

Diversity is Definitely Needed: Improving Model-Agnostic Zero-shot Classification via Stable Diffusion

Jordan Shipard¹, Arnold Wiliem^{1,2}, Kien Nguyen Thanh¹, Wei Xiang³, and Clinton Fookes¹

¹Signal Processing, Artificial Intelligence and Vision Technologies (SAIVT), Queensland University of Technology, Australia

²Sentient Vision Systems, Australia

³School of Computing, Engineering and Mathematical Sciences, La Trobe University, Australia

{jordan.shipard@hdr., k.nguyenthanh@, c.fookes@}qut.edu.au, arnoldw@sentientvision.com, W.Xiang@latrobe.edu.au

Abstract

In this work, we investigate the problem of *Model-Agnostic Zero-Shot Classification (MA-ZSC)*, which refers to training non-specific classification architectures (downstream models) to classify real images without using any real images during training. Recent research has demonstrated that generating synthetic training images using diffusion models provides a potential solution to address MA-ZSC. However, the performance of this approach currently falls short of that achieved by large-scale vision-language models. One possible explanation is a potential significant domain gap between synthetic and real images. Our work offers a fresh perspective on the problem by providing initial insights that MA-ZSC performance can be improved by improving the diversity of images in the generated dataset. We propose a set of modifications to the text-to-image generation process using a pre-trained diffusion model to enhance diversity, which we refer to as our **bag of tricks**. Our approach shows notable improvements in various classification architectures, with results comparable to state-of-the-art models such as CLIP. To validate our approach, we conduct experiments on CIFAR10, CIFAR100, and EuroSAT, which is particularly difficult for zero-shot classification due to its satellite image domain. We evaluate our approach with five classification architectures, including ResNet and ViT. Our findings provide initial insights into the problem of MA-ZSC using diffusion models. All code is available at https://github.com/JordanHS/Diversity_is_Definitely_Needed

1. Introduction

Data is a critical element for the successful training of deep learning models [3, 24]. However, acquiring and curating a high-quality dataset can be challenging, time-

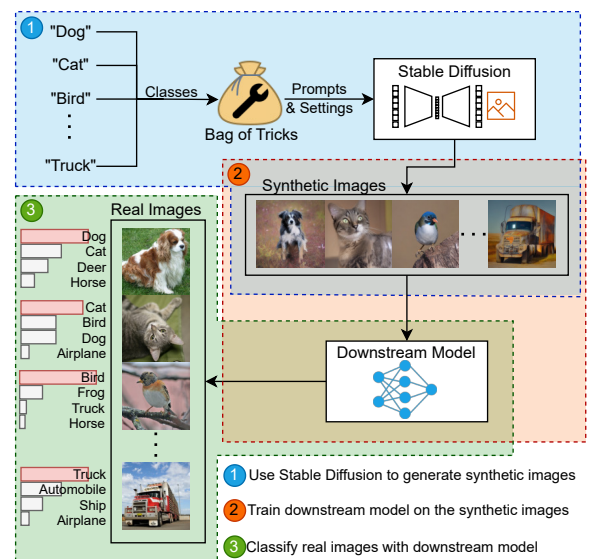


Figure 1. A diagram of our proposed model agnostic zero-shot classification method. We first generate a set of diverse training images using our proposed *bag of tricks*. These generated synthetic images are then used to train a downstream model capable of classifying a set of real images. This method achieves zero-shot performance comparable to CLIP [23].

consuming, and expensive [4, 29]. This is especially true for inherently expensive domains, such as remote sensing where the creation of a dataset requires the use of satellites [32]. The zero-shot learning tackles the data issue by training a model using no data from the downstream task [37]. Generally, zero-shot learning models use specialised architecture [18, 20, 22, 23, 35] or learn auxiliary features to achieve the objective [21]. This limits the array of potential model architectures. On the other hand, the Model-Agnostic Zero-Shot Learning problem aims to

address this issue by considering methods that allow any model architecture to perform the zero-shot learning task. In this work, we specifically investigate image classification tasks, and as such, refer to the problem domain as Model-Agnostic Zero-Shot Classification (MA-ZSC). Solving the MA-ZSC problem alleviates the previously mentioned problems of constructing a high-quality dataset without the need for specialised architecture or techniques applied to the model.

