

Extract More from Less: Efficient Fine-Grained Visual Recognition in Low-Data Regimes

Dmitry Demidov

dmitry.demidov@mbzuai.ac.ae

Abduragim Shtanchaev

abduragim.shtanchaev@mbzuai.ac.ae

Mihail Mihaylov

mihail.mihaylov@mbzuai.ac.ae

Mohammad Almansoori

mohammad.almansoori@mbzuai.ac.ae

Mohamed bin Zayed University
of Artificial Intelligence

United Arab Emirates, Abu Dhabi

Abstract

The emerging task of fine-grained image classification in low-data regimes assumes the presence of low inter-class variance and large intra-class variation along with a highly limited amount of training samples per class. However, traditional ways of separately dealing with fine-grained categorisation and extremely scarce data may be inefficient under both these harsh conditions presented together. In this paper, we present a novel framework, called AD-Net, aiming to enhance deep neural network performance on this challenge by leveraging the power of Augmentation and Distillation techniques. Specifically, our approach is designed to refine learned features through self-distillation on augmented samples, mitigating harmful overfitting. We conduct comprehensive experiments on popular fine-grained image classification benchmarks where our AD-Net demonstrates consistent improvement over traditional fine-tuning and state-of-the-art low-data techniques. Remarkably, with the smallest data available, our framework shows an outstanding relative accuracy increase of up to 45 % compared to standard ResNet-50 and up to 27 % compared to the closest SOTA runner-up. We emphasise that our approach is practically architecture-independent and adds zero extra cost at inference time. Additionally, we provide an extensive study on the impact of every framework's component, highlighting the importance of each in achieving optimal performance. Source code and trained models are publicly available at github.com/demidov98/fgic_lowd.

1 Introduction

Overview. Deep learning models, inherently data-hungry, require vast quantities of annotated data for effective training [15]. However acquiring such extensive datasets, particularly with meticulous annotations, is labour-intensive and challenging. Following this constraint, the task of low-data image classification assumes that only a highly limited amount of annotated samples per class is available with no other information. [58, 61, 62]. Despite the




Figure 1: An overview of our proposed pre-processing pipeline. Three random crops of different sizes are generated from each input image and all of them are further augmented with the same set of transformations. The largest cropped region is further used for classification and the other two (mid- and small-size) are used as target and source in a distillation objective. Only random crop is shown as an applied augmentation for simplicity.

advantages of the existing methodologies, they often come with computationally expensive and architecture-dependent extra modules [69] and in some cases still demand substantial minimally required data quantities [48] or expect the input data to be in a specific domain only [13]. A potentially helpful traditional way to address this bottleneck is to artificially increase data amounts by applying various data augmentations [9, 40, 62]. However, our investigation demonstrates that, when applied within the standard transfer learning pipeline [2], these approaches are not practically efficient under the extremely low amount of data.

Another challenging domain within computer vision is fine-grained visual recognition, where the goal is to learn distinctive features for visually or semantically similar sub-classes of the same meta-category [12, 68]. This task may be difficult even for human annotators due to specific domain knowledge requirements [61, 62], making large-scale data acquisition and labelling impractical. Therefore, the problem of low-data learning is especially important for fine-grained classification tasks where the available data per class is often scarce [61, 62].

Problem Statement. Our work focuses on fine-grained image classification (FGIC) under low-data regimes, where both intra-class variation is high and the data is scarce, so conventional solutions may be limited [40]. This challenge is complicated by the intricate nature of the FGIC task and limited research in transfer learning for low-data problems. Most of the few existing solutions [52, 59] aim to directly introduce learning constraints without addressing the data variation insufficiency and often at the cost of noticeable time and resource increase. Therefore, motivated by the under-explored low-data learning domain, we provide a solution that effectively solves fine-grained classification problems by utilising a soft and indirectly regularised transfer learning technique in order to mitigate overfitting challenges.

Approach. We propose a novel approach which performs two crucial tasks: leverages image augmentations to enrich the feature space and utilises self-distillation to facilitate knowledge refinement among feature layers. Unlike traditional methods, our single-model approach integrates the classification and teacher-student distillation pipelines in one framework. Specifically, in our solution, we leverage a self-distillation objective to maintain close feature distributions between teacher and student outputs for different views of each image, which helps to enhance learned representations. Our detailed study offers the following *contributions*:

- We propose an end-to-end low-data framework which includes two feature distillation branches for model purification by soft and implicit regularisation.
- Our architecture-independent solution allows a model to learn and refine representations by enriching the variability of data through extra augmented image views.
- We demonstrate that our approach achieves state-of-the-art results on popular FGIC ben-

chmarks under extremely low data regimes and helps to reduce overfitting.

