

Robust and Explainable Fine-Grained Visual Classification with Transfer Learning: A Dual-Carriageway Framework

Zheming Zuo¹, Joseph Smith¹, Jonathan Stonehouse², Boguslaw Obara¹

¹School of Computing, Newcastle University, UK ²Procter and Gamble, UK

{zheming.zuo, j.smith57, boguslaw.obara}@newcastle.ac.uk stonehouse.jr@pg.com

Abstract

In the realm of practical fine-grained visual classification applications rooted in deep learning, a common scenario involves training a model using a pre-existing dataset. Subsequently, a new dataset becomes available, prompting the desire to make a pivotal decision for achieving enhanced and leveraged inference performance on both sides: Should one opt to train datasets from scratch or fine-tune the model trained on the initial dataset using the newly released dataset? The existing literature reveals a lack of methods to systematically determine the optimal training strategy, necessitating explainability. To this end, we present an automatic best-suit training solution searching framework, the Dual-Carriageway Framework (DCF), to fill this gap. DCF benefits from the design of a dual-direction search (starting from the pre-existing or the newly released dataset) where five different training settings are enforced. In addition, DCF is not only capable of figuring out the optimal training strategy with the capability of avoiding overfitting but also yields built-in quantitative and visual explanations derived from the actual input and weights of the trained model. We validated DCF's effectiveness through experiments with three convolutional neural networks (ResNet18, ResNet34 and Inception-v3) on two temporally continued commercial product datasets. Results showed fine-tuning pathways outperformed training-from-scratch ones by up to 2.13% and 1.23% on the pre-existing and new datasets, respectively, in terms of mean accuracy. Furthermore, DCF identified reflection padding as the superior padding method, enhancing testing accuracy by 3.72% on average. This framework stands out for its potential to guide the development of robust and explainable AI solutions in fine-grained visual classification tasks.

1. Introduction

Fine-Grained Visual Classification (FGVC) [36], as a specialised task within the realm of computer vision, focuses

on analysing the visual objects from subordinate categories with applications in biological research [35], automotive industry [10], fashion [64] and retail [40, 49], etc. The past two decades have witnessed a decent amount of FGVC research outcomes [4, 8, 9, 12, 31, 59] suppressing human experts since the broad application of Convolutional Neural Networks (CNNs) [29, 30].

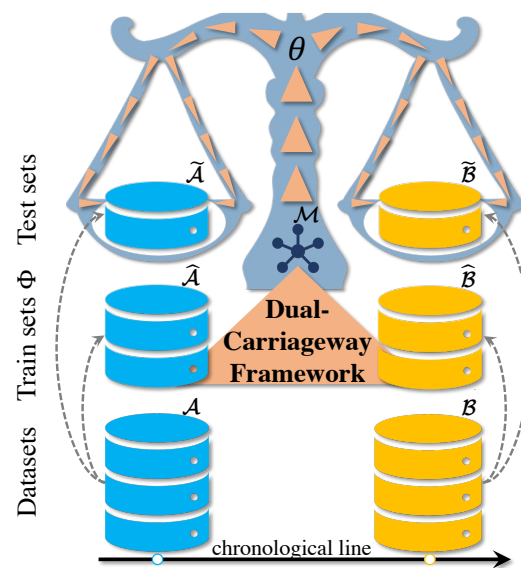


Figure 1. A robust and explainable learning framework should be revealed by not only the leveraged FGVC performance over two temporally continued datasets with the existence of subtle pattern differences, imbalanced data samples and high sparsity within the region of interest but also the associated quantitative (via the actual model input, *i.e.* frequency distribution of the padded image) and visual explanations (through the learned model weights, layer-wise attention) in a *putting-through* manner.

FGVC tends to be more challenging than conventional image classification due to the high degree of similar appearance among subordinate categories [37]. Besides, the coarse granularity of the captured image resulted from illumination variations (*e.g.* low-light [62] or over-exposed

[13] environment) and camera motion associated with unstable hand stability (e.g. blurriness [66] and distortion [28]) may deteriorate FGVC performance.

In a real-world scenario depicted in Fig. 1, suppose that \mathcal{A} and \mathcal{B} are the two available datasets; this paper aims to find an optimal solution to leverage the prediction performance of a model \mathcal{M} on their testing sets, i.e. $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$. Such an optimisation process can be expressed by:

$$\operatorname{argmax}_{\theta} \left(\operatorname{Acc}(\mathcal{M}_{\theta, \Phi}(\tilde{\mathcal{A}})), \operatorname{Acc}(\mathcal{M}_{\theta, \Phi}(\tilde{\mathcal{B}})) \right), \quad (1)$$

where the optimal hyperparameters θ were learned by model \mathcal{M} on the data Φ . The above optimisation problem is usually resolved by training two datasets from scratch in a row [60] or conducting transfer learning [41, 43, 65].

Fundamentally, the choice of training pathway highly relies on the data scale and pattern disparity between datasets \mathcal{A} and \mathcal{B} . However, those characteristics cannot be precisely quantified as the existence of challenging factors mentioned before the antecedent paragraph and various types of uncertainties (e.g. dataset noise and model randomness) [22] raised during the model training phase. The former would be less preferred if those two datasets had very different distributions. In contrast, the latter one suffers from the issue of catastrophic forgetting [3, 55] where the model forgets important features from dataset \mathcal{A} when fine-tuning on dataset \mathcal{B} . Besides, though transfer learning is a more common choice due to its much lower computational cost and possibly better generalisation capability contributed by e.g. the large-scale ImageNet dataset [7], it may distort the pre-trained features [26].

The challenges mentioned above in selecting the appropriate training pathway for leveraging the prediction performance on $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$ that are adjacent in the timeline as depicted in Fig. 1, this work addresses the problem from a different yet explainable perspective. Generally, we present an automatic framework for selecting the best model training scheme, named Dual-Carriageway Framework (DCF), for an FGVC task given two chronologically continued datasets, and our contributions are summarised as follows:

1) We propose an automatic best-suit training solution searching framework (Sec. 3) through five training settings (Sec. 3.3) configured in a dual-direction manner for robust fine-grained visual classification where the prediction performance on test sets of two temporally continued datasets are maximised.

2) We prove the feasibility and efficacy of the proposed framework (Sec. 4.3) through not only the prediction accuracy and degree of confidence yielded by the three trained CNN models but also quantitative and visual explanations in a *putting-through* manner from the actual model input data and layer-wise attention.

3) We show that determining the potentially most appropriate padding scheme (Sec. 3.2) is a reasonable start-

ing point (Sec. 4.2) for deploying the proposed framework, where merely the CNN model trained on the dataset collected earlier in the timeline is needed.

2. Materials

This section introduces the two datasets employed in this work, following the fine-grained classification task process sequence. Then, we combine several representative works to explain our choice and rationale for semantic segmentation models, padding schemes, and visual classifiers.

2.1. Datasets

To facilitate the conduct of this research, our industrial partners delivered two RGB image datasets, \mathcal{A} and \mathcal{B} , of the bottom side of commercial products, partially presented in Fig. 2 in chronological order. Each dataset contained two categories of instances indicating the corresponding manufacturers, labelled ‘F1’ and ‘F2’.



