin Hood effect

669 **The Robin Hood effect on all tasks** For each dataset, we provide the per-class accuracy distribution  
670 of GRIP compared with CLIP, Figure 8. The Robin Hood effect characterizes all the tasks. We  
671 observe that for GRIP the increase in overall accuracy corresponds to consistent improvements in  
672 the predictions of initially poor classes. By comparing Figure 7 with Figure 8, we see that GRIP  
673 reinforces the Robin Hood effect already visible when using FPL in certain cases.

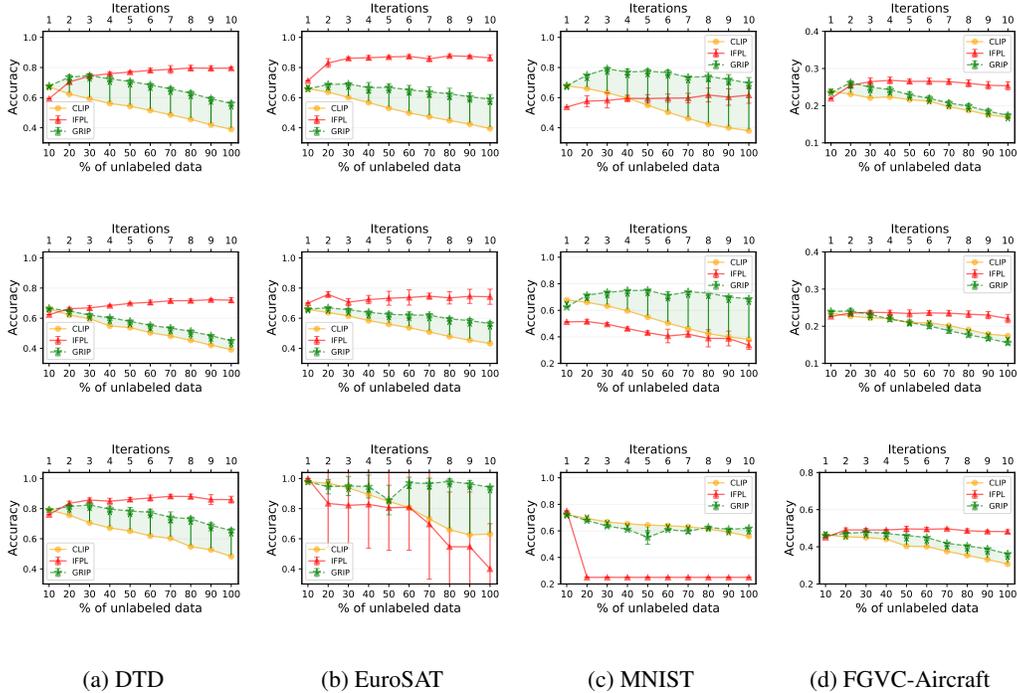


Figure 6: We plot the evolution of dynamic-pseudolabels accuracy during training. The rows refer to SSL, UL, and TRZSL, in order. IFPL refers to the top x-axis, while CLIP and GRIP to the bottom.

674 **The importance of good quality pseudolabels to mitigate the Matthew effect in SSL** In the SSL  
 675 setting, we train a logistic regression on top of the visual feature extracted by CLIP’s image encoder  
 676 (ViT-B/32). In Figure 9, we show the per-class accuracy of the final model trained by combining  
 677 labeled data with either pseudolabels assigned with the conventional scheme (threshold at .95) or 16  
 678 CLIP-generated pseudolabels. We compare the two distribution with the per-class accuracy of the  
 679 model trained solely on the few labeled examples per class (2 instances).

680 **The different impact of prompt tuning and linear probing on the Robin Hood effect** We  
 681 investigate if there is any difference in the Robin Hood effect when adapting CLIP via prompt  
 682 tuning or linear probing. We train both relying on the iterative training strategy that grows the set of  
 683 pseudolabels at each iteration by using the top- $K$  scheme (Section 3). We consider the UL setting.

684 Among the set of target classes, we distinguish between *poor* and *rich* classes. A class is *poor*, if  
 685 CLIP’s accuracy on that class is lower than its overall accuracy on the task. Otherwise, the class is  
 686 considered *rich*. Table 7 reports the accuracy of the two approaches, and the accuracy on the poor and  
 687 rich classes, while highlighting the average effect with respect to CLIP. Training with prompt tuning  
 688 retains more knowledge of the rich classes than linear probing. Prompt tuning reduces the accuracy  
 689 on the rich classes by on average 0.3 points, while linear probing has an average deterioration of 9.4.  
 690 Overall, GRIP works better than linear probing. We note that the lower accuracy of linear probing  
 691 is characterized by a worse ability to correctly predict the rich classes, i.e., “rich get poorer.” This  
 692 is surprising, as we would have expected the errors to concentrate on the poor classes compared to  
 693 CLIP.