- With our framework, the performance gain comes at zero extra cost at inference time.

In Section 4.3, we present an extensive study of the architecture design, objective functions, and augmentation types, contributing to the evolving landscape of both fine-grained image classification and low-data regimes learning.

2 Related Work

Fine-grained Image Classification. There exist multiple approaches for the fine-grained image classification setting. Traditionally, pre-trained R-CNNs [56] and part detectors [40] had been utilised to detect similarities between specific parts of the object in the image. While modern approaches tend to utilise end-to-end architectures [59] where a mutual feature vector is obtained from multiple backbones and further used to compare to unique image representations such that the model can distinguish between difficult classes [48, 51, 58]. Another way to solve the FGIC problem is to focus on robust loss functions [6], for example having both a discriminability component which forces all feature channels belonging to the same class to be discriminative, using a channel-wise attention mechanism, and a diversity component which promotes mutually exclusive channels. Similar methods also focus on better feature extraction, like aggregating the important tokens from each transformer layer to compensate for local, low-level and middle-level information [48]. Some of the recent approaches [8, 50] consider utilising the heavy Vision Transformer (ViT) [45] architecture and semi-supervised settings, however such backbones require a computationally expensive pre-training phase and are noticeably slow at inference time.

Low-Data Learning. The specific task of fine-grained image classification in low-data regimes [23] is not widely explored, and, to our knowledge, there exist only a few approaches that attempt to solve it. One of them is using a self-boosting attention mechanism in CNNs to focus on more descriptive parts of the classes [59]. However, this FGIC method is designed in an architecture-specific way and can not be extended to other backbones. There exist other approaches that attempt to improve the training data quality by utilising saliency maps in CNNs to specifically target low-data regime settings. This can be seen in [46] where the authors incorporate a specific attention mechanism in the original image classification pipeline such that the model focuses on the important distinctive parts of the image when classifying similar classes. In [44] the authors use a multi-scale feature pyramid and a multi-level attention pyramid on a backbone network to progressively aggregate features from different granular spaces. They further present an attention-guided refinement strategy in collaboration with a multi-level attention pyramid to reduce the uncertainty brought by backgrounds conditioned by limited samples. Nevertheless, the existing methods are often sub-optimal or architecture-specific and resource-demanding, leaving room for a more elegant solution.

Augmentations. Recent developments in augmentation strategies for self-supervised models include methods like MultiCrop [2] which employ a mix of views with different resolutions, enhancing encoding space variability without increasing memory or compute requirements. Another powerful augmentation is ScaleMix [50] which generates new views of an image by mixing two views of potentially different scales using binary masking. Similar to ScaleMix, CutMix [54] operates exclusively on views from the same image, producing a single view of standard size, however, the latter also approximates benefits of multiple crops without requiring separate processing of small crops. Both MultiCrop and ScaleMix introduce extra variance to the encoding space, enriching the representations.

Distillation. Another recent approach utilised for efficient learning is model distillation during training. Self-supervised learning methodologies have increasingly embraced distillation objectives as a central principle for robust representation learning in the absence of explicit labels [8, 26, 65]. Empirical studies in the literature, including works such as [4, 19, 65], demonstrate reasonable performance, particularly in the challenging paradigm of zero-shot learning [44]. Drawing motivation from this paradigm, we attempt to integrate the distillation techniques into our approach to target the challenging scenarios of limited data availability.

3 Method

3.1 Baseline Framework

Transfer Learning. In traditional transfer learning, a backbone is usually pre-trained on a huge and visually diverse dataset (such as ImageNet [14]) and is further fine-tuned on a smaller fine-grained downstream task in order to transfer generic low- and mid-level features learned on the large and diverse dataset. This type of initialisation is especially beneficial in the case of using smaller fine-grained datasets [11, 60]. Adopting standard terminology in transfer learning, we define the source domain D_S with data X_S and the target domain D_T with data X_T . Here, $X = \{x_1, \dots, x_n\}$ represents the feature set within domain-specific feature spaces \mathcal{X}_S and \mathcal{X}_T , respectively. The objective in this context is to predict the label y' for a new instance x' in the target domain, utilising a label space Y and learning a prediction function $f: \mathcal{X}_T \rightarrow Y$. This function is derived from training pairs (x_i, y_i) , where $x_i \in X_S \cup X_T$ and $y_i \in Y$, effectively mapping the learned features to the desired outputs.

