

# LAVA: Label-efficient Visual Learning and Adaptation

Islam Nassar<sup>1\*</sup>, Munawar Hayat<sup>1</sup>, Ehsan Abbasnejad<sup>2</sup>, Hamid Rezatofghi<sup>1</sup>,  
Mehrtash Harandi<sup>1</sup>, Gholamreza Haffari<sup>1</sup>

<sup>1</sup> Monash University, Australia    <sup>2</sup> University of Adelaide, Australia

## Abstract

We present LAVA, a simple yet effective method for multi-domain visual transfer learning with limited data. LAVA builds on a few recent innovations to enable adapting to partially labelled datasets with class and domain shifts. First, LAVA learns self-supervised visual representations on the source dataset and ground them using class label semantics to overcome transfer collapse problems associated with supervised pretraining. Secondly, LAVA maximises the gains from unlabelled target data via a novel method which uses multi-crop augmentations to obtain highly robust pseudo-labels. By combining these ingredients, LAVA achieves a new state-of-the-art on ImageNet semi-supervised protocol, as well as on 7 out of 10 datasets in multi-domain few-shot learning on the Meta-dataset.<sup>1</sup>

## 1. Introduction

Using limited data to effectively adapt to new tasks is a challenging but essential requirement for modern deep learning systems. It enables leveraging the power of such systems while avoiding excessive data annotation which is usually costly, time consuming, and often requires domain expertise [9, 17, 52]. A promising direction is to develop methods that are capable of transferring knowledge across collective data of many tasks. In this work, we examine low-resource multi-domain visual transfer: given a visual learner pretrained on a source dataset<sup>2</sup>, our goal is to effectively transfer to a target dataset with potential class and/or domain shift. We focus on low-resource cases where the target dataset is very small but fully labelled (as in few-shot learning - FSL); or cases where it is sufficiently large but only partially labelled (as in semi-supervised learning - SSL). We propose a transfer method which simultaneously addresses such cases and demonstrates a strong performance on multiple transfer benchmarks.

\*corresponding author: islam.nassar@monash.edu

<sup>1</sup>Code: [github.com/islam-nassar/lava.git](https://github.com/islam-nassar/lava.git)

<sup>2</sup>We use ImageNet [40] as the source dataset in all our experiments.

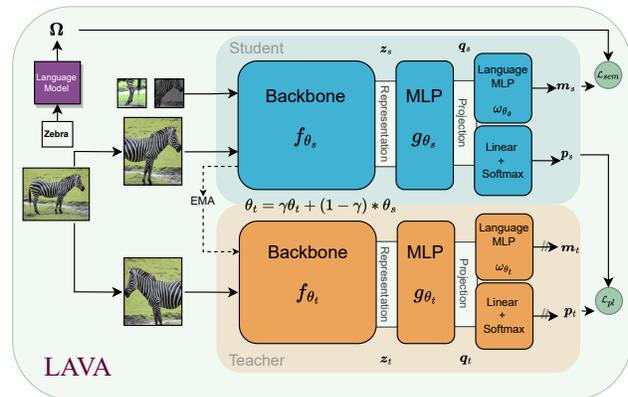


Figure 1: **Method Overview.** LAVA uses the source dataset to learn self-supervised initialisations as well as a mapping between visual features and label semantics. During transfer, LAVA uses a teacher-student setup to adapt by multi-crop pseudo-labeling unlabelled instances, and by matching the semantics of labelled ones.

LAVA’s first design goal is to employ a pretraining strategy which supports generalisation beyond classes and domains with limited labelled data. Hence, we begin by investigating the effect of source pretraining on the visual transfer performance. In line with recent research, we find that supervised pretraining leads to sub-optimal transfer [2, 12, 30]. Supervision with labels, more often than not, is too eager to learn specialised features which can successfully discriminate the source classes/domains but fails to generalise beyond them. We argue that its other limitation is the semantic-free nature of the labels used. Labels are represented using one-hot vectors explicitly encouraging to ignore label semantics. For example, the learner is encouraged to treat “bus” and “school bus” as two unrelated classes. In Fig. 2, we use DomainNet [33] clipart dataset to qualitatively demonstrate two artifacts associated with supervised pretrained representations: class collapse, whereby the pretrained representations collapse into incorrect source classes just because they share superficial similarities with target classes; and domain collapse, where the class semantics are preserved but the visual domain information is disregarded. LAVA’s first ingredient is introduc-



Figure 2: **Class & Domain collapse.** We display two query images from clipart target (left), and their 4 nearest neighbours, from clipart and ImageNet, in LAVA’s representation space (middle) and an ImageNet-supervised learner’s (right) *before training on target instances*. ImageNet instances are bordered in blue. We observe that the supervised learner suffers from two types of representation collapse: (i) class collapse (upper), where the learner picks an irrelevant source class which shares superficial patterns with target class; or (ii) domain collapse (lower), where the learner picks a relevant source class but ignores the visual domain. Please refer to Sec. 5 for more details.

ing a two-fold approach to address the collapse problem: 1) self-supervised source pretraining to learn task-independent features leading to better transfer; and 2) using the language modality to ground the self-supervised representations to an independent semantic space: during pretraining, LAVA uses source class labels to learn a mapping between the visual representations of the source instances and the language semantic representation of their class labels. At transfer time, such mapping is used to infer relations between “seen” and “unseen” classes by virtue of their foreknown semantic similarities. (Sec. 3.1)

LAVA’s second design goal is to leverage unlabelled target data (if available) to improve transfer performance. For that, we employ multi-crop augmentation which was originally proposed to encourage learning representations that are invariant to spatial augmentations for self-supervised contrastive learning [8]. We extend the idea to semi-supervised learning exploiting the observation that images can often contain multiple semantic classes and hence using a single label per image can hurt the performance. Consider the image in Fig. 3, due to random cropping used during training, it is conceivable that the teacher model (performing the pseudo-labeling) receives a view centered around a different object than the student, leading to a label that is not compatible with the image. Therefore, we calculate pseudo-labels based on multiple local and global views of images to account for those containing multiple semantic concepts. We show that enforcing a single pseudo-label per image is sub-optimal. Instead, applying a pair-wise aggregate loss across multiple views enhances the quality of pseudo-labels. (Sec. 3.2)