One avenue to address the MA-ZSC problem is by generating a synthetic dataset which is then used to train a classification model. A recent work [8] explored the use of generated synthetic images powered by the GLIDE diffusion model [19] primarily for improving the zero-shot and few-shot performance of CLIP [23]; however, they briefly explore training classification models, finding synthetic images to be inferior compared to real images. This is thought to be due to a domain gap between real and synthetic images [8]. We revisit the issue and find it possible to improve the MA-ZSC performance by increasing the diversity of the images in the generated datasets. Increasing diversity in training data has been utilised in various domains, such as computer vision domains [41] and robotics [33]. For instance, in domain randomisation methods in the robotics domain [33], increasing diversity in the training data would help to reduce the ‘reality gap’ that separates simulated robotics from experiments on hardware. With enough variability in the simulator, the real world may appear to the model as just another variation. Increasing training data variability is also a common practice in training classification models by employing data augmentation methods [30]. The variability of the training data is increased by simply performing data augmentation operations such as flipping, and rotations. Figure 1 presents our overall method, and we present our results as follows.

Using the latent diffusion model, Stable Diffusion [27], we start by generating baseline synthetic datasets using the prompt “an image of a {class}” for CIFAR10 [13], CIFAR100 [13] and EuroSAT [9], which is particularly difficult for zero-shot classification due to its satellite image domain, as noted in [23]. For each dataset, {class} is replaced iteratively with the class labels to generate images belonging to that class. We then train a ResNet50 model [7] on these datasets and obtain a zero-shot top-1 accuracy of **60.1%**, **29.72%** and **36.18%** respectively. For reference, the CLIP zero-shot performance with a ResNet50 backbone achieves 75.6%, 41.6%, 41.1%.

In order to improve the diversity of images in the generated dataset we propose a set of modifications to the text-to-image generation process, which we refer to as our *bag of tricks*. Each trick only improves image diversity, with no mitigation of the image domain. In fact, one of our tricks (*multi-domain*) specifically generates out-of-domain examples. After applying the proposed *bag of tricks* and find-

ing the best tricks for each dataset, we obtain top-1 zero-shot classification accuracy of **81%** (**↑20.5**) on CIFAR10, **45.63%** (**↑15.91**) on CIFAR100 and **42.5%** (**↑6.41**) on EuroSAT. Surprisingly, these results surpass the performance of CLIP-ResNet50 [23].

We list our **contributions** as follows:

1. Equipped with insights found in our work, we demonstrate the feasibility of addressing the MA-ZSC problem by training classification models on a high-quality Stable Diffusion [27] generated synthetic dataset.
2. We show that improving the diversity of images in a generated synthetic dataset improves the MA-ZSC performance.
3. We provide a *bag of tricks* for improving diversity during latent diffusion image generation.

We continue our paper as follows. In Section 2, we provide a summary of previous works related to zero-shot learning, image generation, and training with synthetic data. We then describe the problem of model-agnostic zero-shot classification, and introduce our solution in Section 3, before describing the process of how latent diffusion models generate images in Section 4. In Section 4.2, We investigate how we can improve the quality of our generated synthetic datasets. Following this, we propose our *bag of tricks* to improve the diversity of our generated synthetic datasets in Section 4.3. Finally, we present our experimental results on CIFAR10 [13], CIFAR100 [13], and EuroSAT [9] in Section 5, demonstrating the impact of our proposed *bag of tricks* on zero-shot performance across five classification models.

2. Related Work

In this section, we first show previous works relating to zero-shot learning and the introduction of CLIP; before covering the recent history of image generation and diffusion models. Lastly, we will cover previous works where training was done using generated synthetic images.