Limitations. Although the traditional fine-tuning strategy is useful for standard datasets with relative abundance of labelled samples, it may be sub-optimal on tasks with highly limited data. In low-data regimes, this practically leads to dramatic overfitting to the few given samples per class. Moreover, previous works on low-data learning [69, 48] have demonstrated that the unconstrained baselines mostly focus on the most obvious but less distinctive visual features. Taking this into account, we address these challenges of the low-data setting with our proposed approach, which is able to reduce harmful overfitting by refining learned representations through self-distillation on multiple augmented views of the input image.

3.2 Our Approach

Overall Architecture. In this paper, we propose a framework, called AD-Net, which aims to achieve the enhancement of deep neural network performance through the combination of Augmentation and Distillation techniques. Such a combination leads to substantial quality improvement of learned information per image, which allows for the neural network to learn better representations with fewer images. Our primary focus is on the application of self-distillation techniques, incorporating an additional distillation loss applied to augmented input samples to facilitate knowledge enhancement within the network. More specifically, we explore the utilisation of various data augmentations to further improve the model’s robustness and experiment with different distillation loss functions and components, providing a comprehensive analysis of their effects on the model’s performance. Taking this into account, our proposed framework incorporates a multi-branch configuration, functioning in a manner akin to Siamese networks [26, 65]. This design choice allows for weight sharing between branches, enhancing computational efficiency and facilitating the consolidation of

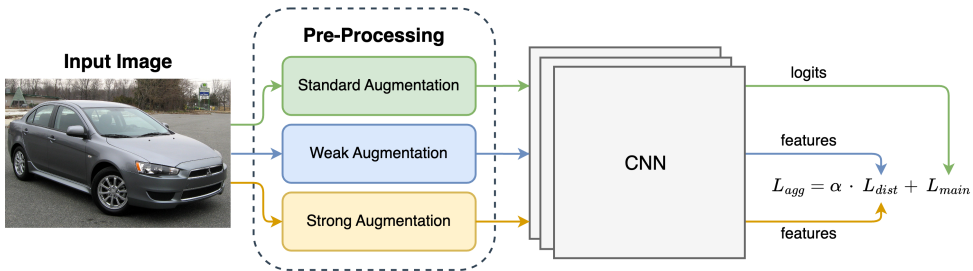


Figure 2: A proposed architecture. A random crop is applied as an input augmentation for images which are subsequently fed into the same model with shared weights. Further, we compare the feature outputs of different crops with a distillation loss \mathcal{L}_{dist} . The aggregated loss \mathcal{L}_{agg} is a combination of the traditional cross-entropy loss \mathcal{L}_{main} and \mathcal{L}_{dist} .

learned features. Despite the shared weights, each branch is uniquely controlled by distinct loss objectives, emphasising their individual contributions to the overall learning process. Specifically, the branches can be delineated as the classification branch and the distillation branches for simplicity. Figure 2 provides an overview of our proposed framework.

Classification Branch. This branch replicates the standard supervised training process, which follows the conventional fine-tuning procedure [57] with an ordinary classification objective. Full-sized images with traditional data augmentation techniques applied are fed into this branch in order to separately obtain the class prediction distribution.

Distillation Branches. This part of the architecture performs knowledge distillation with two differently augmented views of the input image. More specifically, one “source” branch and one “target” branch are utilised to generate corresponding source and target representations of the original images, which are further used for self-distillation [9, 45]. This encourages the model to do feature refinement towards more useful class-specific patterns. In more detail, both branches utilise the same shared-weight backbone from the main classification branch, but it is now fed with random mid-sized and small crops of the original image. Specifically, following [50], we randomly sample two categories of patches, larger target crops ranging between 65-85% of the original image content and smaller source crops with an area range of 1-65% in order to capture fine-grained class-specific patterns [4]. This kind of image space limitation can help to increase data variation and introduce implicit feature pruning via information dropout [44]. Next, all sampled regions are resized to the fixed resolution and augmented with the same set of data augmentations used for the classification branch. Thus, for each input image, one cropped region from both categories is fed into the corresponding branches of the model to obtain two different representations. Finally, inspired by [45], a self-distillation technique is applied to the corresponding feature outputs of previously sampled image crops, which encourages the model to learn better class-specific patterns. This helps to overcome quick overfitting towards explicit but weak features.

Loss Formulation. In our approach, we deploy a combination of two distinct loss functions to optimise learning efficiency. The primary classification objective, denoted as \mathcal{L}_{main} , explicitly facilitates the probability distribution prediction towards a pre-defined class. Concurrently, secondary distillation objective \mathcal{L}_{dist} serves as an implicit regularisation mechanism [50]. More specifically, the classification loss is applied to the output logits of the classification branch, while the distillation loss is applied to the feature maps derived from the target and source distillation branches (refer to Section 3.2). For our main classification objective,

we employ traditionally used Cross-Entropy loss [25, 84, 57]:

$$\mathcal{L}_{main} = -\frac{1}{N} \sum_{n=1}^N \sum_i^C y_{ni} \log(\hat{y}_{ni}) \quad (1)$$

where y_{ni} is the ground truth label and \hat{y}_{ni} is the prediction of the model.