Figure 2. Sample images of the bottom side of commercial products produced by two manufacturers, ‘F1’ and ‘F2’, in pre-existing (\mathcal{A}) and newly released (\mathcal{B}) datasets. Target regions (dotted parts) to be segmented by U-Net [48] in Sec. 2.2) with ever-evolving anti-counterfeit code patterns embody spatial variations led by varying illumination, camera bias [20], and noisy textures [21] (e.g. dust and stains), affecting the overall FGVC accuracy.

Datasets \mathcal{A} and \mathcal{B} contain 5870 and 3112 RGB images, each of which comes with different resolutions, as the data-capturing process was performed by varying types of cameras in uncontrolled scenes. Slightly different from the commonly adopted data division approach, we use the data of the tail time block in both datasets as the test set, i.e. the test set of \mathcal{A} ($\tilde{\mathcal{A}}$) contains 1716 images in which $|\tilde{\mathcal{A}}_{F1}| = 1389$ and $|\tilde{\mathcal{A}}_{F2}| = 327$, and the test set of \mathcal{B} ($\tilde{\mathcal{B}}$) includes 822 images where $|\tilde{\mathcal{B}}_{F1}| = 220$ and $|\tilde{\mathcal{B}}_{F2}| = 602$.

The remaining part is divided into a training set and a validation set in a ratio of 4:1.

2.2. Semantic Segmentation Models

A substantial body of literature exists on the employment of deep neural networks for image semantic segmentation. Given that semantic segmentation is not the primary focus of this work, we briefly review key network architectures in the order of their introduction and describe our choice.

Fully Convolutional Network (FCN) [38], as an early seminal contribution to semantic segmentation, is a simple yet efficient architecture that preserves spatial information through end-to-end pixel-wise predictions due to the substitution of fully connected layers using the convolutional ones. As the evolutionary successor to FCN, U-Net [48] demonstrates a sophisticated capacity for capturing coarse and fine-grained features, achieving a favourable balance between temporal efficiency and richer feature representation. Its distinctive U-shaped architectural configuration integrates contracting and expansive pathways [14, 61], culminating in enhanced accuracy in spatial localisation.

DeepLab [6] utilises atrous convolutions and spatial pyramid pooling to capture multi-scale contextual information, thus enhancing the model’s ability to consider a broader context and overcome the restricted receptive field limitation associated with U-Net. High-Resolution Network (HRNet) [58] mitigates DeepLab’s limitations by devising high-resolution feature fusion, enhancing accuracy in fine details without compromising computational efficiency. Such a fusion addresses challenges related to longer training times and limited generalisation observed in DeepLab.

As the region of interest in this work is the anti-counterfeit code (with significant difference relative to background) in each of the given raw RGB images (depicted in the dotted part of Fig. 2), a U-Net [48], which is pre-trained on 1479 images with bounding boxes that annotated by the Amazon Mechanical Turk (MTurk) [1] crowdsourcing platform and refined by the Maximally Stable Global Region [21], is employed in this work for semantic segmentation (*i.e.* the input of the padding schemes partially introduced in Sec. 2.3 and systematically formulated in Sec. 3.2), as exemplified in the leftmost side of Fig. 3, to facilitate the nature of an FGVC task.

2.3. Padding Schemes and Classification Models

There is a limited amount of literature that proposes or systematically compares [19] padding schemes in FGVC, except for the most commonly adopted zero padding [45] in the past decade. Noteworthy, while many works focus on how to make neural networks deeper or wider, the padding scheme, which is usually neglected, is also closely related to the FGVC performance [39]. In CNNs [27], zero padding maintains the spatial size of the output feature maps by

adding layers of zeros around the input image before performing the convolution operation. However, it may introduce artefacts into the data that the model learns, leading to possible overfitting or reduced model generalisability [23]. The `reflection` padding addresses this by mirroring the image at the borders, maintaining the continuity of image features and reducing the edge artefacts [5, 11]. However, it may introduce redundancy and may not be suitable for all types of images. Similarly, padding schemes based on partial convolution [32, 33] and repetition [44] are robust to boundary artefacts [34].

On the contrary, the learning-based [2, 23], interpolation-based [16], and distribution-based [42] padding schemes focus on leveraging the connectivity of image borders for more effective padding to reflect the natural extension of the image content and suppress the risk of distorting the convolution. In addition, a padding scheme dynamically adjusted based on the contextual information around the border pixels is devised explicitly for semantic segmentation [18].

Furthermore, padding with a non-zero constant value is also a viable option, albeit infrequently implemented in practical applications [42]. Exploring whether segmented rectangular images and padded parts are complementary or mutually exclusive regarding model perceptibility would be interesting. To this end, considering that the segmented images are relatively sparse, and their width is much greater than their height, we systematically formulate six padding schemes, including zero padding and three schemes of padding with a non-zero constant in Eq. (2) as well as the `reflection` padding Eq. (4) in Sec. 3.2.

The padded images with the resolution of 1024-by-1024 pixels, as the actual input to CNNs, are then readily fed forward for convolution operations. To perform the binary classification task, we adopt three CNNs for binary classification, *i.e.* two residual networks (ResNet18 and ResNet34) [17] and one Inception network (Inception-v3) [57] due to their advantages of overcoming vanishing gradients [46] and capturing multi-scale features [63]. To take advantage of large-scale model pretraining, all three models initialise the training process based on the weights derived from ImageNet [7]. Detailed training pathways are elaborated in Sec. 3.3, and setups are configured in Sec. 4.1.

3. The Dual-Carriageway Framework

In this section, we present the proposed DCF for FGVC in detail. First, we briefly explain the DCF workflow and introduce its two key components.

3.1. Workflow

The workflow of the proposed DCF is depicted in Fig. 3. Briefly, it contains two consecutive components: a Padding