LAVA’s main contributions are: **1)** a practical method which combines and extends a few recent innovations to address various transfer learning regimes; **2)** provides empirical insights about transfer collapse problems associated with supervised pretraining and proposes a strategy to ad-

dress them; **3)** extends the multi-crop augmentation strategy to the semi-supervised setting via pseudo-labeling; **4)** sets a new state-of-the-art on the ImageNet [40] SSL benchmark and demonstrates strong performance on other challenging SSL and FSL benchmarks including the Meta-Dataset [49].

## 2. Related Work

**Few-Shot learning (FSL)** Existing FSL approaches can be categorized into metric-learning methods that learn an embedding space to compare query and support samples [23, 32, 43, 51], meta-learning approaches that adapt a base learner to new classes [7, 14, 20, 25, 27, 37, 41], or a combination of both [49]. Most existing FSL methods work well when the train and test sets are from the same domain *e.g.* subsets of ImageNet (mini-ImageNet [51], and tiered-ImageNet [38]). They lack out-of-domain (OOD) generalization, once there is a distribution mismatch between train and test data. Recently introduced meta-dataset [49] provides a challenging benchmark for multi-domain FSL. FSL methods developed on meta-dataset therefore aim to tackle OOD generalization [12, 28]. For instance, Transformer module is employed to capture relationships between different domains in [12, 28]. Even though our approach is generic across different low-label regimes, our results on meta-dataset show that we perform favorably against recent approaches which are specifically developed for FSL.

**Semi-supervised learning.** A common approach for SSL is to train the model with a joint loss formulation *i.e.* a supervised cross-entropy loss for the labelled samples, and a un/self-supervised regularization term for unlabelled samples. Examples include UDA [53], S4L [55] and [16, 29, 50]. Another approach to SSL is using pseudo-labels which are generated by either training the model on the labelled samples and pruning the confident predictions on unlabelled data [5, 6, 24, 30, 39, 44], or by us-

### Effect of cropping augmentations



### Images with multiple semantic objects



Figure 3: We display the effect of cropping (top), commonly used in vision training, to highlight our multi-crop pseudo-labelling motivation: cropping can potentially change the main focus to a different semantic object. Since in a teacher-student setup, each network receives a different view of the image, images with multiple semantic objects (bottom)<sup>3</sup> can hurt performance if we rely on a single pseudo-label per image.

ing a teacher-student configuration, where a slowly updating teacher model is used to generate soft-predictions on the unlabelled samples, which serve as a supervisory signal for the student model [35, 46, 54]. LAVA leverages the latter paradigm but improves pseudo-labels using a multi-crop augmentation strategy.

**Semantics and Self-supervision for FSL.** Rich semantics [1, 15, 36] and self-supervision [12, 21] have been explored to help FSL. [1] introduces an auxiliary task to produce class-level semantic descriptions and showed improvements on fine-grained tasks, while [36] learns rich features from large number of (image, text) pairs. On the other hand, recent work [12, 47] recognised the usefulness of self-supervised features and their ability to generalise, [21] adjusted the instance discrimination SimCLR [10] method to use the few support labels for positive and negative mining, while [12] explicitly adds SimCLR [10] episodes to the training pipeline. Unlike existing approaches, LAVA only uses class-level label semantics and employs self-supervision [9] as a pretraining step rather than incorporating it into the FSL task.

## 3. Method

We consider the problem of adapting a classifier pretrained on a set of source classes  $C_{src}$  using labelled samples  $\mathcal{D}_{src}$  to a target dataset of  $C_{tgt}$  classes by using labelled instances  $\mathcal{D}_{tgt} = \{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^n$ , and unlabelled instances  $\mathcal{U}_{tgt} = \{\mathbf{u}^j\}_{j=1}^m$ , with  $\mathbf{u}, \mathbf{x}$  denoting an unlabelled, and labelled image respectively, and  $\mathbf{y}^i$  is the class label. Note

<sup>3</sup>These images are not cherry-picked. They were identified in our analysis (see Sec. 5) to have inconsistent pseudo-labels.

that such setup suits both SSL and FSL settings. However, in FSL,  $C_{src}$  and  $C_{tgt}$  are strictly disjoint and the few-shot transfer utilises a fully labelled support set.

LAVA employs a teacher-student setup with a teacher identical in architecture to an online student (see Fig. 1). The student is trained to match the “soft” label generated by the teacher when each receives different views of a given image. The student and teacher networks, parametrized by  $\theta_s$  and  $\theta_t$  respectively, are updated alternatively: given a fixed teacher, the student is first updated by gradient descent to minimize the network loss; subsequently, the teacher parameters are updated as an exponential moving average (EMA) of the student’s, i.e.  $\theta_t \leftarrow \gamma\theta_t + (1 - \gamma)\theta_s$ , where  $\gamma$  is the momentum parameter.

### 3.1. Generalising beyond domains and classes

Our first aim is to use the source dataset to pretrain our teacher and student with good initial representations to support out-of-distribution transfer while avoiding the collapse problems mentioned in Sec. 1.

**Self-supervised Pretraining.** We employ the recently proposed DINO [9] method (without modification) to learn self-supervised initialisations from  $\mathcal{D}_{src}$  after discarding the labels. DINO, like other self-supervised methods [3, 10, 13, 18, 19, 26], learn visual features which are invariant to common factors of variation (e.g. colour distortion, pose, scale) while not being tied to a specific set of classes or visual domains. Therefore, they encode richer information which better supports generalisation. At transfer time, we use the target instances without their labels to further fine-tune the DINO representations to the target dataset. We provide in the appendix a detailed procedure of fine-tuning DINO fea-

tures to target instances in low-resource cases. We note that our method is not specifically tied to DINO. However, we chose it due to its demonstrated performance and its similar teacher-student setup making it seamless to integrate with our method. We ablate such choice in our experiments.