Meanwhile, for the distillation objective, we utilise Kullback–Leibler divergence [27]:

$$\mathcal{L}_{dist} = \frac{1}{N} \sum_{n=1}^N \sum_i^C P'_{ni} \log \left(\frac{P'_{ni}}{Q_{ni}^s} \right) \quad (2)$$

where P'_{ni} and Q_{ni}^s are feature distributions from target and source distillation branches.

In our experiments we apply the softmax function on feature outputs to convert them into normalised probability distributions. The main intuition behind this decision is to enforce the model to produce similar feature distributions for different regions of the same image.

Finally, following a widely accepted way of objectives combination [13, 29, 39], our aggregated loss function \mathcal{L}_{agg} , consisting of the classification and distillation components, is:

$$\mathcal{L}_{agg} = \mathcal{L}_{main} + \alpha \cdot \mathcal{L}_{dist} \quad (3)$$

where α is a hyper-parameter determining the weight of the distillation loss in the final objective (for more details and ablation analysis refer to Section 4.3).

4 Experiments

4.1 Experimental Details

Datasets. We explore the properties of our approach on the following popular FGIC datasets: CUB-200-2011 [47], Stanford Cars [28], FGVC-Aircraft [33]. The following datasets have been chosen due to their balanced class sizes and a similar number of images per class, which allow for a more fair evaluation (for details refer to Appendix B.1). In order to imitate low-data regimes, we strictly follow the sampling proposed in [39], where samples are randomly drawn from the training set providing 10%, 15%, 30%, and 50% percentages of data.

Baselines. In order to provide a fair comparison, in our experiments we utilise methods based on popular CNN and ViT architectures, where for all models (including ours) we perform transfer learning by fine-tuning. For the vanilla baselines we also consider two setups: traditional fine-tuning and advanced fine-tuning with our proposed training recipe designed for low-data learning. For the comparison, we first consider methods specifically designed for the fine-grained classification task, such as: Full Bilinear Pooling (FBP) [31], Compact Bilinear Pooling with Tensor Sketch projection (CBP-TS) [18], Hierarchical Bilinear Pooling (HBP) [53], Deep Bilinear Transformation (DBTNet-50) [58], and FFVT [48]. As a main competitor, we evaluate the first and currently the only specifically designed approach for low-data in FGIC [39], which is a CNN-based model with a self-boosting attention mechanism, by preserving the authors’ codebase and training hyper-parameters. For our approach, we similarly use a family of ResNet models [21] as a main backbone along with such other popular baselines as: GoogLeNet [42], DenseNet [24], Inception v3 [43], and ViT [15]. All considered models were standardly pre-trained on ImageNet and further fine-tuned on the above-mentioned fine-grained datasets for both vanilla baselines and our AD-Net.

Implementation Details. Compared to the traditional fine-tuning procedure, initially designed for a sufficient amount of samples per class, our proposed training recipe includes several adaptations for the low-data regimes (for application guidelines refer to Appendix D.1).

Table 1: Comparison of different approaches using various percentages of the data on three popular FGIC datasets. Our proposed solution achieves consistent improvement in performance over other methods across low data settings. Best results are highlighted in bold.

Dataset	Method	Training data percentage			
		10%	15%	30%	50%
CUB-200-2011 [47]	ResNet-50	36.99	48.88	62.60	73.23
	FBP	37.88	49.12	63.27	73.70
	CBP-TS	37.12	47.82	62.24	72.37
	HBP	38.57	50.12	63.86	74.18
	DBTNet-50	37.67	49.52	63.16	73.28
	SAM (ResNet-50)	40.24	52.05	64.07	73.92
	SAM (FBP)	41.83	52.35	65.19	74.54
	<i>Ours (ResNet-50)</i>	47.51	60.08	71.11	77.67
Stanford Cars [28]	ResNet-50	37.45	53.01	75.26	83.56
	FBP	40.13	55.07	76.42	85.10
	CBP-TS	37.77	54.87	75.51	84.80
	HBP	40.02	55.82	76.81	85.31
	DBTNet-50	39.48	55.24	76.52	86.52
	SAM (ResNet-50)	39.96	55.02	76.69	84.85
	SAM (FBP)	43.19	57.42	77.63	85.71
	<i>Ours (ResNet-50)</i>	55.09	67.42	81.53	87.41
FGVC-Aircraft [33]	ResNet-50	43.52	53.17	71.32	78.61
	FBP	45.16	55.06	72.12	79.93
	CBP-TS	44.63	54.79	71.32	79.60
	HBP	45.28	56.12	72.58	81.47
	DBTNet-50	45.35	56.36	73.06	81.26
	SAM (ResNet-50)	46.73	56.02	72.59	79.21
	SAM (FBP)	47.97	57.47	73.43	80.86
	<i>Ours (ResNet-50)</i>	55.81	62.59	74.44	81.73