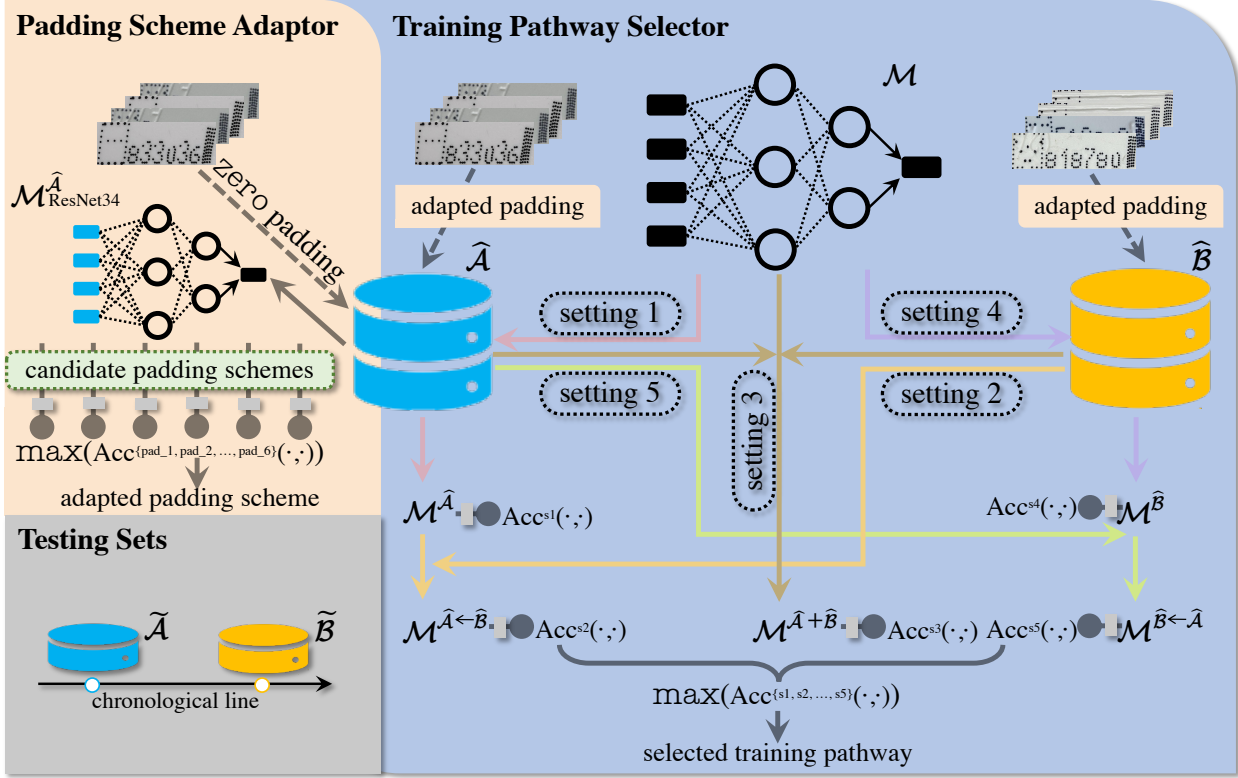


Figure 3. Workflow of the proposed DCF for robust and explainable fine-grained visual classification. This figure is separated into two consecutive components and a shared part by colours corresponding to 1) the Padding Scheme Adaptor (Sec. 3.2), 2) the Training Pathway Selector (Sec. 3.3), and 3) testing sets with pattern variations derived from two chronologically continued datasets (Sec. 2.1).

Scheme Adaptor (PSA) (detailed in the upper left side of Fig. 3) and a Training Pathway Selector (TPS) (depicted in the right-hand side of Fig. 3). Concretely, PSA is proposed to figure out the appropriate padding scheme to be adopted in the subsequent TPS where the FGVC accuracies on both temporally continued datasets (visualised in the lower left side of Fig. 3) are maximised and leveraged.

In PSA, we first train a ResNet34 [17] model (one of the three visual classification models adopted in this work and mentioned in the tail part of Sec. 2.3) on the earlier training set $\hat{\mathcal{A}}$ (splitter from its full dataset, \mathcal{A} detailed in Sec. 2.1) with the most commonly employed zero padding (revisited in the first half of Sec. 2.3). Such a trained model is tested on each of the six padding schemes (visualised in Fig. 4) that are systematically formulated in Sec. 3.2, using Eqs. (2) and (4) in general and Eqs. (3) and (5) in specific. Given that subtle changes in terms of the anti-counterfeit code pattern existed in datasets \mathcal{A} and \mathcal{B} , we can retain the most suitable padding scheme in our latter TPS phase by looking for the maximised inference performance on both testing sets. In TPS, for each of three visual classifiers (ResNet18, ResNet34 and Inception-v3 [57] adopted, we perform model training using five settings that

are introduced in Sec. 3.3 and summarised in Eq. (6) in a dual-carriageway manner, *i.e.* train on $\hat{\mathcal{A}}$ only, fine-tune the model learned from $\hat{\mathcal{A}}$ using $\hat{\mathcal{B}}$, train on both $\hat{\mathcal{A}}$ and $\hat{\mathcal{B}}$, train on $\hat{\mathcal{B}}$ only, and fine-tune the model learned from $\hat{\mathcal{B}}$ using $\hat{\mathcal{A}}$.

3.2. Padding Scheme Adaptor

As the output of a pre-trained semantic segmentation model [21, 48] mentioned in Sec. 2.2, each cropped image is \mathcal{I}_c in the shape of a rectangle whose width is much larger than its height. Given the characteristics of a range of padding schemes that have been systematically introduced in Sec. 2.3, we determine that a total of six padding schemes (in Fig. 4), *i.e.* three constant value-based ones (zero, white [45] and grey), two mean value-based ones (RGB-mean and LAB-mean), and one mirroring-based method (reflection) (a variant of [11] and [5]), are evenly performed in the vertical direction merely (except the reflection padding) with three practical advantages: a) preserving different aspect ratios; b) avoiding unnecessary distortions and c) reducing the computational cost.

Concretely, given a cropped image $\mathcal{I}_c \in \mathbb{R}^{H \times W \times C}$ where $H \ll W$, the zero, RGB-mean, LAB-mean,

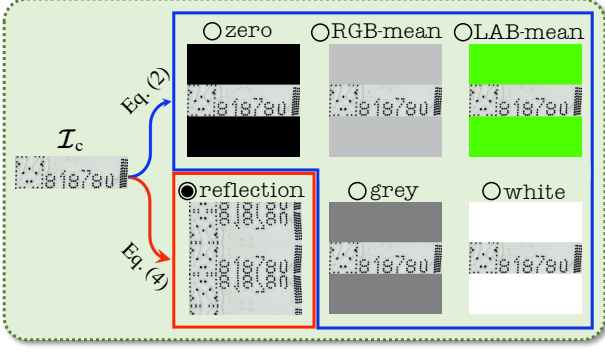


Figure 4. Candidate padding schemes available in the PSA component of the proposed DCF presented in Fig. 3.

white and grey padding schemes are performed evenly on both top and bottom side of \mathcal{I}_c , *i.e.* $P_t = P_b$, and thus the padded image $\mathcal{I}_p \in \mathbb{R}^{W \times W \times C}$ is defined by:

$$\mathcal{I}_p[i', j, :] = \begin{cases} \mathcal{I}_c[i - P_t, j, :] & \text{if } i \in [P_t, H + P_t), \\ (v_1, v_2, v_3) & \text{otherwise,} \end{cases} \quad (2)$$

where $i \in [0, H - 1]$, $j \in [0, W - 1]$, $i' \in [0, H + P_t - 1]$ s.t. $P_t = P_b = (W - H)/2$, and (v_1, v_2, v_3) is computed as:

$$\begin{cases} (0, 0, 0) & \text{zero,} \\ \frac{1}{HW} \sum_{i=0}^{H-1} \sum_{j=0}^{W-1} \mathcal{I}_c[i, j, :] & \text{RGB-mean,} \\ \frac{1}{HW} \sum_{i=0}^{H-1} \sum_{j=0}^{W-1} \mathcal{I}_{c \rightarrow \text{LAB}}[i, j, :] & \text{LAB-mean,} \\ (255, 255, 255) & \text{white,} \\ (128, 128, 128) & \text{grey.} \end{cases} \quad (3)$$

Particularly, the reflection padding scheme is expressed as:

$$\mathcal{I}_p[i', j, :] = \mathcal{I}_c[\text{reflect}(i, P_t, H), j, :], \quad (4)$$

in which the function $\text{reflect}(\cdot)$ is conducted when the index i meets in the following criteria in the vertical direction of \mathcal{I}_c :

$$\text{reflect}(i, P_t, H) = \begin{cases} i & \text{if } i \in [0, H), \\ 2P_t - i - 1 & \text{if } i \in [H, H + P_t]. \end{cases} \quad (5)$$

This ensures that reflection padding occurs when the index i is beyond the top border of \mathcal{I}_c . Additionally, if $H < P_t$, multiple times of reflections will occur, simulating the effect of extending the height of \mathcal{I}_c .