**Semantic Grounding.** To combat the class collapse problem and to aid generalisation to unseen classes, we employ language semantics as an independent modality to ground visual features. During the source pretraining, we additionally learn a semantic MLP module  $\omega_{\theta_*}$  (\* denotes s and t for student and teacher respectively) mapping the projection  $\mathbf{q}_*^i$  for a given labelled image  $\mathbf{x}^i$  to an embedding  $\mathbf{m}_*^i = \omega_{\theta_*}(\mathbf{q}_*^i)$ . We apply a hinge loss which minimizes the discrepancy between the semantic projection  $\mathbf{m}_s^i$ <sup>4</sup> and the corresponding embedding  $\Omega^i$  obtained by applying a pre-trained language model on the class label, as per:

$$\mathcal{L}_{sem} = \frac{1}{|\mathcal{D}_{src}|} \sum_{\mathbf{x}^i \in \mathcal{D}_{src}} \max(0, \eta - \langle \mathbf{m}_s^i, \Omega^i \rangle + \langle \mathbf{m}_s^i, \hat{\Omega}^i \rangle), \quad (1)$$

where  $\langle \cdot, \cdot \rangle$  denotes cosine similarity,  $\eta > 0$  is a scalar hinge margin, and  $\Omega^i$  and  $\hat{\Omega}^i$  are the language embeddings of the true and a randomly-sampled false class respectively. In effect, LAVA learns how to map the visual representations of the source instances to the language model representation space so that each instance is mapped closer to its true class language embedding and further away (up to a margin) from all other class embeddings. At transfer time,  $\omega_{\theta_*}$  is fine-tuned (together with the backbones) without re-initialisation. This is in contrast to the classifier head which must be re-initialised to match the target classes.

Unlike the discrete nature of one-hot class labels, using language embeddings to ground visual representations acts as a continuous space to represent class labels. In such space, “bench” is close to “park bench” and “zebra” is closer to “horse” than it is to “car”. This is intuitively useful to learn visual-semantic relations which enhance generalisation to novel concepts. However, it also implicitly assumes that linguistic similarity is always a good proxy for visual similarity, which is sometimes not true *e.g.* “wine glass” and “wine bottle” while linguistically similar, are usually visually distinct. Hence, we explored few alternatives for the source of semantics including knowledge graph embeddings [30], “Glove” [34] and other variants of language models (*see* appendix for details). We find that language models trained on paraphrasing tasks [45] provide the best performance in our setup. We conjecture that this is because it helps the model to unify similar visual concepts which appear under different names across datasets (*e.g.* “airplane” vs “plane”, “horse cart” vs “carriage”, *etc.*).

<sup>4</sup>We use the student semantic projection  $\mathbf{m}_s^i$  to apply the loss, while  $\mathbf{m}_t^i$  is only used during inference.

### 3.2. Multi-crop pseudo-labelling

When transferring to a partially labelled dataset, LAVA leverages unlabelled samples by using the teacher to iteratively produce pseudo-labels to expand the labelled samples used to train the student. We differ from previous similar SSL approaches (*e.g.* MeanTeacher [46]) primarily in the way we generate the pseudo-labels: we encourage more robust pseudo-labels via multi-crop augmentation. We generate pseudo-labels based on a set  $\mathcal{T}^i$  of multiple sized crops of a given unlabelled image  $\mathbf{u}^i$ ; similarly, student predictions are generated based on another set  $\mathcal{S}^i$ <sup>5</sup>. The pseudo-label is then aggregated over the combined views.

Formally, LAVA uses the backbone  $f_{\theta_*}$  to map an unlabelled image  $\mathbf{u}^i$  to  $\mathbf{z}_*^i = f_{\theta_*}(\mathbf{u}^i)$ , followed by an MLP  $g_{\theta_*}$  mapping  $\mathbf{z}_*^i$  to a projection  $\mathbf{q}_*^i = g_{\theta_*}(\mathbf{z}_*^i)$ . Finally, using a linear layer followed by a temperature sharpened softmax,  $\mathbf{q}_*^i$  is normalised into a probability distribution  $\mathbf{p}_*^i \in \mathbb{R}^{\|\mathcal{C}_{tgt}\|}$ . Then, we apply our loss as per:  $\mathcal{L}_{pl} = \frac{1}{|\mathcal{U}|} \sum_{\mathbf{u}^i \in \mathcal{U}} \mathcal{L}_{multi}^i$ , where  $\mathcal{L}_{multi}^i$  represents the aggregate loss over targets  $\mathbf{p}_t^i | \mathbf{u}^i \in \mathcal{T}^i$  and predictions  $\mathbf{p}_s^i | \mathbf{u}^i \in \mathcal{S}^i$ .

**Design choices.** Aggregating multi-crop losses involves a few design alternatives such as: using a single pseudo-label (*e.g.* via voting) or averaging across the different crops; using hard pseudo-labels (*e.g.* using argmax or sampling) or soft pseudo-labels; and finally, the count, scale, and size of the crops are important hyperparameters as they respectively impact the diversity of pseudo-labels, the locality of the crops, and the memory consumption during training. Our empirical study found that using soft sharpened pseudo-labels and averaging over pairs of crops yields the best performance across different domains (refer to appendix for more details). More concretely, we eventually opted for:  $\mathcal{L}_{multi}^i = \frac{1}{|\mathcal{A}^i|} \sum_{(\tilde{\mathbf{u}}_s^j, \tilde{\mathbf{u}}_t^j) \in \mathcal{A}^i} -\mathbf{p}_t^j \log \mathbf{p}_s^j$ , where  $(\tilde{\mathbf{u}}_s^j, \tilde{\mathbf{u}}_t^j)$  is a pair of crops of  $\mathbf{u}^i$  passed to the student and teacher respectively; and  $\mathcal{A}^i = \{(\tilde{\mathbf{u}}_s^j, \tilde{\mathbf{u}}_t^j) | \tilde{\mathbf{u}}_s^j \in \mathcal{S}^i, \tilde{\mathbf{u}}_t^j \in \mathcal{T}^i\}$  is the set of all crop pairs.