First, following a common idea of increasing the main learning rate value lr by a fixed ratio for the last classification layer lr_{cls} , we automatically update this ratio depending on the percentage of the available data \mathcal{D}_{aval} with the following rule:

$$lr_{cls} = lr \cdot \left(\text{round}\left(\frac{100}{\mathcal{D}_{aval}} \cdot 2\right) / 2 \right). \quad (4)$$

This assumption is valid in the simulated setup due to similar dataset sizes, but it can be further extended by replacing the available data percentage with the number of images per class adjusted to the number of classes. Additionally, we utilise a learning rate scheduler with downscaling lr for 5 iterations $iter \in [1 : 5]$ by 0.1 at each $(steps_{max} \cdot 0.5^{iter})$ training step. For other implementation details and training hyper-parameters refer to Appendix B.2.

4.2 Results and Analysis

Quantitative Analysis. In order to thoroughly investigate our proposed method, we conduct experiments on different label proportions of the popular FGIC benchmarks and further

compare the results with the state-of-the-art approaches. Results in Table 1 demonstrate that our self-distillation component significantly improves the performance of the vanilla ResNet-50 across all of the sampled percentages on all considered datasets, while also outperforming other SOTA methods in the low-data settings. Specifically, with the smallest data available, our framework shows an outstanding relative accuracy increase of up to 45 % compared to standard ResNet-50 and up to 27 %, compared to the closest SOTA runner-up. Our solution also outperforms the specifically designed for low-data fine-grained classification SAM framework [39], while not introducing any extra resource-consuming and architecture-specific module. It can be observed that the most significant improvement is achieved with the smallest number of samples per class available, while the relative performance increase goes down with more data used for training. This could be explained by the overall data saturation on larger data percentages available, which is the same behaviour for every considered method. Nevertheless, the benchmark analysis proves the ability of our solution to adapt to the setups with limited data by performing feature refinement through self-distillation on the existing samples. For a detailed comparison including other low-performing baselines refer to Appendix C.4.1. Additionally, we demonstrate the high transferability of our approach by utilising it on top of most of the popular CNN- and ViT-based backbones (see Appendix C.4.2 for details). The absolute improvement varies between 3-10 % showcasing that distilling local augmentations is practically an architecture-independent technique.

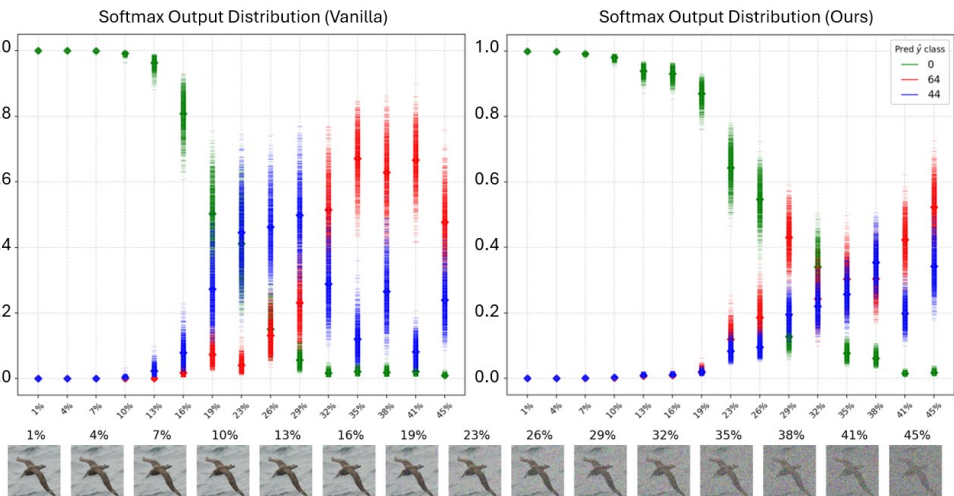


Figure 3: Model prediction probability distribution over 1000 forward passes with Monte Carlo Dropout. X-axis stands for the amount of Gaussian noise added (see bottom row) and Y-axis is the probability of a predicted class. Each dash (—, —, —) represents a single model’s prediction, where green is the correct class and others are the Top-2 following classes. Each diamond (◆, ◆, ◆) stands for the mean probability over all predictions.