3.3. Training Pathway Selector

To resolve the optimisation problem defined in Eq. (1), given the training sets $\hat{\mathcal{A}}$ and $\hat{\mathcal{B}}$ split by the ratio defined in Sec. 2.1, optimal hyperparameters θ are commonly learned through the following dual-carriageway forms:

$$\theta := \begin{cases} \mathcal{M}^{\hat{\mathcal{A}}} & \text{setting 1,} \\ \mathcal{M}^{\hat{\mathcal{A}} \leftarrow \hat{\mathcal{B}}} & \text{setting 2,} \\ \mathcal{M}^{\hat{\mathcal{A}} + \hat{\mathcal{B}}} & \text{setting 3,} \\ \mathcal{M}^{\hat{\mathcal{B}}} & \text{setting 4,} \\ \mathcal{M}^{\hat{\mathcal{B}} \leftarrow \hat{\mathcal{A}}} & \text{setting 5,} \end{cases} \quad (6)$$

in which $\mathcal{M}^{\hat{\mathcal{A}} + \hat{\mathcal{B}}}$ denotes the model trained on $\hat{\mathcal{A}}$ and $\hat{\mathcal{B}}$ together, whereas $\mathcal{M}^{\hat{\mathcal{A}} \leftarrow \hat{\mathcal{B}}}$ represents the model trained on $\hat{\mathcal{A}}$ first and then fine-tuned using $\hat{\mathcal{B}}$, and vice versa. In addition, the alias of each training setting is used in the bottom right side of Fig. 3, *e.g.* s1 denotes setting 1.

Specifically, we adopt three relatively light CNN models in this work to justify the proposed framework of seeking an optimal training pathway for robust and explainable FGVC. Noteworthy, for training settings 2 and 4 of Eq. (6), we fine-tune each of the CNNs by layer freezing for simplicity and time-efficiency [50, 52, 53]. Detailed experimental setups are presented in Sec. 4.1.

4. Experiments

4.1. Experimental Setups and Evaluation Metrics

We implemented DCF on an NVIDIA RTX 3090 Ti GPU. To ensure uniform input image size of deep neural classifiers and compensate for various aspect ratios of the segmented images, we scale them to a resolution of 1024×1024 pixels with six padding schemes (detailed in Sec. 3.2) applied. For training, each of the three image classifiers (Sec. 2.2), we adopt Adam optimiser [24] and categorical cross-entropy loss [56] with a learning rate of 0.0001 and a weight decay of 0. Besides, we adopt random rotation for data augmentation to improve model generalisation.

The result was reported as mean accuracy \pm standard deviation by training the same deep image classifier five times for a fair comparison of FGVC performance. As a valuable add-on to the proposed DCF, we provide quantitative and visual explanations of the padded input (*i.e.* frequency distribution) and weights (*i.e.* layer-wise attention) associated with the trained classifier via mean pixel value and Gradient-weighted Class Activation Mapping (GradCAM) [51]. More experimental results are available in our *supplementary materials*.

4.2. Optimal Padding Scheme Determination

As briefly introduced (Sec. 2.3) and systematically designed (Sec. 3.2), we trained a ResNet34 model with the most com-

Table 1. Prediction performance (in %) comparisons on the testing sets of datasets \mathcal{A} (i.e. $\tilde{\mathcal{A}}$) and \mathcal{B} (i.e. $\tilde{\mathcal{B}}$) using the baseline models (ResNet34) trained on the training set of dataset \mathcal{A} (i.e. $\tilde{\mathcal{A}}$) that was processed by various padding schemes. Results yielded by the optimal padding scheme are marked in bold and highlighted with (detailed in Table 2).

Padding Scheme	$\text{Acc}(\mathcal{M}_{\text{ResNet34}}^{\tilde{\mathcal{A}}}(\tilde{\mathcal{A}}))$	$\text{Acc}(\mathcal{M}_{\text{ResNet34}}^{\tilde{\mathcal{A}}}(\tilde{\mathcal{B}}))$
zero	94.05 ± 2.03	71.79 ± 3.15
RGB-mean	95.07 ± 0.53	74.59 ± 0.86
LAB-mean	95.59 ± 0.48	72.61 ± 1.21
white	94.89 ± 0.20	75.78 ± 1.41
grey	95.10 ± 0.41	77.10 ± 1.06
reflection	96.58 ± 0.35	79.85 ± 1.18

Table 2. Detailed prediction performance (in %) of ResNet34 trained on $\tilde{\mathcal{A}}$ with reflection padding. The optimal model is marked in bold and highlighted by (adopted for tuning purposes in Sec. 4.3).

run	$\text{Acc}(\mathcal{M}_{\text{ResNet34}}^{\tilde{\mathcal{A}}}(\tilde{\mathcal{A}}))$	$\text{Acc}(\mathcal{M}_{\text{ResNet34}}^{\tilde{\mathcal{A}}}(\tilde{\mathcal{B}}))$
1	97.09	79.81
2	96.68	78.98
3	96.60	81.87
4	96.19	79.04
5	96.32	79.56

Table 3. Quantitative explanation of the relationship between input pixel value and confidence of correct predictions given by ResNet34 trained with zero padding under six padding schemes. $\text{avg}(\cdot; \cdot)$ calculates the average for each of the two elements regarding the correctly predicted samples of ‘F1’ and ‘F2’ within the testing set, \bar{p} (normalised to the range of [0, 1]) represents the mean pixel value, \mathcal{M}_c denotes the model confidence and \mathcal{M}_{acc} corresponds to the model inference accuracy measured in %. The best adapted $\text{avg}(\bar{p})$ values and their associated \mathcal{M}_c ones are marked in bold and highlighted with . Further visual explanation is delivered in Fig. 6.