## 4. Experiments

We evaluate LAVA against state-of-the-art (SOTA) baselines in three regimes: **1)** SSL transfer with domain shift on DomainNet [33]; **2)** SSL without domain shift on ImageNet [40], and **3)** multi-domain FSL on Meta-dataset [49].

**Training.** Unless otherwise stated, we use a batch size of 256 with a learning rate of  $5e^{-4}$  and Adam [22] optimizer with a cosine scheduler. For our multi-crop pseudo-labelling, we use 6 small scale crops and 2 large scale crops (different for teacher than student) following same scales in [9], and a teacher momentum  $\gamma = 0.996$ . We use *mpnet-base-v2* [45] language model<sup>6</sup> to obtain the label embed-

<sup>5</sup>Spatial augmentations (*e.g.* color jittering, random flips) are also applied on all the crops.

<sup>6</sup>[github.com/UKPLab/sentence-transformers](https://github.com/UKPLab/sentence-transformers)

Table 1: Comparison to semi-supervised baselines on four domains of the DomainNet dataset. We report average accuracy over 3 runs for three amounts of labelled instances. All the baselines are implemented in the same codebase using the same network backbone.

	Real			Clipart			Sketch			Quickdraw		
	2-shot	4-shot	8-shot									
FixMatch [44]	23.06	34.68	42.14	30.21	41.21	51.29	12.73	21.65	33.07	24.51	32.98	43.91
SemCo [30]	24.38	40.03	51.13	28.39	46.96	55.48	15.71	28.62	41.06	26.17	34.17	44.12
MeanTeacher++ [46]	51.44	66.16	68.77	46.02	52.43	63.09	25.8	38.79	51.16	29.78	39.12	47.11
LAVA (supervised init.)	57.47	<b>69.51</b>	<b>75.41</b>	38.45	53.05	64.74	36.15	45.15	52.15	32.61	41.67	48.44
LAVA (no semantic loss)	58.57	67.88	72.12	48.57	58.75	<b>65.18</b>	38.76	47.55	53.91	35.95	44.01	<b>54.91</b>
LAVA	<b>58.79</b>	68.04	72.19	<b>48.65</b>	<b>59.05</b>	65.08	<b>39.12</b>	<b>47.63</b>	<b>54.39</b>	<b>36.66</b>	<b>44.12</b>	54.75

dings for our semantic loss. We report accuracy based on the softmax classifier (*see* Fig. 1), but when relevant, we compare it with the K-Nearest Neighbour (KNN with  $K=20$ ) accuracy based on the representation  $z_t$  and/or the semantic accuracy obtained by applying a cosine classifier on the semantic projection  $m_t$ .

**SSL on DomainNet.** This dataset includes images from 6 visual domains spanning 345 object classes. We examine LAVA’s ability to transfer from ImageNet to 4 domains with decreasing similarity: *real*, *clipart*, *sketch*, and *quickdraw*. To ensure fixed settings across all baselines (*e.g.* labelled splits, backbone, learning rate schedule, *etc.*), we follow recommendations in [31] and re-implement the three closest baselines in our codebase. FixMatch [44] uses consistency regularization with confidence-based pseudo-labelling, SemCo [30] builds on FixMatch but leverages label semantics (via a knowledge graph) to account for pre-known similarities among classes, and MeanTeacher [46] uses momentum teachers for SSL. We extend it to MeanTeacher++ where we employ the same spatial augmentations we use (instead of the original gaussian noise) and we use the same backbone as ours (ViT-S/16 [48]). For all experiments, we initialise models with pretrained ImageNet weights and follow the SSL standard approach: we use a fraction of the labelled data (expressed in images/class) together with all the unlabelled data to adapt to target. We fix the training to 70 epochs <sup>7</sup> (w.r.t unlabelled data) for LAVA and we use early stopping for all the baseline methods using a validation set. For each of the 4 domains, we examine low-, and moderate-shot transfer scenarios using 2, 4 and 8 images/class. We explore both self-supervised initialisation (*i.e.* DINO) and supervised initialisation [48] for all the baselines. We report the best results among the two for the baselines and both results for LAVA in Tab. 1. Finally, to examine the contribution of our semantic loss to SSL, we report results for LAVA when switching it off.

We observe LAVA outperforms baselines consistently, sometimes, with large margins. Interestingly, FixMatch and SemCo obtains their best results using supervised source

<sup>7</sup>The 70 epochs are split into 50 epochs of DINO target pretraining and 20 epochs for training LAVA

pretraining rather than self-supervised (*see* self-supervised results in appendix). One possible explanation is that this is due to the very different method of augmentations used in DINO pretraining compared to FixMatch and SemCo. As conjectured, we see that the impact of self-supervised initialisation for LAVA becomes more significant when the visual domain is different from that of ImageNet. For example, we observe an impressive 10% boost in the clipart 2-shot setting from 38.4% to 48.6%, proving that self-supervised features helps generalisation beyond domains. Among the baselines, MeanTeacher++ is the closest to LAVA; the main two differences are our multi-crop pseudo-labelling strategy and the semantic loss. We witness a significant boost of LAVA over MeanTeacher especially with fewer labelled samples. This confirms the usefulness of our multi-crop pseudo-labelling strategy in low-data regimes. Finally, we get a marginal boost when using semantic loss in SSL across almost all experiments.