Qualitative Analysis. In order to analyse the motivation behind the significant performance improvement with our approach, we examine the model’s performance under different levels of noise added to the input images. Specifically, we evaluate the model’s prediction uncertainty by using Monte Carlo Dropout [17]. We first sample a random image from the validation set and create a set of 1000 perturbed copies of the image by adding Gaussian noise with varying intensity. We then perform a forward pass using vanilla ResNet and our

AD-Net for each image in the set (comparison illustrated in Figure 3). We can observe that while the vanilla’s predictions have a high false positive rate on noisy samples, the predictions of our AD-Net have lower dispersion and are more stable due to more robust purified features. For more qualitative analysis and visualisations refer to Appendix C.1 and C.6.

4.3 Ablation Study

Architecture Design. In order to justify the chosen multi-branch pipeline, we provide an architecture design ablation for our framework (see Table 2). Specifically, we first measure the performance of the standard classification branch with the naive fine-tuning procedure. Next, we check the performance of this model with our advanced low-data training recipe, which includes an adaptive learning rate ratio for the final classification layer and a scheduler. Further, we add a single distillation branch which is fed with a small cropped region and then feature distillation is performed between the classification and distillation feature outputs. Finally, we show the performance of our best solution with one target and one source distillation branches which are fed with corresponding larger target crop and smaller source crop in order to perform a separate self-distillation procedure.

Table 2: Ablation results on our framework architecture design and training recipe. The experiments are performed on ResNet-50 and the CUB dataset with 10% of the data.

Distillation	Our recipe	Augmentation	Acc, %
-	×	Basic	36.99
	✓	Basic	40.05
Single-branch	✓	ScaleMix	45.98
	✓	MultiCrop	46.22
	✓	Basic	47.09
Double-branch (Ours)	✓	ScaleMix	46.11
	✓	MultiCrop	45.91
	✓	AsymAug	47.18
	✓	Basic	47.51

Interestingly, application of advanced augmentation techniques, such as ScaleMix [60], MultiCrop [4], and AsymAug [50], within our self-distillation branches demonstrated some gains, but is not superior. We assume that this could be explained by the strong nature of these augmentations, which may not be as efficient due to the harsh perturbations. For more ablation analysis and limitations refer to Appendix C.2, C.3, and C.5.

5 Conclusion

Our proposed AD-Net framework demonstrates exceptional results on the fine-grained image classification task under low-data regimes where both intra-class variation is high and the data is scarce, which typically requires the identification of local and subtle class-specific features. It is achieved by providing a solution that effectively solves fine-grained classification problems by utilising a soft and indirectly regularised transfer learning technique in order to mitigate overfitting challenges, while also increasing data variation and introducing implicit feature pruning via information dropout.

Specifically, our core architectural modification is the integration of advanced augmentation techniques and a distinctive self-distillation strategy, wherein a single model processes both the original image and its transformed views. Such a multi-input approach enhances the model’s ability to capture diverse fine-grained patterns, which further allows it to recognise and differentiate intricate features unique to each class.

Through comprehensive experiments, we have established that our model surpasses existing state-of-the-art techniques by a significant margin, particularly excelling in scenarios with the most limited data (see Appendix C.6). Specifically, with the smallest data available, our AD-Net shows an outstanding relative accuracy increase of up to 45 % compared to standard ResNet-50 and up to 27 % compared to the closest SOTA runner-up. This success is attributed to the synergistic combination of augmentations and a tailored objective function, which collectively improve the model’s learning quality. In detail, we explain this noticeable performance gain by the effect of the self-distillation objective, which provides a separate and more detailed source of information by enforcing feature space alignment for the augmented views of the same input image. Such a setup allows our AD-Net to avoid quick overfitting typical for the standard fine-tuning procedure (refer to Appendix C.1).

We also emphasise that our proposed framework is practically architecture-independent since it operates on the final feature representation output from the utilised baseline. Along with zero extra costs at inference time, the above-mentioned benefits make our AD-Net framework highly suitable for low-data fine-grained image classification problems.