Testing Set Label		zero	RGB-mean	LAB-mean	white	grey	reflection
		$\text{avg}(\bar{p}; \mathcal{M}_c)$ \mathcal{M}_{acc}	$\text{avg}(\bar{p}; \mathcal{M}_c)$ \mathcal{M}_{acc}	$\text{avg}(\bar{p}; \mathcal{M}_c)$ \mathcal{M}_{acc}	$\text{avg}(\bar{p}; \mathcal{M}_c)$ \mathcal{M}_{acc}	$\text{avg}(\bar{p}; \mathcal{M}_c)$ \mathcal{M}_{acc}	$\text{avg}(\bar{p}; \mathcal{M}_c)$ \mathcal{M}_{acc}
$\tilde{\mathcal{A}}$	F1	0.18; 0.989 95.22	0.60; 0.979 95.37	0.39; 0.986 94.98	0.88; 0.990 95.07	0.53; 0.976 94.91	0.60; 0.990 96.60
	F2	0.18; 0.976 95.22	0.60; 0.967 95.37	0.33; 0.984 94.98	0.88; 0.973 95.07	0.53; 0.972 94.91	0.60; 0.983 96.60
$\tilde{\mathcal{B}}$	F1	0.20; 0.958 73.68	0.70; 0.941 75.66	0.58; 0.951 74.66	0.91; 0.961 77.37	0.56; 0.959 78.22	0.70; 0.961 81.87
	F2	0.23; 0.909 73.68	0.71; 0.917 75.66	0.60; 0.932 74.66	0.90; 0.932 77.37	0.57; 0.952 78.22	0.70; 0.940 81.87

monly employed zero padding scheme on the training split of a pre-existing dataset $\tilde{\mathcal{A}}$ as the backbone model, i.e. training setting 1 defined in Eq. (6) and visualised in Fig. 3. From there, we evaluate the transferrable capability of such a model by adopting six padding schemes (include zero padding) on the testing split of the pre-existing and lastly released dataset, i.e. $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$.

Testing performance on these two tests is summarised in Table 1 and visualised in Fig. 5. It could be generally observed from Table 1 that reflection padding contributed the most to the inference performance compared to the rest, which is consistent with the findings presented in [19]. When testing above-mentioned backbone model on $\tilde{\mathcal{A}}$, i.e. the column of $\text{Acc}(\mathcal{M}_{\text{ResNet34}}^{\tilde{\mathcal{A}}}(\tilde{\mathcal{A}}))$ in Table 1, the reflection padding outperforms zero, RGB-mean, LAB-mean, white and grey padding schemes by 2.53%,

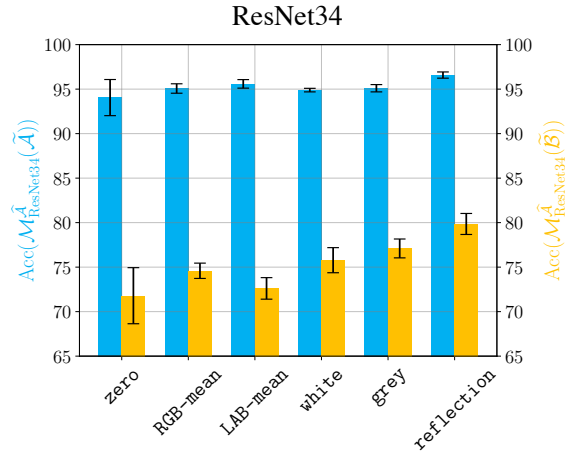


Figure 5. Prediction performance (in %) comparisons on the testing sets of datasets \mathcal{A} (in the y -axis on the left) and \mathcal{B} (in the y -axis on the right) using the baseline models (ResNet34) trained on the training set of dataset \mathcal{A} that was processed by six padding schemes summarised in Table 1. Results yielded by the optimal padding scheme – reflection – are detailed in Table 2. Quantitative and visual explanations of the six padding schemes are shown in Fig. 6. Colour codes are consistent with Figs. 1 and 3.

1.51%, 0.99%, 1.69% and 1.48% in terms of mean accuracy. Impressively, the reflection padding exceeds its counterparts with a more significant margin, i.e. 8.06%, 5.26%, 7.24%, 4.07% and 2.75% improvements, respectively, for mean accuracy.

Furthermore, for selecting a single model to continue the conduction of the proposed DCF, we decompose the highlighted performance yielded by reflection padding (available in Table 1) in Table 2. One step further, in line with the columns containing ‘ \mathcal{M}_{acc} ’ values as summarised in Table 3, the peak inference performance of the backbone model on two testing sets resulting from reflection padding is 96.60% and 81.87%, as outperformed zero padding for 1.38% and 8.19%, RGB-mean padding for 1.23% and 6.21%, LAB-mean padding for 1.71% and 7.21%, white padding for 1.53% and 4.5%, and grey pa-

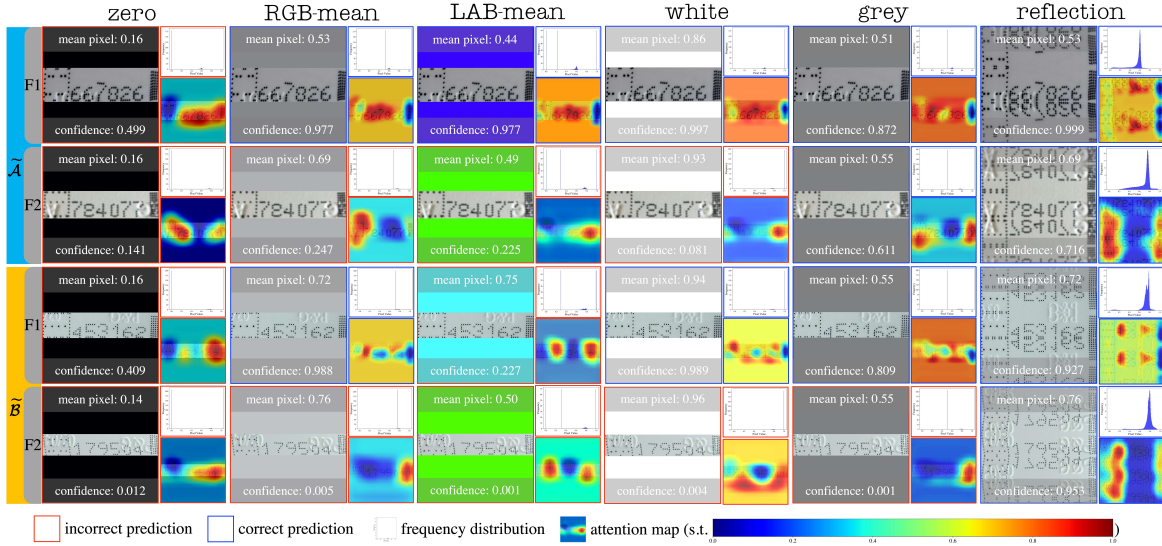


Figure 6. The relationship between the input image with six different padding schemes applied and the model (ResNet34) predictions in the binary FGVC task is revealed by the frequency distribution and model prediction confidence in the testing sets of \mathcal{A} and \mathcal{B} (Sec. 2.1), *i.e.* $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$ with consistent colour schemes adopted in Figs. 1 and 3. Best viewed in colour and zoomed mode.




Figure 7. Prediction performance (in %) comparisons on the testing sets of datasets \mathcal{A} (*i.e.* $\tilde{\mathcal{A}}$) and \mathcal{B} (*i.e.* $\tilde{\mathcal{B}}$) using the three deep neural networks under five training settings with *reflection* padding scheme applied. This figure is a visual presentation of Table 4. Colour codes are consistent with Figs. 1 and 3.

dding for 1.78% and 3.65%. Noteworthy, the above-peak performances are qualified following the criteria defined in Eq. (1), *i.e.* the leveraged performance over two testing sets using a single model as opposed to ones that performed exceptionally well on one of the two testing sets and marginally poorly on another.