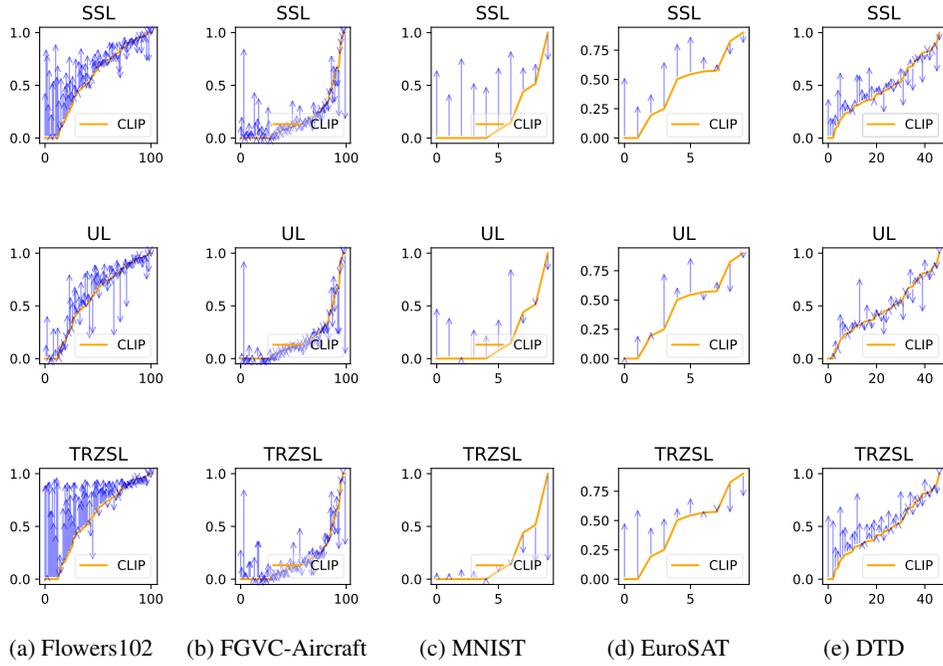


Figure 7: Per-class accuracy of FPL compared to CLIP’s per-class accuracy on Flowers102, FGVC-Aircraft, MNIST, EuroSAT, DTD. **X-axis** is the ranked class index, while the **y-axis** is the accuracy.

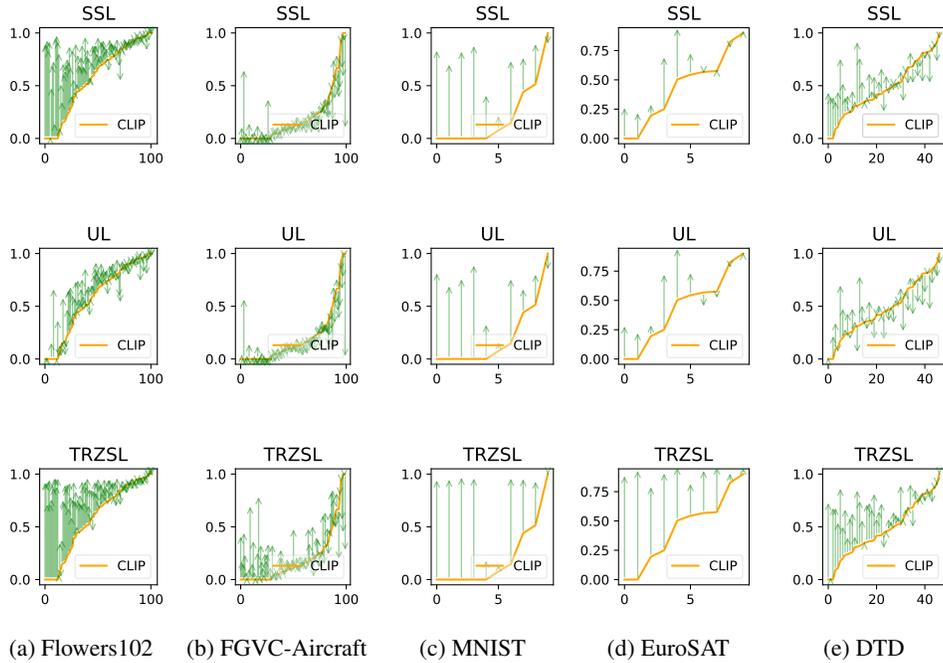


Figure 8: Per-class accuracy of GRIP compared to CLIP’s per-class accuracy on Flowers102, FGVC-Aircraft, MNIST, EuroSAT, and DTD. **X-axis** is the ranked class index, while the **y-axis** is the accuracy.