**SSL on ImageNet.** To examine SSL transfer under the same domain, we follow ImageNet evaluation protocol by using 1% and 10% of the labels to train LAVA. Due to the demanding computational requirements for running experiments on ImageNet, we opted to only re-run SOTA method PAWS [2] with the same ViT-S backbone as ours. PAWS combines self-supervised learning with a non-parametric method of generating pseudo-labels based on a small labelled set. For all the other baselines, unless ViT-S results are reported in the original work, we only report the Resnet50 results. We note however, that the ViT-S model has less parameters (21M) compared to Resnet50 (24M), yet recent work [9, 48] showed an approximate 1-2% improvement in favour of ViT-S. Again, we see (Tab. 2) a significant boost for LAVA against other methods. Interestingly, as opposed to DINO, LAVA achieves large gains. Note, however, that DINO reports Linear evaluation results on frozen features and does not fine-tune end-to-end, so (64.5% and 72.2%) are not directly comparable to the (69.3% and 76.4%) of LAVA. However, K-NN results can be directly compared to measure the differential between LAVA and DINO.

**FSL on Meta-dataset.** We use the “ImageNet-only” pro-

Table 2: SSL results on ImageNet with 1% & 10% of the labels.

Method	Arch.	Epochs	1%	10%
<i>Different Architecture:</i>				
FixMatch [44]	RN50	300	-	71.5
MPL [35]	RN50	800	-	73.9
SwAV [8]	RN50	800	53.9	70.2
SimCLRv2++ [11]	RN50	1200	60.0	70.5
<i>Same Architecture:</i>				
DINO-NN [9]	ViT-S	300	61.3	69.1
DINO [9]	ViT-S	350	64.5	72.2
PAWS-NN [2]	ViT-S	300	63.5	72.3
PAWS [2]	ViT-S	300	68.9	75.2
LAVA-NN	ViT-S	300	<b>67.2</b>	<b>73.3</b>
LAVA	ViT-S	350	<b>69.3</b>	<b>76.4</b>

to evaluate multi-domain FSL on Meta-dataset [49]. Specifically, we use images of 712 out of 1000 classes of ImageNet as our source dataset, choose our hyperparameters by validating on 158 classes and evaluate using episodes coming from the remaining 130 classes of ImageNet as well as other 9 datasets. Meta-dataset measures cross-domain FSL by evaluating on datasets including fine-grained tasks (such as Birds, Aircrafts, Fungi, Flowers, Textures, Traffic Signs), characters and symbols (Omniglot), and real and quickly drawn objects (ImageNet, MSCoCO and Quickdraw). During source pretraining, we use all the instances coming from the 712 train classes without their labels for source pretraining<sup>8</sup>. And, we use the same instances with their labels to train the Semantic MLP ( $\omega_\theta$ ). During transfer, we freeze the backbone  $f_{\theta_s}$  and finetune LAVA using the support set for 300 epochs. We use the standard Meta-dataset setting where each episode contains varying number of ways (*i.e.* classes) and imbalanced shots (*i.e.* images per class). As per the common practice, we report results averaged over 600 episodes for each dataset.

As seen in Tab. 3, with such a simple strategy and without resorting to any meta-learning techniques, LAVA outperforms the closest baseline [12] on 7 out of 10 datasets and for the other 3, it exceeds the second best with a large margin. It is important to note here that it is not straight forward to directly compare between the FSL baselines: primarily, because different methods employ different styles of training (*e.g.* meta-training [12, 14, 43] vs finetuning), different initialisation (self-supervised [12] vs supervised [14, 43]) and different backbones (Resnet18 [42, 49] vs Resnet34 [12]). However, we think of our method as orthogonal to other methods: since we are using language semantics, which is not available for other methods, we are cautious about directly comparing with them and hence we report the results for reference and not comparison. A possible future direction is to explore how semantics (or other

<sup>8</sup>Unlike meta-learning methods, we do not use episodes during the training.

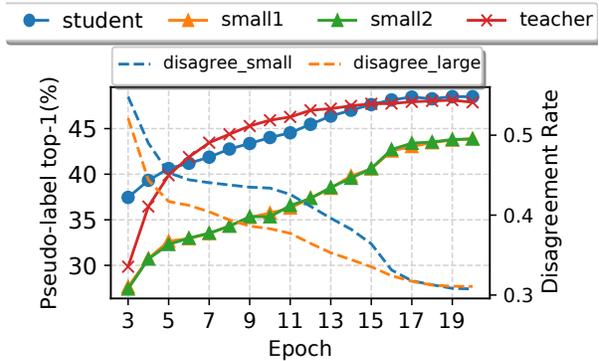


Figure 4: Multi-crop Pseudo-labelling analysis.

possible ways of grounding class relationships) help other strong methods such as “Cross Transformers”.

## 5. Analysis and Ablations

We gain key insights by analysing the following: **1)** the multi-crop strategy dynamics and its usefulness to produce higher quality pseudo-labels; **2)** the importance of self-supervised learning for out-of-domain generalisation; **3)** the role of language semantics towards class generalisation; and finally **4)** the effect of key hyperparameters.