Besides, Table 3 presented a quantitative explanation of the rationale that *reflection* padding outperforms its counterpart schemes. Concretely, the averaged mean pixel value of images padded using *reflection* padding in both subsets of ‘F1’ and ‘F2’ are constantly identical in the testing sets of $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$, *i.e.* 0.60 and 0.78, respectively. Such an observation confirms the robustness of the *reflection* padding schemes in the context of changing anti-counterfeit code patterns. We also provide visual explanations in Fig. 6 to justify this. Overall, the competitive performance is mainly due to the backbone model’s high de-

gree of generalisation capability. This, in turn, reveals the practical feasibility and advantages of dense input regarding the more robust averaged mean pixel value, the corresponding more even frequency distribution, and the correctness of model attention.

4.3. Optimal Training Pathway Selection

In this experiment phase, we assess the feasibility of identifying the most suitable training pathway selection from two sequential datasets, as delineated in Sec. 3.3. Prediction performances on two datasets are summarised in Table 4 and visualised in Fig. 7.

Noteworthy, the training pathways 1 and 2 are similar to 4 and 5, where 1 and 3 are the starting point of the proposed DCF, each of which was trained on with *zero* padding scheme. The third one is different among the five training pathways as it was trained on a combination of both training

Table 4. Summary of prediction performance yielded by three fine-grained visual classifiers under five training pathways that defined in Eq. (6). Results of the optimal training pathway are marked in bold and highlighted with .

Model	setting	Acc($\mathcal{M}(\tilde{\mathcal{A}})$)	Acc($\mathcal{M}(\tilde{\mathcal{B}})$)
ResNet18	1	96.80 ± 0.07	81.81 ± 0.53
	2	96.38 ± 0.32	92.32 ± 0.85
	3	94.25 ± 0.83	91.49 ± 0.44
	4	88.02 ± 1.82	90.15 ± 0.82
	5	96.11 ± 0.66	86.04 ± 0.79
ResNet34	1	96.58 ± 0.35	79.85 ± 1.18
	2	96.51 ± 0.40	92.80 ± 0.23
	3	94.65 ± 1.07	91.67 ± 0.62
	4	90.94 ± 1.15	91.16 ± 0.43
	5	96.81 ± 0.06	88.73 ± 0.58
Inception-v3	1	96.14 ± 0.93	83.33 ± 1.27
	2	96.77 ± 0.40	93.11 ± 0.48
	3	95.96 ± 0.68	91.88 ± 0.38
	4	88.69 ± 0.90	90.90 ± 0.69
	5	96.97 ± 0.27	89.02 ± 1.13

Table 5. Detailed summary of peak prediction performance yielded by three fine-grained visual classifiers under five training pathways that defined in Eq. (6). Accuracy is measured in %. Peak performance under each training pathway is marked in bold and highlighted with . The visual explanations of ResNet34 models with the peak performances are provided in Fig. 8 as an example.

Model	setting	Acc($\mathcal{M}^{\tilde{\mathcal{A}}}(\tilde{\mathcal{A}}_{F1}, \tilde{\mathcal{A}}_{F2})$)	Acc($\mathcal{M}^{\tilde{\mathcal{A}}}(\tilde{\mathcal{A}})$)	Acc($\mathcal{M}^{\tilde{\mathcal{B}}}(\tilde{\mathcal{B}}_{F1}, \tilde{\mathcal{B}}_{F2})$)	Acc($\mathcal{M}^{\tilde{\mathcal{B}}}(\tilde{\mathcal{B}})$)
ResNet18	1	(94.60, 98.78)	96.69	(78.64, 86.21)	82.42
	2	(94.82, 98.17)	96.50	(90.91, 96.35)	93.63
	3	(91.79, 99.39)	95.59	(83.64, 98.84)	91.24
	4	(91.14, 88.69)	89.91	(84.55, 98.34)	91.44
	5	(94.74, 97.55)	96.14	(79.55, 92.86)	86.20
ResNet34	1	(95.03, 98.17)	96.60	(78.18, 85.55)	81.87
	2	(95.46, 98.17)	96.81	(87.27, 98.84)	93.06
	3	(94.53, 96.64)	95.59	(84.09, 99.17)	91.63
	4	(92.30, 92.97)	92.63	(83.64, 98.67)	91.16
	5	(94.74, 98.78)	96.76	(80.45, 98.34)	89.40
Inception-v3	1	(94.89, 99.69)	97.29	(80.00, 89.37)	84.69
	2	(95.54, 98.17)	96.86	(90.91, 95.85)	93.38
	3	(95.25, 96.94)	96.09	(86.36, 98.50)	92.43
	4	(87.69, 92.0)	89.87	(84.55, 98.34)	91.44
	5	(94.96, 99.08)	97.02	(84.09, 96.51)	90.30

sets with reflection padding scheme applied directly.

Apart from ResNet34, which was trained as the backbone model in training setting 1, we also employed ResNet18 and Inception-v3 in such a setting. The detailed performance is reported in Table 4, where the model with peak prediction performances is highlighted and used for the fine-tuning process required by training setting 2.

Unlike training settings 2 and 4, we train a model using both training sets from scratch with the qualified padding scheme (*i.e.* reflection padding). In the meantime, as the peak performance presented in Table 5 and partially explained in Fig. 8, the prediction performance resulting from training setting 2 is much better than that of setting 4, emphasising that the order of training datasets matters.

Overall, the outperformance of training setting 2 not only reveals that different padding schemes could compensate with each other but also addresses the possible overfitting of training data [15] and catastrophic forgetting of model [25].

5. Conclusion

In this paper, we proposed an optimal training pathway selection framework DCF for robust and explainable fine-grained visual classification in the context where two temporal continued datasets are provided, and a model has been pre-trained on the dataset released earlier. The DCF consists of two adjacent steps, *i.e.* determining the optimal padding scheme and selecting the optimal training pathway, each equipped with a quantitative and visual explanation in a *putting-through* fashion. The experimental results first indicate that the best-suited padding scheme determination is a feasible starting point for deploying the proposed DCF, as it reveals great potential in enhancing the eventual classification

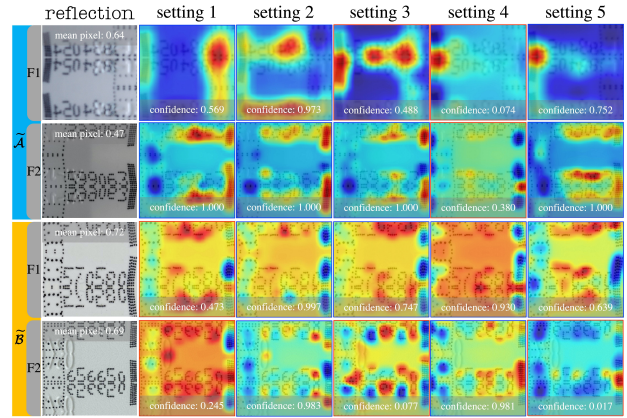


Figure 8. Visual explanations supported by five ResNet34 models trained using various training pathways as summarised in Table 5. The adopted colour codes are consistent with Fig. 6. The second training pathway (*i.e.* setting 2) has demonstrated its compatibility with blurry and clean images (first two rows) and illustrated its discriminability in coping with finer patterns, *e.g.* printed dents (last row) in the testing sets of \mathcal{A} and \mathcal{B} (*i.e.* $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{B}}$).

performance. The efficiency and efficacy of DCF are then confirmed by its second step of figuring out the optimal training pathway, where the classification accuracy over two datasets is maximally enhanced and leveraged. Experimental results also inspire further investigations on image-wise quality separation [54] of the optimal learning pathway selection for boosted image categorisation with mitigated data bias [47] and reduced computational cost.