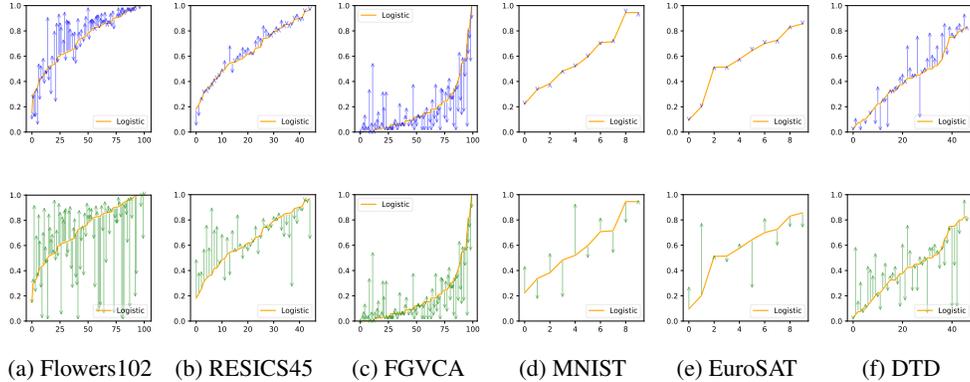


Figure 9: Per-class accuracy of a logistic classifier using conventional pseudolabels (first row) and CLIP-based pseudolabels (second row). The solid orange line represents the per-class accuracy of a logistic regression trained on 2-shots per class. **X-axis** is the ranked class index, while the **y-axis** is the accuracy. We present results for Flowers102, RESICS45, FGVC-Aircraft, MNIST, EuroSAT, and DTD, in order.

	Flowers102	RESICS45	FGVC-Aircraft	MNIST	EuroSAT	DTD	Avg. $\Delta$
Linear probe (LP)	41.01	58.79	61.94	50.52	51.37	10.17	-
GRIP	<b>46.09</b>	<b>70.55</b>	<b>69.84</b>	<b>57.21</b>	<b>67.88</b>	<b>15.22</b>	-
Rich CLIP	<b>67.81</b>	75.47	85.16	65.26	65.14	<b>45.93</b>	-
Rich LP	52.87	69.01	79.55	67.53	50.34	29.12	-
Rich GRIP	56.05	<b>78.81</b>	<b>86.40</b>	<b>71.73</b>	<b>77.84</b>	31.95	-
$\Delta$ LP	$\downarrow 14.92$	$\downarrow 6.47$	$\downarrow 5.61$	$\uparrow 2.26$	$\downarrow 14.79$	$\downarrow 16.81$	$\downarrow 9.39$
$\Delta$ GRIP	$\downarrow 11.76$	$\uparrow 3.33$	$\uparrow 1.24$	$\uparrow 6.46$	$\uparrow 12.70$	$\downarrow 13.98$	$\downarrow 0.33$
Poor CLIP	25.63	35.60	27.98	11.10	3.18	5.35	-
Poor LP	26.50	42.77	36.25	28.34	56.76	4.77	-
Poor GRIP	<b>35.03</b>	<b>56.85</b>	<b>42.82</b>	<b>39.88</b>	<b>65.08</b>	<b>6.31</b>	-
$\Delta$ LP	$\uparrow 0.87$	$\uparrow 7.18$	$\uparrow 8.27$	$\uparrow 17.24$	$\uparrow 53.58$	$\downarrow 0.58$	$\uparrow 14.43$
$\Delta$ GRIP	$\uparrow 9.4$	$\uparrow 21.26$	$\uparrow 14.84$	$\uparrow 28.78$	$\uparrow 61.9$	$\uparrow 0.96$	$\uparrow 22.86$

Table 7: For each task we report the overall accuracy of linear probing (LP) and GRIP textual along with the accuracy on *poor* and *rich* classes.  $\Delta$  METHOD is the difference between the accuracy of CLIP and METHOD. For an overall evaluation of the difference between linear probing and prompt tuning, we report the average difference of LP and GRIP with respect to CLIP on poor and rich classes.