References

- [1] Steve Branson, Grant Van Horn, Serge Belongie, and Pietro Perona. Bird species categorization using pose normalized deep convolutional nets. *arXiv preprint arXiv:1406.2952*, 2014.
- [2] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33:9912–9924, 2020.
- [3] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS ’20*, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers, 2021.
- [5] Dongliang Chang, Yifeng Ding, Jiyang Xie, Ayan Kumar Bhunia, Xiaoxu Li, Zhanyu Ma, Ming Wu, Jun Guo, and Yi-Zhe Song. The devil is in the channels: Mutual-channel loss for fine-grained image classification. *IEEE Transactions on Image Processing*, 29: 4683–4695, 2020.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML’20*. JMLR.org, 2020.

- [7] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners, 2020.
- [8] Po-Yung Chou, Cheng-Hung Lin, and Wen-Chung Kao. A novel plug-in module for fine-grained visual classification. *arXiv preprint arXiv:2202.03822*, 2022.
- [9] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018.
- [10] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020.
- [11] Mahdi Darvish, Mahsa Pouramini, and Hamid Bahador. Towards fine-grained image classification with generative adversarial networks and facial landmark detection. In *2022 International Conference on Machine Vision and Image Processing (MVIP)*, pages 1–6. IEEE, 2022.
- [12] Dmitry Demidov., Muhammad Sharif., Aliakbar Abdurahimov., Hisham Cholakkal., and Fahad Khan. Salient mask-guided vision transformer for fine-grained classification. In *Proceedings of the 18th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2023) - Volume 4: VISAPP*, pages 27–38. INSTICC, SciTePress, 2023. ISBN 978-989-758-634-7. doi: 10.5220/0011611100003417.
- [13] Dmitry Demidov, Roba Al Majzoub, Amandeep Kumar, and Fahad Khan. Distilling local texture features for colorectal tissue classification in low data regimes. In Xiaohuan Cao, Xuanang Xu, Islem Rekik, Zhiming Cui, and Xi Ouyang, editors, *Machine Learning in Medical Imaging*, pages 357–366, Cham, 2024. Springer Nature Switzerland. ISBN 978-3-031-45676-3.
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [16] Carola Figueroa Flores, Abel Gonzalez-Garcia, Joost van de Weijer, and Bogdan Raducanu. Saliency for fine-grained object recognition in domains with scarce training data. *Pattern Recognition*, 94:62–73, 2019.
- [17] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning, 2016.
- [18] Yang Gao, Oscar Beijbom, Ning Zhang, and Trevor Darrell. Compact bilinear pooling. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 317–326, 2016. doi: 10.1109/CVPR.2016.41.

- [19] 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, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning, 2020.
- [20] Ju He, Jieneng Chen, Shuai Liu, Adam Kortylewski, Cheng Yang, Yutong Bai, Changhu Wang, and Alan Loddon Yuille. Transfg: A transformer architecture for fine-grained recognition. In *AAAI Conference on Artificial Intelligence*, 2021. URL <https://api.semanticscholar.org/CorpusID:232233178>.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [22] Elad Hoffer, Tal Ben-Nun, Itay Hubara, Niv Giladi, Torsten Hoefer, and Daniel Soudry. Augment your batch: better training with larger batches. *arXiv preprint arXiv:1901.09335*, 2019.
- [23] Grant van Horn and Pietro Perona. The devil is in the tails: Fine-grained classification in the wild. *arXiv preprint arXiv:1709.01450*, 2, 2017.
- [24] Gao Huang, Zhuang Liu, and Kilian Q. Weinberger. Densely connected convolutional networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269, 2016. URL <https://api.semanticscholar.org/CorpusID:9433631>.
- [25] Like Hui and Mikhail Belkin. Evaluation of neural architectures trained with square loss vs cross-entropy in classification tasks, 2021.
- [26] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning, 2021.
- [27] Taehyeon Kim, Jaehoon Oh, NakYil Kim, Sangwook Cho, and Se-Young Yun. Comparing kullback-leibler divergence and mean squared error loss in knowledge distillation, 2021.
- [28] Jonathan Krause, Jia Deng, Michael Stark, and Li Fei-Fei. Collecting a large-scale dataset of fine-grained cars. 2013. URL <https://api.semanticscholar.org/CorpusID:16632981>.
- [29] Emirhan Kurtulus, Zichao Li, Yann Dauphin, and Ekin D. Cubuk. Tied-augment: controlling representation similarity improves data augmentation. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org, 2023.
- [30] Manuel Lagunas, Brayan Impata, Victor Martinez, Virginia Fernandez, Christos Georgakis, Sofia Braun, and Felipe Bertrand. Transfer learning for fine-grained classification using semi-supervised learning and visual transformers. *arXiv preprint arXiv:2305.10018*, 2023.
- [31] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnn models for fine-grained visual recognition. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1449–1457, 2015. doi: 10.1109/ICCV.2015.170.