Acknowledgment. This work was supported by Procter & Gamble (P&G) United Kingdom Technical Centres Ltd.

References

- [1] Amazon Mechanical Turk. <https://www.mturk.com/>. 3
- [2] Fahad Alrasheedi, Xin Zhong, and Pei-Chi Huang. Padding module: Learning the padding in deep neural networks. *IEEE Access*, 11:7348–7357, 2023. 3
- [3] Matteo Boschini, Lorenzo Bonicelli, Angelo Porrello, Giovanni Bellitto, Matteo Pennisi, Simone Palazzo, Concetto Spampinato, and Simone Calderara. Transfer without forgetting. In *Proc. ECCV*, 2022. 2
- [4] Dongliang Chang, Kaiyue Pang, Ruoyi Du, Yujun Tong, Yi-Zhe Song, Zhanyu Ma, and Jun Guo. Making a bird AI expert work for you and me. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(10):12068–12084, 2023. 1
- [5] Chaoran Chen, Mai Xu, Shengxi Li, Tie Liu, Minglang Qiao, and Zhuoyi Lv. Residual based hierarchical feature compression for multi-task machine vision. In *Proc. ICME*, pages 1463–1468, 2023. 3, 4
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2017. 3
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. CVPR*, 2009. 2, 3
- [8] Ruoyi Du, Dongliang Chang, Ayan Kumar Bhunia, Jiyang Xie, Zhanyu Ma, Yi-Zhe Song, and Jun Guo. Fine-grained visual classification via progressive multi-granularity training of jigsaw patches. In *Proc. ECCV*, 2020. 1
- [9] Ruoyi Du, Jiyang Xie, Zhanyu Ma, Dongliang Chang, Yi-Zhe Song, and Jun Guo. Progressive learning of category-consistent multi-granularity features for fine-grained visual classification. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(12):9521–9535, 2021. 1
- [10] Ruoyi Du, Wenqing Yu, Heqing Wang, Ting-En Lin, Dongliang Chang, and Zhanyu Ma. Multi-view active fine-grained visual recognition. In *Proc. ICCV*, 2023. 1
- [11] Logan Engstrom, Brandon Tran, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. Exploring the landscape of spatial robustness. In *Proc. ICLR*, 2019. 3, 4
- [12] Jianlong Fu, Heliang Zheng, and Tao Mei. Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In *Proc. CVPR*, 2017. 1
- [13] Ying Fu, Yang Hong, Yunhao Zou, Qiankun Liu, Yiming Zhang, Ning Liu, and Chenggang Yan. Raw image based over-exposure correction using channel-guidance strategy. *IEEE Trans. Circuits Syst. Video Technol.*, 34(4):2749–2762, 2024. 2
- [14] Zhigang Fu and Rui Zhang. x-net: Dual supervised medical image segmentation with multi-dimensional self-attention and diversely-connected multi-scale convolution. *Neurocomputing*, 500:177–190, 2022. 3
- [15] Yunhui Guo, Yandong Li, Liqiang Wang, and Tajana Rosing. Adafilter: Adaptive filter fine-tuning for deep transfer learning. In *Proc. AAAI*, pages 4060–4066, 2020. 8
- [16] Mahdi Hashemi. Enlarging smaller images before inputting into convolutional neural network: zero-padding vs. interpolation. *J. Big Data*, 6(1):1–13, 2019. 3
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. CVPR*, pages 770–778, 2016. 3, 4
- [18] Wei He, Meiqing Wu, Mingfu Liang, and Siew-Kei Lam. Cap: Context-aware pruning for semantic segmentation. In *Proc. WACV*, 2021. 3
- [19] Md Amirul Islam, Matthew Kowal, Sen Jia, Konstantinos G Derpanis, and Neil DB Bruce. Position, padding and predictions: A deeper look at position information in cnns. *arXiv preprint arXiv:2101.12322*, 2021. 3, 6
- [20] Philip T. Jackson, Stephen Bonner, Ning Jia, Christopher Holder, Jon Stonehouse, and Boguslaw Obara. Camera bias in a fine grained classification task. In *Proc. IJCNN*, 2021. 2
- [21] Ning Jia, Christopher Holder, Stephen Bonner, and Boguslaw Obara. Coarse annotation refinement for segmentation of dot-matrix batchcodes. In *Proc. ICMLA*, pages 2001–2007, 2019. 2, 3, 4
- [22] Paul Jungmann, Julia Poray, and Akash Kumar. Analytical uncertainty propagation in neural networks. *IEEE Trans. Neural Networks Learn. Syst.*, pages 1–14, 2024. 2
- [23] Xiao Ke, Xinru Lin, and Wenzhong Guo. S-pad: self-learning padding mechanism. *Mach. Vision Appl.*, 34(6):102, 2023. 3
- [24] Diederik P Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *Proc. ICLR*, 2014. 5
- [25] Yajing Kong, Liu Liu, Huanhuan Chen, Janusz Kacprzyk, and Dacheng Tao. Overcoming catastrophic forgetting in continual learning by exploring eigenvalues of hessian matrix. *IEEE Trans. Neural Networks Learn. Syst.*, 2023. 8
- [26] Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *Proc. ICLR*, 2022. 2
- [27] Yann Lecun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86(11):2278–2324, 1998. 3
- [28] Chunlei Li, Xiao Li, Xueping Wang, Di Huang, Zhoufeng Liu, and Liang Liao. Fg-agr: Fine-grained associative graph representation for facial expression recognition in the wild. *IEEE Trans. Circuits Syst. Video Technol.*, 34(2):882–896, 2023. 2
- [29] Zewen Li, Fan Liu, Wenjie Yang, Shouheng Peng, and Jun Zhou. A survey of convolutional neural networks: analysis, applications, and prospects. *IEEE Trans. Neural Networks Learn. Syst.*, 33(12):6999–7019, 2021. 1
- [30] Mingbao Lin, Bohong Chen, Fei Chao, and Rongrong Ji. Training compact CNNs for image classification using dynamic-coded filter fusion. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(8):10478–10487, 2023. 1
- [31] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear CNN models for fine-grained visual recognition. In *Proc. ICCV*, 2015. 1
- [32] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for