Split 1									
Method	Flowers102			RESICS45			FGVCAircraft		
	S	U	H	S	U	H	S	U	H
CLIP	64.26 <sub>0.00</sub>	62.56 <sub>0.00</sub>	63.40 <sub>0.00</sub>	54.85 <sub>0.00</sub>	54.08 <sub>0.00</sub>	54.46 <sub>0.00</sub>	16.27 <sub>0.00</sub>	19.79 <sub>0.00</sub>	17.86 <sub>0.00</sub>
CoOp	<b>91.52</b> <sub>0.36</sub>	48.35 <sub>2.96</sub>	63.22 <sub>2.60</sub>	<b>84.66</b> <sub>1.01</sub>	50.73 <sub>3.28</sub>	63.37 <sub>2.23</sub>	<b>34.18</b> <sub>1.56</sub>	16.28 <sub>3.69</sub>	21.70 <sub>3.45</sub>
GRIP	90.31 <sub>0.51</sub>	<b>82.57</b> <sub>1.26</sub>	<b>86.26</b> <sub>0.81</sub>	82.68 <sub>0.47</sub>	<b>79.53</b> <sub>0.72</sub>	<b>81.07</b> <sub>0.37</sub>	22.25 <sub>0.07</sub>	<b>31.51</b> <sub>0.59</sub>	<b>26.08</b> <sub>0.25</sub>
$\Delta$ CLIP	$\uparrow$ 26.05	$\uparrow$ 20.01	$\uparrow$ 22.86	$\uparrow$ 27.83	$\uparrow$ 25.45	$\uparrow$ 26.61	$\uparrow$ 5.98	$\uparrow$ 11.72	$\uparrow$ 8.22
$\Delta$ CoOp	$\downarrow$ 1.21	$\uparrow$ 34.22	$\uparrow$ 23.04	$\downarrow$ 1.98	$\uparrow$ 28.8	$\uparrow$ 17.7	$\downarrow$ 11.93	$\uparrow$ 15.23	$\uparrow$ 4.38
Split 2									
Method	Flowers102			RESICS45			FGVCAircraft		
	S	U	H	S	U	H	S	U	H
CLIP	31.74 <sub>0.00</sub>	15.43 <sub>0.00</sub>	20.77 <sub>0.00</sub>	22.33 <sub>0.00</sub>	48.30 <sub>0.00</sub>	30.54 <sub>0.00</sub>	42.50 <sub>0.00</sub>	44.44 <sub>0.00</sub>	43.45 <sub>0.00</sub>
CoOp	<b>94.68</b> <sub>5.64</sub>	15.43 <sub>7.75</sub>	21.15 <sub>12.18</sub>	82.91 <sub>8.81</sub>	46.02 <sub>9.23</sub>	58.64 <sub>5.86</sub>	<b>69.67</b> <sub>1.17</sub>	34.81 <sub>3.44</sub>	46.32 <sub>9.92</sub>
GRIP	<b>95.13</b> <sub>0.11</sub>	<b>60.63</b> <sub>0.44</sub>	<b>74.06</b> <sub>0.29</sub>	<b>91.75</b> <sub>0.53</sub>	<b>92.91</b> <sub>0.91</sub>	<b>92.33</b> <sub>0.70</sub>	<b>68.26</b> <sub>0.69</sub>	<b>62.61</b> <sub>1.87</sub>	<b>65.30</b> <sub>1.03</sub>
$\Delta$ CLIP	$\uparrow$ 63.39	$\uparrow$ 45.2	$\uparrow$ 53.29	$\uparrow$ 69.42	$\uparrow$ 44.61	$\uparrow$ 61.79	$\uparrow$ 25.76	$\uparrow$ 18.17	$\uparrow$ 21.85
$\Delta$ CoOp	$\uparrow$ 0.45	$\uparrow$ 45.2	$\uparrow$ 52.91	$\uparrow$ 8.84	$\uparrow$ 46.89	$\uparrow$ 33.69	$\downarrow$ 1.41	$\uparrow$ 27.8	$\uparrow$ 19.00
Split 2									
Method	Flowers102			RESICS45			FGVCAircraft		
	S	U	H	S	U	H	S	U	H
CLIP	65.38 <sub>0.00</sub>	60.64 <sub>0.00</sub>	62.92 <sub>0.00</sub>	59.50 <sub>0.00</sub>	47.06 <sub>0.00</sub>	52.55 <sub>0.00</sub>	17.30 <sub>0.00</sub>	18.12 <sub>0.00</sub>	17.70 <sub>0.00</sub>
CoOp	<b>91.8</b> <sub>1.32</sub>	47.75 <sub>3.86</sub>	62.77 <sub>3.31</sub>	<b>86.54</b> <sub>1.92</sub>	48.00 <sub>3.01</sub>	61.70 <sub>2.17</sub>	<b>33.59</b> <sub>4.12</sub>	19.57 <sub>1.37</sub>	<b>24.63</b> <sub>0.63</sub>
GRIP	88.84 <sub>0.75</sub>	<b>70.93</b> <sub>2.08</sub>	<b>78.80</b> <sub>1.26</sub>	84.47 <sub>0.41</sub>	<b>84.09</b> <sub>1.01</sub>	<b>84.28</b> <sub>0.73</sub>	22.13 <sub>0.24</sub>	28.32 <sub>0.33</sub>	<b>24.84</b> <sub>0.05</sub>
$\Delta$ CLIP	$\uparrow$ 23.46	$\uparrow$ 10.29	$\uparrow$ 15.94	$\uparrow$ 27.83	$\uparrow$ 25.45	$\uparrow$ 26.61	$\uparrow$ 4.83	$\uparrow$ 10.20	$\uparrow$ 7.14
$\Delta$ CoOp	$\downarrow$ 2.96	$\uparrow$ 23.18	$\uparrow$ 16.09	$\downarrow$ 2.07	$\uparrow$ 36.09	$\uparrow$ 22.58	$\downarrow$ 11.46	$\uparrow$ 8.75	$\uparrow$ 0.21