**Multi-crop Pseudo-labelling.** To study LAVA’s dynamics, we conduct an “oracle” experiment where we use the ground truth labels (originally hidden to emulate an SSL setup) to calculate the true pseudo-labels accuracy of the multiple crops seen by the student and the teacher networks as the training proceeds. We calculate accuracies based on the argmax of the soft predictions  $p_s$  and  $p_t$  for the student and teacher respectively. Fig. 4 provides few interesting observations. First, based on the large crops, the teacher model exhibits a consistently better performance compared to the student after an initial ramp up phase (due to the EMA). This demonstrates the usefulness of the teacher-student setup, whereby the student model is always guided by a slightly better teacher. Second, as expected, we observe that the small crops (only seen by the student) have lower accuracy, on average, compared to large crops; but interestingly, the disagreement among the predictions associated with the small crops decrease as training proceeds, suggesting that the model is learning from the different small views of an image a consistent representation which truly captures its main object. Next, we repeat the same process for our closest baseline (MeanTeacher++) to examine the difference in pseudo-labelling quality obtained by each of the methods. Note that with the modifications we introduce to MeanTeacher++, the difference between the two methods in such setting is only the multi-crop pseudo-labelling. We observe in Tab. 4 that indeed, the multi-crop strategy brings a significant benefit across three different domains

Table 3: Results on Meta-Dataset when only pre-training on ImageNet train split. Mean accuracy, 95% confidence interval are reported over 600 test episodes. Our method outperforms the best method in 7 out of 10 datasets.

	K-NN [49]	MatchNet [49]	ProtoNet [49]	Finetune [49]	fo-Proto MAML [49]	BOHB [42]	ProtoNet-L [49]	ALFA-fo-Proto MAML[4]	CTX [12]	LAVA (ours)
ImageNet	41.03±1.01	45.00±1.10	50.50±1.08	45.78±1.10	49.53±1.05	51.92±1.05	53.69±1.07	52.80±1.11	62.76±0.99	<b>68.75±0.54</b>
Omniglot	37.07±1.15	52.27±1.28	59.98±1.35	60.85±1.58	63.37±1.33	67.57±1.21	68.50±1.27	61.87±1.51	<b>82.21±1.00</b>	77.92±0.50
Aircraft	46.81±0.89	48.97±0.93	53.10±1.00	68.69±1.26	55.95±0.99	54.12±0.90	58.04±0.96	63.43±1.10	79.49±0.89	<b>81.14±0.49</b>
Birds	50.13±1.00	62.21±0.95	68.79±1.01	57.31±1.26	68.66±0.96	70.69±0.90	74.07±0.92	69.75±1.05	80.63±0.88	<b>84.88±0.77</b>
Textures	66.36±0.75	64.15±0.85	66.56±0.83	69.05±0.90	66.49±0.83	68.34±0.76	68.76±0.77	70.78±0.88	75.57±0.64	<b>82.05±0.50</b>
Quick-Draw	32.06±1.08	42.87±1.09	48.96±1.08	42.60±1.17	51.52±1.00	50.33±1.04	53.30±1.06	59.17±1.16	<b>72.68±0.82</b>	68.44±0.57
Fungi	36.16±1.02	33.97±1.00	39.71±1.11	38.20±1.02	39.96±1.14	41.38±1.12	40.73±1.15	41.49±1.17	51.58±1.11	<b>55.02±0.67</b>
Flowers	83.10±0.68	80.13±0.71	85.27±0.77	85.51±0.68	87.15±0.69	87.34±0.59	86.96±0.73	85.96±0.77	<b>95.34±0.37</b>	<b>95.43±0.66</b>
Traffic-Sign	44.59±1.19	47.80±1.14	47.12±1.10	66.79±1.31	48.83±1.09	51.80±1.04	58.11±1.05	60.78±1.29	<b>82.65±0.76</b>	69.24±1.06
MSCOCO	30.38±0.99	34.99±1.00	48.03±0.99	41.00±1.10	34.86±0.97	43.74±1.12	41.70±1.08	48.11±1.14	59.90±1.02	<b>63.75±0.45</b>

Table 4: **Initialisation study.** Cross-domain performance when using different initialisation and different fine-tuning methods.

	Real		Clipart		Quickdraw	
	2-shot	8-shot	2-shot	8-shot	2-shot	8-shot
Fully-supervised	73.74	76.55	71.87	72.75	61.22	67.23
<i>Initialisation:</i>						
Sup. (ImNet)*	54.92	64.81	22.49	35.39	10.68	18.78
DINO (ImNet)	46.17	59.42	15.19	26.97	8.98	16.81
DINO (Target)**	50.45	62.54	40.35	55.03	30.05	43.64
<i>Fine-tuning from *:</i>						
Linear Probing	49.03	64.89	21.5	36	6.08	13.9
MeanTeacher++	54.26	70.60	34.58	60.22	16.38	35.75
LAVA	57.47	<b>75.41</b>	38.45	64.74	32.61	48.44
<i>Fine-tuning from **: </i>						
Linear Probing	49.68	64.42	38.86	56.6	28.73	47.5
MeanTeacher++	51.44	68.77	46.02	63.09	29.78	47.11
LAVA	<b>58.79</b>	72.19	<b>48.65</b>	<b>65.08</b>	<b>36.66</b>	<b>54.91</b>

especially when less labelled data is available. Additionally, we captured a fine-grained view of the pseudo-labels to examine what are the images that most differ in pseudo-labels among the two methods. As expected (*see* Fig. 3 - bottom), those are the images which contain multiple semantic objects. We provide further examples in the appendix.

**Initialisation Study.** Here, we are interested to examine the effect of pretraining when training LAVA using few labels across different visual domains. In Tab. 4, we report results using different initialisations and different fine-tuning settings across *real*, *clipart*, and *quickdraw* domains. In the top section, we display KNN accuracy based on the model representation  $z_t$  when initialised with **1)** Supervised ImageNet features [48]; **2)** DINO ImageNet features [9]; and **3)** DINO features when trained on the target dataset without labels (the standard LAVA initialisation). Note that for those results, the labels are only used to obtain the KNN accuracy on the validation set of the respective target dataset but never for fine-tuning so they are only meant to compare the quality of “off-the-shelf” pretrained features. First, we observe that using only ImageNet data, supervised pre-

training is more useful than self-supervised across the three domains and shots with a degrading performance as the domain deviates from ImageNet. Note that 23% of Domain-Net classes also exists in ImageNet, explaining why class-specific features might be helpful in such case. However, once we have access to the target instances (without their labels), we observe that self-supervised target training (*i.e.* LAVA initialisation) dramatically improves the representations to become more suited to the target domain without any labeling expense. Even when the target domain is very close to ImageNet (*e.g.* *real*), we see a 4% gain in the 2-shot regime (46.17 to 50.45). This boost is even more pronounced in highly dis-similar domains, *e.g.* 25% and 21% boosts for *clipart* and *quickdraw* with only 2 shots per class.