- irregular holes using partial convolutions. In *Proc. ECCV*, 2018. 3
- [33] Guilin Liu, Kevin J. Shih, Ting-Chun Wang, Fitsum A. Reda, Karan Sapra, Zhiding Yu, Andrew Tao, and Bryan Catanzaro. Partial convolution based padding. In *arXiv preprint arXiv:1811.11718*, 2018. 3
- [34] Guilin Liu, Aysegul Dundar, Kevin J Shih, Ting-Chun Wang, Fitsum A Reda, Karan Sapra, Zhiding Yu, Xiaodong Yang, Andrew Tao, and Bryan Catanzaro. Partial convolution for padding, inpainting, and image synthesis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(5):6096–6110, 2022. 3
- [35] Hai Liu, Cheng Zhang, Yongjian Deng, Bochen Xie, Tingting Liu, Zhaoli Zhang, and You-Fu Li. Transifc: invariant cues-aware feature concentration learning for efficient fine-grained bird image classification. *IEEE Trans. Multimedia*, pages 1–14, 2023. 1
- [36] Man Liu, Chunjie Zhang, Huihui Bai, Riquan Zhang, and Yao Zhao. Cross-part learning for fine-grained image classification. *IEEE Trans. Image Process.*, 31:748–758, 2021. 1
- [37] Yang Liu, Lei Zhou, Pengcheng Zhang, Xiao Bai, Lin Gu, Xiaohan Yu, Jun Zhou, and Edwin R Hancock. Where to focus: Investigating hierarchical attention relationship for fine-grained visual classification. In *Proc. ECCV*, 2022. 1
- [38] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proc. CVPR*, pages 3431–3440, 2015. 3
- [39] Angela Lopez-del Rio, Maria Martin, Alexandre Perera-Lluna, and Rabie Saidi. Effect of sequence padding on the performance of deep learning models in archaeal protein functional prediction. *Sci. Rep.*, 10(1):14634, 2020. 3
- [40] Weiqing Min, Zhiling Wang, Yuxin Liu, Mengjiang Luo, Liping Kang, Xiaoming Wei, Xiaolin Wei, and Shuqiang Jiang. Large scale visual food recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(8):9932–9949, 2023. 1
- [41] Behnam Neyshabur, Hanie Sedghi, and Chiyuan Zhang. What is being transferred in transfer learning? In *Proc. NeurIPS*, 2020. 2
- [42] Anh-Duc Nguyen, Seonghwa Choi, Woojae Kim, Sewoong Ahn, Jinwoo Kim, and Sanghoon Lee. Distribution padding in convolutional neural networks. In *Proc. ICIP*, pages 4275–4279, 2019. 3
- [43] Shuteng Niu, Yongxin Liu, Jian Wang, and Houbing Song. A decade survey of transfer learning (2010–2020). *IEEE Trans. Artif. Intell.*, 1(2):151–166, 2020. 2
- [44] Xolisani Nkwentsha, Anicet Hounkanrin, and Fred Nicolls. Automatic classification of medical X-Ray images with convolutional neural networks. In *Proc. ISRPC*, 2020. 3
- [45] Mohammad Norouzi, Mani Ranjbar, and Greg Mori. Stacks of convolutional restricted boltzmann machines for shift-invariant feature learning. In *Proc. CVPR*, pages 2735–2742, 2009. 3, 4
- [46] Oyebede K Oyedotun, Kassem Al Ismaeil, and Djamila Aouada. Why is everyone training very deep neural network with skip connections? *IEEE Trans. Neural Networks Learn. Syst.*, 34(9):5961–5975, 2022. 3
- [47] Maan Qraitem, Kate Saenko, and Bryan A Plummer. Bias mimicking: A simple sampling approach for bias mitigation. In *Proc. CVPR*, pages 20311–20320, 2023. 8
- [48] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proc. MICCAI*, pages 234–241, 2015. 2, 3, 4
- [49] Ryosuke Sakai, Tomokazu Kaneko, and Soma Shiraishi. Framework for fine-grained recognition of retail products from a single exemplar. In *Proc. ICKST*, 2023. 1
- [50] Amir M Sarfi, Zahra Karimpour, Muawiz Chaudhary, Nasir M Khalid, Mirco Ravanelli, Sudhir Mudur, and Eugene Belilovsky. Simulated annealing in early layers leads to better generalization. In *Proc. CVPR*, pages 20205–20214, 2023. 5
- [51] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proc. ICCV*, pages 618–626, 2017. 5
- [52] Zhiqiang Shen, Zechun Liu, Jie Qin, Marios Savvides, and Kwang-Ting Cheng. Partial is better than all: Revisiting fine-tuning strategy for few-shot learning. In *Proc. AAAI*, pages 9594–9602, 2021. 5
- [53] Ensheng Shi, Yanlin Wang, Hongyu Zhang, Lun Du, Shi Han, Dongmei Zhang, and Hongbin Sun. Towards efficient fine-tuning of pre-trained code models: An experimental study and beyond. In *Proc. SIGSOFT*, pages 39–51, 2023. 5
- [54] Joseph Smith, Zheming Zuo, Jonathan Stonehouse, and Boguslaw Obara. How quality affects deep neural networks in fine-grained image classification. In *Proc. VISAPP*, pages 448–457, 2024. 8
- [55] Shikhar Srivastava, Mohammad Yaqub, and Karthik Nandakumar. Lifelong learning of task-parameter relationships for knowledge transfer. In *Proc. CVPR*, 2023. 2
- [56] Sainbayar Sukhbaatar and Rob Fergus. Learning from noisy labels with deep neural networks. *arXiv preprint arXiv:1406.2080*, 2(3):4, 2014. 5
- [57] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proc. CVPR*, pages 2818–2826, 2016. 3, 4
- [58] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(10):3349–3364, 2020. 3
- [59] Yaming Wang, Vlad I Morariu, and Larry S Davis. Learning a discriminative filter bank within a cnn for fine-grained recognition. In *Proc. CVPR*, 2018. 1
- [60] Yulong Wang, Xiaolu Zhang, Lingxi Xie, Jun Zhou, Hang Su, Bo Zhang, and Xiaolin Hu. Pruning from scratch. In *Proc. AAAI*, 2020. 2
- [61] Zhenghua Xu, Shijie Liu, Di Yuan, Lei Wang, Junyang Chen, Thomas Lukasiewicz, Zhigang Fu, and Rui Zhang. ω -net: Dual supervised medical image segmentation with multi-dimensional self-attention and diversely-connected multi-scale convolution. *Neurocomputing*, 500:177–190, 2022. 3

- [62] Wenhan Yang, Ye Yuan, Wenqi Ren, Jiaying Liu, Walter J. Scheirer, Zhangyang Wang, Taiheng Zhang, et al. Advancing image understanding in poor visibility environments: A collective benchmark study. *IEEE Trans. Image Process.*, 29: 5737–5752, 2020. [1](#)
- [63] Ye Zhang, Yi Hou, Kewei OuYang, and Shilin Zhou. Multi-scale signed recurrence plot based time series classification using inception architectural networks. *Pattern Recognit.*, 123:108385, 2022. [3](#)
- [64] Shumin Zhu, Xingxing Zou, Jianjun Qian, and Wai Keung Wong. Learning structured relation embeddings for fine-grained fashion attribute recognition. *IEEE Trans. Multimedia*, 26:1652–1664, 2023. [1](#)
- [65] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proc. IEEE*, 109(1):43–76, 2020. [2](#)
- [66] Zheming Zuo, Xinyu Chen, Han Xu, Jie Li, Wenjuan Liao, Zhi-Xin Yang, and Shizheng Wang. Idea-net: Adaptive dual self-attention network for single image denoising. In *Proc. WACV*, 2022. [2](#)