Split 2									
Method	Flowers102			RESICS45			FGVCAircraft		
	S	U	H	S	U	H	S	U	H
CLIP	15.99 <sub>0.00</sub>	39.18 <sub>0.00</sub>	22.71 <sub>0.00</sub>	32.47 <sub>0.00</sub>	33.10 <sub>0.00</sub>	32.78 <sub>0.00</sub>	45.43 <sub>0.00</sub>	39.72 <sub>0.00</sub>	42.39 <sub>0.00</sub>
CoOp	<b>90.6</b> <sub>13.02</sub>	18.77 <sub>9.12</sub>	30.29 <sub>12.38</sub>	86.43 <sub>3.23</sub>	47.16 <sub>11.17</sub>	60.53 <sub>8.42</sub>	<b>70.4</b> <sub>1.99</sub>	32.53 <sub>4.58</sub>	44.42 <sub>4.63</sub>
GRIP	<b>95.71</b>	<b>97.50</b>	<b>96.59</b>	<b>91.08</b> <sub>0.02</sub>	<b>92.02</b> <sub>0.98</sub>	<b>91.55</b> <sub>0.47</sub>	66.69 <sub>0.53</sub>	<b>56.19</b> <sub>1.18</sub>	<b>60.99</b> <sub>0.69</sub>
$\Delta$ CLIP	$\uparrow$ 85.12	$\uparrow$ 50.76	$\uparrow$ 79.32	$\uparrow$ 58.61	$\uparrow$ 58.92	$\uparrow$ 58.77	$\uparrow$ 21.26	$\uparrow$ 16.47	$\uparrow$ 18.6
$\Delta$ CoOp	$\uparrow$ 6.11	$\uparrow$ 71.20	$\uparrow$ 57.19	$\uparrow$ 4.65	$\uparrow$ 44.86	$\uparrow$ 31.02	$\downarrow$ 3.71	$\uparrow$ 23.66	$\uparrow$ 16.57
Split 3									
Method	Flowers102			RESICS45			FGVCAircraft		
	S	U	H	S	U	H	S	U	H
CLIP	68.29 <sub>0.00</sub>	57.25 <sub>0.00</sub>	62.28 <sub>0.00</sub>	56.02 <sub>0.00</sub>	52.32 <sub>0.00</sub>	54.10 <sub>0.00</sub>	17.55 <sub>0.00</sub>	17.71 <sub>0.00</sub>	17.63 <sub>0.00</sub>
CoOp	<b>91.52</b> <sub>0.35</sub>	48.35 <sub>2.95</sub>	63.22 <sub>2.60</sub>	<b>87.61</b> <sub>2.17</sub>	43.64 <sub>4.97</sub>	58.14 <sub>4.12</sub>	<b>37.77</b> <sub>1.92</sub>	16.46 <sub>3.23</sub>	22.77 <sub>3.09</sub>
GRIP	90.09 <sub>0.53</sub>	<b>69.00</b> <sub>2.44</sub>	<b>78.13</b> <sub>1.71</sub>	85.19 <sub>0.15</sub>	<b>75.58</b> <sub>3.17</sub>	<b>80.07</b> <sub>1.79</sub>	22.07 <sub>0.23</sub>	<b>28.72</b> <sub>0.76</sub>	<b>24.95</b> <sub>0.20</sub>
$\Delta$ CLIP	$\uparrow$ 21.8	$\uparrow$ 11.75	$\uparrow$ 15.85	$\uparrow$ 29.17	$\uparrow$ 23.26	$\uparrow$ 25.97	$\uparrow$ 4.52	$\uparrow$ 11.01	$\uparrow$ 7.32
$\Delta$ CoOp	$\downarrow$ 1.43	$\uparrow$ 20.65	$\uparrow$ 14.91	$\downarrow$ 2.42	$\uparrow$ 31.94	$\uparrow$ 21.93	$\downarrow$ 15.70	$\uparrow$ 12.26	$\uparrow$ 2.18
Split 3									
Method	Flowers102			RESICS45			FGVCAircraft		
	S	U	H	S	U	H	S	U	H
CLIP	10.59 <sub>0.00</sub>	46.74 <sub>0.00</sub>	17.27 <sub>0.00</sub>	41.47 <sub>0.00</sub>	19.60 <sub>0.00</sub>	26.62 <sub>0.00</sub>	45.52 <sub>0.00</sub>	39.58 <sub>0.00</sub>	42.34 <sub>0.00</sub>
CoOp	89.6 <sub>8.08</sub>	26.3 <sub>12.88</sub>	39.4 <sub>16.61</sub>	79.33 <sub>9.37</sub>	43.38 <sub>12.49</sub>	55.06 <sub>8.62</sub>	<b>70.53</b> <sub>3.11</sub>	24.94 <sub>5.37</sub>	36.63 <sub>5.57</sub>
GRIP	<b>95.8</b>	<b>96.06</b>	<b>95.93</b>	<b>90.57</b> <sub>0.13</sub>	<b>94.25</b> <sub>1.10</sub>	<b>92.37</b> <sub>0.60</sub>	67.28 <sub>0.74</sub>	<b>58.94</b> <sub>2.78</sub>	<b>62.81</b> <sub>1.75</sub>
$\Delta$ CLIP	$\uparrow$ 79.81	$\uparrow$ 56.88	$\uparrow$ 73.22	$\uparrow$ 49.1	$\uparrow$ 74.65	$\uparrow$ 65.75	$\uparrow$ 21.76	$\uparrow$ 19.36	$\uparrow$ 20.47
$\Delta$ CoOp	$\uparrow$ 5.20	$\uparrow$ 77.29	$\uparrow$ 65.64	$\uparrow$ 11.24	$\uparrow$ 50.87	$\uparrow$ 37.31	$\downarrow$ 3.25	$\uparrow$ 34.00	$\uparrow$ 26.18