On the other hand, in the middle and bottom sections of Tab. 4, we report transfer results based on the supervised initialisation and the LAVA initialisation. We also report results using two other transfer methods: 1) Linear Probing on top of frozen representations [9]; and 2) MeanTeacher++ described in Sec. 4. Finally, as an upper bound, we report the fully-supervised results obtained when using the entire target dataset (with labels) to train ViT in a supervised manner, then using the few shots to obtain the KNN accuracy reported. Here, we observe that LAVA benefits from self-supervised DINO initialisation in all domains, but the gain is more clear when the target domain is different than ImageNet and the available labels are less. For example, we witness an impressive 10%, and 4% boosts in the 2-shot scenario for *clipart* and *quickdraw*, respectively. Additionally, to quantify what LAVA brings on top of the DINO initialisation, we compare LAVA with MeanTeacher++ and observe that LAVA outperforms it in all cases thanks to our multi-crop pseudo-labelling strategy. Finally, we notice that LAVA is almost closing the gap to fully-supervised training using all the target labels: by only using 8-shots, LAVA achieves 75.4% on the *real* domain compared to the 76.5% obtained when all the labels are used for training.

**Language Semantics.** Now, we examine the role of la-

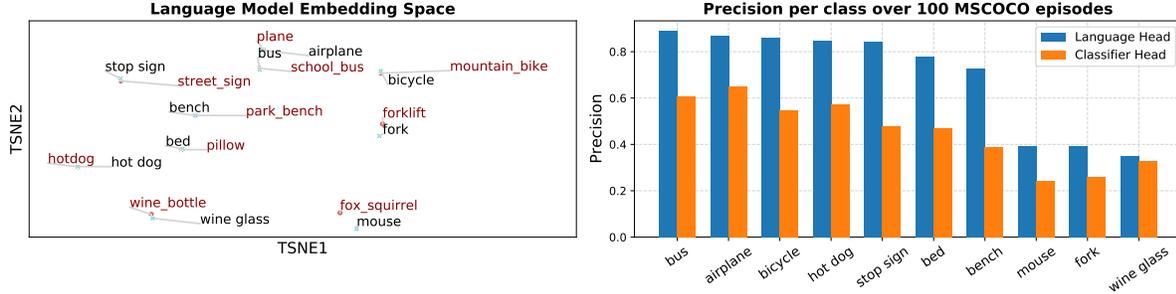


Figure 5: **Left:** we display the t-SNE visualization of the language model embeddings for a subset of MSCOCO classes (black) and their nearest neighbours in ImageNet train split (red). **Right:** we report precision per class over 100 FSL episodes for both the semantic and softmax predictions.

bel semantics towards generalisation to unseen classes and we investigate if LAVA indeed mitigates the “class collapse” problem. To study the impact of label semantics, we consider the 100 FSL episodes of MSCOCO in the Meta-dataset experiment. Recall that we first pretrain on instances of 712 classes of ImageNet  $C_{src}$  then transferred to a different set of classes in MSCOCO  $C_{tgt}$ . In Fig. 5-left, we display a subset of MSCOCO  $C \in C_{tgt}$  together with their nearest neighbors among  $C_{src}$ , when they are projected into the language model semantic space. This space is pretrained using language and so it captures the linguistic semantic similarity between different sentences/words. On the right plot, we display the average precision per class  $c \in C$  based on the softmax classification head as well as the semantic projection head<sup>9</sup>. We observe that for any given MSCOCO class, when the most semantically-similar class in ImageNet is also visually similar (e.g. “bus” and “school\_bus”), the language head has significantly higher precision per class. In contrast, in cases when the nearest neighbour is not visually similar (e.g. “wine glass” and “wine bottle”), both heads have comparable performance. This suggests the benefit of the learnt semantic mapping module  $\omega_\theta$ : during test time and without any further training, when the model receives an image sharing similar visual features associated with one of the source classes  $C_{src}$ , the semantic head maps it to the most closely related linguistic concept in  $C_{tgt}$ .

**Collapse Analysis.** We follow a similar setup as [12] to investigate LAVA’s ability to avoid transfer collapse (see Sec. 1): we begin by uniformly sampling 100 images per class from ImageNet  $C_{src}$  as well as 1000 query images from clipart dataset. Subsequently, we compute the representation  $z_t$  for all of the sampled images. Besides LAVA, we also compute representations obtained by a supervised learner; and a DINO-initialised learner. In Fig. 2, we report examples of query images with their 4 nearest neighbours among all the sampled images. The representation of a given query image is said to be collapsed if its nearest

<sup>9</sup>For a given image, semantic predictions are obtained by finding the class whose language embedding is nearest to  $m_t$ .

neighbours mostly belong to source classes. To quantify the collapse in the three methods, we calculate the percentage of the 10 nearest neighbours which are instances of source classes over the 1000 query images. We find that this figure is 25% for the supervised learner, 21.7% for DINO, and 17% for LAVA. We provide further details in the appendix.

**Hyperparameters sensitivity.** During preliminary experiments, we identified few important design choices to tune: number of large scale and small scale crops, the pseudo-labelling aggregation loss function, the teacher momentum update rate ( $\gamma$ ), and the source of label semantics embeddings. We used a held-out validation set on each of *real* and *clipart* domains to tune the hyperparameters (except the source of label semantics for which we used MSCOCO validation in the FSL regime), then we obtained a single set of parameters which we used across all experiments in this paper. We refer the reader to the appendix for a complete list of hyperparameters of LAVA in addition to a study to demonstrate their effect on LAVA’s performance.