- [32] Ziqian Lu, Zheming Lu, Yunlong Yu, and Zonghui Wang. Learn more from less: Generalized zero-shot learning with severely limited labeled data. *Neurocomputing*, 477:25–35, 2022. ISSN 0925-2312. doi: <https://doi.org/10.1016/j.neucom.2022.01.007>. URL <https://www.sciencedirect.com/science/article/pii/S0925231222000078>.
- [33] S. Maji, J. Kannala, E. Rahtu, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. Technical report, 2013.
- [34] Anqi Mao, Mehryar Mohri, and Yutao Zhong. Cross-entropy loss functions: Theoretical analysis and applications, 2023.
- [35] Pratik Mazumder, Pravendra Singh, and Vinay P. Namboodiri. Fair visual recognition in limited data regime using self-supervision and self-distillation, 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. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [37] Ade Ramdan, Ana Heryana, Andria Arisal, R. Budiarianto S. Kusumo, and Hilman F. Pardede. Transfer learning and fine-tuning for deep learning-based tea diseases detection on small datasets. In *2020 International Conference on Radar, Antenna, Microwave, Electronics, and Telecommunications (ICRAMET)*, pages 206–211, 2020.
- [38] Lars Schmarje, Monty Santarossa, Simon-Martin Schröder, and Reinhard Koch. A survey on semi-, self-and unsupervised learning for image classification. *IEEE Access*, 9:82146–82168, 2021.
- [39] Yangyang Shu, Baosheng Yu, Haiming Xu, and Lingqiao Liu. Improving fine-grained visual recognition in low data regimes via self-boosting attention mechanism. In *European Conference on Computer Vision*, pages 449–465. Springer, 2022.
- [40] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [41] Maryam Sultana, Muzammal Naseer, Muhammad Haris Khan, Salman Khan, and Fahad Shahbaz Khan. Self-distilled vision transformer for domain generalization, 2022.
- [42] C. Szegedy, Wei Liu, Yangqing Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, Los Alamitos, CA, USA, jun 2015. IEEE Computer Society. doi: 10.1109/CVPR.2015.7298594. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2015.7298594>.
- [43] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016. doi: 10.1109/CVPR.2016.308.

- [44] Hao Tang, Chengcheng Yuan, Zechao Li, and Jinhui Tang. Learning attention-guided pyramidal features for few-shot fine-grained recognition. *Pattern Recognition*, 130: 108792, 2022.
- [45] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10347–10357. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/touvron21a.html>.
- [46] Hugo Touvron, Matthieu Cord, and Hervé Jégou. Deit iii: Revenge of the vit, 2022.
- [47] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. Technical report. the caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- [48] Jun Wang, Xiaohan Yu, and Yongsheng Gao. Feature fusion vision transformer for fine-grained visual categorization. In *British Machine Vision Conference*, 2021. URL <https://api.semanticscholar.org/CorpusID:235742913>.
- [49] Xiao Wang and Guo-Jun Qi. Contrastive learning with stronger augmentations. *IEEE transactions on pattern analysis and machine intelligence*, 45(5):5549–5560, 2022.
- [50] Xiao Wang, Haoqi Fan, Yuandong Tian, Daisuke Kihara, and Xinlei Chen. On the importance of asymmetry for siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16570–16579, 2022.
- [51] Mengde Xu, Zheng Zhang, Han Hu, Jianfeng Wang, Lijuan Wang, Fangyun Wei, Xiang Bai, and Zicheng Liu. End-to-end semi-supervised object detection with soft teacher. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3060–3069, 2021.
- [52] Xiangli Yang, Zixing Song, Irwin King, and Zenglin Xu. A survey on deep semi-supervised learning. *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [53] Chaojian Yu, Xinyi Zhao, Qi Zheng, Peng Zhang, and Xinge You. Hierarchical bilinear pooling for fine-grained visual recognition. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 595–610, Cham, 2018. Springer International Publishing. ISBN 978-3-030-01270-0.
- [54] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019.
- [55] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction, 2021.

- [56] Ning Zhang, Jeff Donahue, Ross Girshick, and Trevor Darrell. Part-based r-cnns for fine-grained category detection. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I 13*, pages 834–849. Springer, 2014.
- [57] Zhilu Zhang and Mert R. Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels, 2018.
- [58] Heliang Zheng, Jianlong Fu, Zheng-Jun Zha, and Jiebo Luo. *Learning deep bilinear transformation for fine-grained image representation*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- [59] Peiqin Zhuang, Yali Wang, and Yu Qiao. Learning attentive pairwise interaction for fine-grained classification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 13130–13137, 2020.