Table 8: In the TRZSL settings, for each dataset and split, we compare the accuracy of GRIP textual with CLIP zero-shot (ViT-B/32), and CoOp. Results show the accuracy on seen ( $S$ ) and unseen classes ( $U$ ), and the harmonic mean ( $H$ ). We average the accuracy on 5 seeds and report the standard deviation.  $\Delta$  METHOD is the difference between the accuracy of GRIP and METHOD.

Split 1	Seen classes ( <i>S</i> )	Unseen classes ( <i>U</i> )
Flowers102	canna lily, petunia, silverbush, prince of wales feathers, pincushion flower, bird of paradise, frangipani, hard-leaved pocket orchid, bearded iris, passion flower, tiger lily, lenten rose, cape flower, air plant, mexican petunia, common dandelion, magnolia, foxglove, hibiscus, camellia, orange dahlia, clematis, anthurium, bougainvillea, ruby-lipped cattleya, stemless gentian, oxeye daisy, spring crocus, king protea, cyclamen, fritillary, californian poppy, wild pansy, desert-rose, sunflower, rose, grape hyacinth, pink primrose, red ginger, corn poppy, watercress, colt's foot, blanket flower, monkshood, morning glory, siam tulip, barbeton daisy, bolero deep blue, carnation, tree poppy, globe thistle, english marigold, primula, wallflower, blackberry lily, fire lily, love in the mist, moon orchid, sweet pea, mallow, pelargonium, mexican aster, poinsettia	canterbury bells, snapdragon, spear thistle, yellow iris, globe flower, purple coneflower, peruvian lily, balloon flower, giant white arum lily, artichoke, sweet william, garden phlox, alpine sea holly, great masterwort, daffodil, sword lily, marigold, buttercup, bishop of llandaff, gaura, geranium, pink and yellow dahlia, cauleya spicata, japanese anemone, black-eyed susan, osteospermum, windflower, gazania, azalea, water lily, thorn apple, lotus, toad lily, columbine, tree mallow, hippeastrum, bee balm, bromelia, trumpet creeper
RESICS45	beach, palace, roundabout, railway station, railway, thermal power station, river, airplane, island, bridge, basketball court, desert, runway, ground track field, sea ice, sparse residential, cloud, dense residential, wetland, mountain, meadow, baseball diamond, parking lot, storage tank, tennis court, commercial area, mobile home park	airport, ship, snowberg, chaparral, church, circular farmland, stadium, terrace, forest, freeway, golf course, harbor, industrial area, intersection, lake, medium residential, overpass, rectangular farmland
FGVC-Aircraft	Tu-134, Spitfire, Challenger 600, 737-700, F-A-18, E-170, 727-200, A300B4, Falcon 2000, DR-400, MD-87, CRJ-700 ERJ 145, Falcon 900, MD-80, DC-10, Il-76, Global Express, Gulfstream IV, Saab 340, Yak-42, CRJ-900, L-1011, A330-200, A321, 747-300, DC-3, A310, ATR-42, CRJ-200, Hawk T1, Fokker 100, ATR-72, PA-28, A319, 707-320, A318, A320, BAE-125, 747-200, ERJ 135, 737-800, SR-20, BAE 146-300, Beechcraft 1900, Cessna 172, A340-300, EMB-120, 737-900, 737-400, Cessna 208, MD-90, 777-300, A340-600, 737-600, 737-300, DHC-1, DC-6, A380, C-47, 767-200, BAE 146-200	737-200, 737-500, 747-100, 747-400, 757-200, 757-300, 767-300, 767-400, 777-200, A330-300, A340-200, A340-500, An-12, Boeing 717, C-130, Cessna 525, Cessna 560, DC-8, DC-9-30, DH-82, DHC-6, DHC-8-100, DHC-8-300, Dornier 328, E-190, E-195, Embraer Legacy 600, Eurofighter Typhoon, F-16A-B, Fokker 50, Fokker 70, Gulfstream V, MD-11, Metroliner, Model B200, Saab 2000, Tornado, Tu-154
MNIST	4, 2, 9, 3, 0, 5	8, 1, 6, 7
EuroSAT	industrial buildings or commercial buildings, brushland or shrubland, lake or sea, highway or road, annual crop land, pasture land	river, forest, permanent crop land, residential buildings or homes or apartments