## 6. Conclusion

We introduced a unified strategy for multi-domain visual transfer with limited target data. LAVA employs label semantics and self-supervised pretraining to learn initial representations which support generalisation; and uses multi-crop augmentation to maximise the gains from unlabeled data via pseudo-labeling. We demonstrated LAVA’s success in image classification over multiple benchmarks. We believe our approach is generic and can be extended to other visual learning tasks such as object detection and action recognition. We leave these explorations to future work.

## Acknowledgement

This work was partly supported by DARPA’s Learning with Less Labels (LwLL) program under agreement FA8750-19-2-0501 and by the Australian Government Research Training Program (RTP) Scholarship.

## References

- [1] Mohamed Afham, Salman Khan, Muhammad Haris Khan, Muzammal Naseer, and Fahad Shahbaz Khan. Rich semantics improve few-shot learning. *arXiv preprint arXiv:2104.12709*, 2021.
- [2] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Armand Joulin, Nicolas Ballas, and Michael Rabbat. Semi-supervised learning of visual features by non-parametrically predicting view assignments with support samples. *arXiv preprint arXiv:2104.13963*, 2021.
- [3] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *Advances in neural information processing systems*, 32, 2019.
- [4] Sungyong Baik, Myungsub Choi, Janghoon Choi, Heewon Kim, and Kyoung Mu Lee. Meta-learning with adaptive hyperparameters. *arXiv preprint arXiv:2011.00209*, 2020.
- [5] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785*, 2019.
- [6] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. *arXiv preprint arXiv:1905.02249*, 2019.
- [7] Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. *arXiv preprint arXiv:1805.08136*, 2018.
- [8] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882*, 2020.
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *arXiv preprint arXiv:2104.14294*, 2021.
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [11] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020.
- [12] Carl Doersch, Ankush Gupta, and Andrew Zisserman. Crosstransformers: spatially-aware few-shot transfer. *arXiv preprint arXiv:2007.11498*, 2020.
- [13] Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(9):1734–1747, 2015.
- [14] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017.
- [15] Andrea Frome, Greg Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. 2013.
- [16] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *Proceedings of the 17th International Conference on Neural Information Processing Systems*, pages 529–536, 2004.
- [17] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020.
- [19] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*, 2018.
- [20] Muhammad Abdullah Jamal and Guo-Jun Qi. Task agnostic meta-learning for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11719–11727, 2019.
- [21] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [23] Gregory Koch, Richard Zemel, Ruslan Salakhutdinov, et al. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille, 2015.
- [24] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896, 2013.
- [25] Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10657–10665, 2019.
- [26] Chunyuan Li, Jianwei Yang, Pengchuan Zhang, Mei Gao, Bin Xiao, Xiyang Dai, Lu Yuan, and Jianfeng Gao. Efficient self-supervised vision transformers for representation learning. *arXiv preprint arXiv:2106.09785*, 2021.
- [27] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017.
- [28] Lu Liu, William Hamilton, Guodong Long, Jing Jiang, and Hugo Larochelle. A universal representation transformer layer for few-shot image classification. *arXiv preprint arXiv:2006.11702*, 2020.
- [29] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization

- method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018.
- [30] Islam Nassar, Samitha Herath, Ehsan Abbasnejad, Wray Buntine, and Gholamreza Haffari. All labels are not created equal: Enhancing semi-supervision via label grouping and co-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7241–7250, 2021.
- [31] Avital Oliver, Augustus Odena, Colin Raffel, Ekin D Cubuk, and Ian J Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. *arXiv preprint arXiv:1804.09170*, 2018.
- [32] Boris N Oreshkin, Pau Rodriguez, and Alexandre Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. *arXiv preprint arXiv:1805.10123*, 2018.
- [33] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019.
- [34] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [35] Hieu Pham, Zihang Dai, Qizhe Xie, and Quoc V Le. Meta pseudo labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11557–11568, 2021.
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021.
- [37] Avinash Ravichandran, Rahul Bhotika, and Stefano Soatto. Few-shot learning with embedded class models and shot-free meta training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 331–339, 2019.
- [38] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. *arXiv preprint arXiv:1803.00676*, 2018.
- [39] Ellen Riloff. Automatically generating extraction patterns from untagged text. In *Proceedings of the national conference on artificial intelligence*, pages 1044–1049, 1996.
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [41] Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. *arXiv preprint arXiv:1807.05960*, 2018.
- [42] Tonmoy Saikia, Thomas Brox, and Cordelia Schmid. Optimized generic feature learning for few-shot classification across domains. *arXiv preprint arXiv:2001.07926*, 2020.
- [43] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 4080–4090, 2017.
- [44] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020.
- [45] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. MpNet: Masked and permuted pre-training for language understanding. *arXiv preprint arXiv:2004.09297*, 2020.
- [46] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1195–1204, 2017.
- [47] Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B Tenenbaum, and Phillip Isola. Rethinking few-shot image classification: a good embedding is all you need? In *European Conference on Computer Vision*, pages 266–282. Springer, 2020.
- [48] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021.
- [49] Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, et al. Meta-dataset: A dataset of datasets for learning to learn from few examples. *arXiv preprint arXiv:1903.03096*, 2019.
- [50] Vikas Verma, Kenji Kawaguchi, Alex Lamb, Juho Kannala, Arno Solin, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. *Neural Networks*, 2021.
- [51] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29:3630–3638, 2016.
- [52] Longhui Wei, Lingxi Xie, Jianzhong He, Jianlong Chang, Xiaopeng Zhang, Wengang Zhou, Houqiang Li, and Qi Tian. Can semantic labels assist self-supervised visual representation learning? *arXiv preprint arXiv:2011.08621*, 2020.
- [53] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33, 2020.
- [54] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10687–10698, 2020.
- [55] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1476–1485, 2019.