DTD	knitted, pitted, studded, bumpy, spiralled, scaly, polka-dotted, veined, wrinkled, banded, flecked, stained, chequered, sprinkled, bubbly, grid, lined, crystalline, fibrous, meshed, zigzagged, pleated, braided, perforated, potholed, waffled, dotted, matted, gauzy	blotchy, smeared, cobwebbed, cracked, crosshatched, stratified, striped, swirly, woven, freckled, frilly, grooved, honeycombed, interlaced, lacelike, marbled, paisley, porous
Split 2		
Flowers102	prince of wales feathers, air plant, canterbury bells, bishop of llandaff, bee balm, desert-rose, purple coneflower, spring crocus, pelargonium, windflower, sunflower, bougainvillea, rose, spear thistle, bird of paradise, carnation, fritillary, grape hyacinth, mexican aster, monkshood, poinsettia, black-eyed susan, sweet pea, anthurium, wallflower, oxeye daisy, moon orchid, blackberry lily, hibiscus, frangipani, cauleya spicata, camellia, canna lily, passion flower, wild pansy, stemless gentian, balloon flower, gaura, thorn apple, morning glory, hard-leaved pocket orchid, japanese anemone, sword lily, daffodil, english marigold, globe flower, peruvian lily, barbeton daisy, siam tulip, tiger lily, foxglove, pink and yellow dahlia, pink primrose, alpine sea holly, artichoke, petunia, colt's foot, ruby-lipped cattleya, red ginger, primula, snapdragon, garden phlox, mexican petunia	globe thistle, king protea, yellow iris, giant white arum lily, fire lily, pincushion flower, corn poppy, sweet william, love in the mist, cape flower, great masterwort, lenten rose, bolero deep blue, marigold, buttercup, common dandelion, geranium, orange dahlia, silverbush, californian poppy, osteospermum, bearded iris, tree poppy, gazania, azalea, water lily, lotus, toad lily, clematis, columbine, tree mallow, magnolia, cyclamen, watercress, hippeastrum, mallow, bromelia, blanket flower, trumpet creeper
RESICS45	railway station, snowberg, palace, beach, commercial area, mountain, parking lot, dense residential, sparse residential, rectangular farmland, railway, island, tennis court, baseball diamond, thermal power station, industrial area, golf course, meadow, ground track field, storage tank, circular farmland, forest, bridge, harbor, river, freeway, sea ice	airplane, airport, roundabout, basketball court, runway, ship, chaparral, church, stadium, cloud, terrace, desert, wetland, intersection, lake, medium residential, mobile home park, overpass
FGVC-Aircraft	A321, MD-80, 737-200, DC-8, Falcon 900, Saab 340, 767-200, F-A-18, DC-6, SR-20, DC-3, Saab 2000, Fokker 70, 747-400, 737-700, A340-300, A310, A319, A380, 737-800, C-47, Dornier 328, 737-300, Eurofighter Typhoon, Cessna 208, Challenger 600, 737-600, Yak-42, Hawk T1, Fokker 100, DHC-8-100, Gulfstream IV, Model B200, Embraer Legacy 600, CRJ-900, A330-200, 767-400, DC-9-30, DR-400, Falcon 2000, 727-200, DHC-8-300, C-130, Boeing 717, 737-400, 757-300, 767-300, Beechcraft 1900, BAE 146-300, 737-500, PA-28, DHC-6, 707-320, An-12, A330-300, CRJ-700, 747-200, ATR-42, A318, DC-10, 747-100, A340-500	737-900, 747-300, 757-200, 777-200, 777-300, A300B4, A320, A340-200, A340-600, ATR-72, BAE 146-200, BAE-125, Cessna 172, Cessna 525, Cessna 560, CRJ-200, DH-82, DHC-1, E-170, E-190, E-195, EMB-120, ERJ 135, ERJ 145, F-16A-B, Fokker 50, Global Express, Gulfstream V, Il-76, L-1011, MD-11, MD-87, MD-90, Metroliner, Spitfire, Tornado, Tu-134, Tu-154
MNIST	2, 8, 4, 9, 1, 6	0, 3, 5, 7
EuroSAT	brushland or shrubland, river, industrial buildings or commercial buildings, lake or sea, forest, permanent crop land	annual crop land, highway or road, pasture land, residential buildings or homes or apartments
DTD	pitted, scaly, polka-dotted, bumpy, honeycombed, fibrous, veined, porous, lined, dotted, perforated, potholed, pleated, waffled, braided, wrinkled, paisley, gauzy, meshed, grid, studded, knitted, swirly, crosshatched, freckled, chequered, grooved, smeared, frilly	banded, blotchy, bubbly, spiralled, sprinkled, cobwebbed, cracked, stained, crystalline, stratified, striped, flecked, woven, zigzagged, interlaced, lacelike, marbled, matted
Split 3		
Flowers102	oxeye daisy, canterbury bells, clematis, siam tulip, cape flower, black-eyed susan, air plant, californian poppy, globe thistle, giant white arum lily, cyclamen, snapdragon, frangipani, buttercup, common dandelion, hippeastrum, columbine, spring crocus, bolero deep blue, spear thistle, barbeton daisy, poinsettia, peruvian lily, alpine sea holly, artichoke, sunflower, tiger lily, toad lily, magnolia, lenten rose, great masterwort, camellia, mallow, morning glory, lotus, sweet william, thorn apple, carnation, daffodil, corn poppy, cauleya spicata, marigold, hibiscus, tree poppy, balloon flower, osteospermum, english marigold, king protea, azalea, foxglove, watercress, blackberry lily, bearded iris, monkshood, mexican aster, orange dahlia, water lily, mexican petunia, sweet pea, pink primrose, primula, silverbush, pincushion flower	hard-leaved pocket orchid, moon orchid, bird of paradise, colt's foot, yellow iris, globe flower, purple coneflower, fire lily, fritillary, red ginger, grape hyacinth, prince of wales feathers, stemless gentian, garden phlox, love in the mist, ruby-lipped cattleya, sword lily, wallflower, petunia, wild pansy, pelargonium, bishop of llandaff, gaura, geranium, pink and yellow dahlia, japanese anemone, windflower, gazania, rose, passion flower, anthurium, desert-rose, tree mallow, canna lily, bee balm, bougainvillea, bromelia, blanket flower, trumpet creeper
RESICS45	railway, parking lot, wetland, meadow, harbor, island, mobile home park, storage tank, industrial area, bridge, baseball diamond, sea ice, runway, airplane, thermal power station, circular farmland, basketball court, roundabout, commercial area, railway station, terrace, forest, rectangular farmland, lake, medium residential, snowberg, river	airport, beach, ship, chaparral, church, sparse residential, cloud, stadium, dense residential, desert, tennis court, freeway, golf course, ground track field, intersection, mountain, overpass, palace
FGVC-Aircraft	An-12, 737-200, F-16A-B, BAE 146-200, MD-80, E-170, Gulfstream IV, DR-400, 737-900, 777-200, Boeing 717, 747-100, Saab 340, Cessna 525, Challenger 600, MD-90, DHC-8-100, Cessna 172, C-47, 747-400, BAE-125, MD-11, 767-300, Cessna 560, A330-300, E-195, 737-500, Fokker 50, ATR-72, BAE 146-300, Fokker 70, Falcon 900, Falcon 2000, Spitfire, A340-200, DC-3, A340-300, Beechcraft 1900, A320, Hawk T1, E-190, Gulfstream V, Tu-134, 767-400, CRJ-200, 737-400, 747-300, Eurofighter Typhoon, PA-28, MD-87, Yak-42, DHC-1, 737-800, A380, Model B200, ERJ 135, SR-20, 737-300, 707-320, DC-10, Dornier 328, A300B4	727-200, 737-600, 737-700, 747-200, 757-200, 757-300, 767-200, 777-300, A310, A318, A319, A321, A330-200, A340-500, A340-600, ATR-42, C-130, Cessna 208, CRJ-700, CRJ-900, DC-6, DC-8, DC-9-30, DH-82, DHC-6, DHC-8-300, EMB-120, Embraer Legacy 600, ERJ 145, F-A-18, Fokker 100, Global Express, Il-76, L-1011, Metroliner, Saab 2000, Tornado, Tu-154
MNIST	8, 3, 5, 6, 1, 7	0, 9, 2, 4
EuroSAT	river, highway or road, pasture land, permanent crop land, forest, residential buildings or homes or apartments	annual crop land, lake or sea, brushland or shrubland, industrial buildings or commercial buildings
DTD	pitted, pleated, polka-dotted, sprinkled, grooved, knitted, matted, wrinkled, honeycombed, chequered, braided, zigzagged, spiralled, banded, waffled, crosshatched, bubbly, smeared, dotted, porous, woven, freckled, lined, potholed, lacelike, marbled, stratified, scaly, studded	blotchy, bumpy, stained, cobwebbed, cracked, striped, crystalline, swirly, fibrous, flecked, veined, frilly, gauzy, grid, interlaced, meshed, paisley, perforated

Table 9: For each dataset, we report the class names of seen and unseen classes in each of the splits used for TRZSL.