2.1. Zero-shot Learning and CLIP

Zero-shot learning (ZSL) refers to the ability of a trained model to classify classes it was not trained on [37]. For image classification, this is traditionally achieved by learning auxiliary attributes instead of pre-defined classes [21, 38]. However, a more recent and effective approach is to train an image and text encoder to learn joint representations, such as with CLIP [23]. By encoding an image and names or descriptors of the target datasets classes, zero-shot predictions can be made by finding which descriptions text embedding is the most similar to the image embedding. Using this CLIP achieves zero-shot performance on datasets such as ImageNet [4] (76.2%) comparable to supervised training of high-quality models [5, 7]. Several works look to improve the ZSL performance of CLIP by improving the text description used [18, 20, 22, 35]. In this work, we focus on improving the zero-shot learning performance of non-specific

architectures, specifically focusing on image classification. We refer to this problem as Model Agnostic Zero-Shot Classification (MA-ZSC).

2.2. Image Generation and Diffusion Models

Text-to-image generative models saw a significant performance improvement with DALL-E 1 [26]. The next generation of generative models, GLIDE [19], Latent Diffusion Models (LDM) [25], DALL-E 2 [27], and Imagen [28], use text encoders paired with diffusion models, in contrast to the discrete variational autoencoder and autoregressive transformer used by DALL-E 1. Diffusion models produce an image by denoising Gaussian noise according to some provided conditioning, such as a text prompt. Latent diffusion models perform the denoising in the latent space, making them more computationally efficient. This work uses a latent diffusion model to generate synthetic images.

2.3. Training with Synthetic Images

Since our research’s inception, several papers have investigated utilizing diffusion models to extend real datasets for domain adaptation [40], semi-supervised learning [39], or generating data augmentations [1, 34]. In our work, we concentrate solely on training downstream models on entirely synthetic generated datasets. This differs from the two closest works to ours [2, 8], as He et al. [8] studies if synthetic images from a diffusion model, GLIDE [19], can be used to fine-tune CLIPs zero-shot and few-shot performance. With respect to training downstream models on synthetic data they conclude synthetic images are $5\times$ less data efficient than real images. Additionally, He et al. [8] show synthetic images can effectively pre-train a classifier, on par with ImageNet pre-training. While Besnier et al. [2] use a GAN pre-trained on ImageNet [4] and propose strategies for improving the training quality of the generated images. With their improvements they were able to achieve 88.8% accuracy on ImageNet-10 when training on synthetic images, compared to 88.4% on real images. In our work, we generate training images using a diffusion model trained on a much larger dataset and distribution of images.

3. Model-Agnostic Zero-Shot Classification

Zero-Shot Learning (ZSL), as traditionally defined by [37], aims to learn a classifier $f^u(\cdot) : \mathcal{X} \rightarrow \mathcal{U}$, trained on labelled training instances D^{tr} belonging to a set of seen classes \mathcal{S} ($\mathcal{S} = \{c_i^s | i = 1, \dots, N_s\}$), to classify testing instances X^{te} belonging to a set of unseen classes \mathcal{U} ($\mathcal{U} = \{c_i^u | i = 1, \dots, N_u\}$). Where $\mathcal{S} \cap \mathcal{U} = \emptyset$. One method to address this problem is to describe each unseen class c_i^u with a set of image attributes extracted from the images in the seen classes \mathcal{S} [14]. To classify an unseen image, we first extract image attributes from the image. Then the classification is done in the attribute feature space by comparing the image attributes of the unseen image with each unseen class image




Figure 2. Examples of images generated from the prompt “a {caption} of a car”, where {caption} is the caption of each sub-figure. Each row shares a common initial Gaussian.

attributes. With the advent of vision-language models such as CLIP [23], the image attributes are replaced by the natural language text which is more expressive and can more accurately describe the unseen classes. In the CLIP model, the text features and image features are correlated. Thus, we can classify an unseen image by measuring the dot product between its image features, and the text features extracted from each class description.

Whilst the CLIP model has shown impressive zero-shot performance, one still needs to use the CLIP model and its zero-shot methodology. We argue that this limits the applicability of the zero-shot classification. For instance, it is non-trivial to deploy the CLIP model into edge devices due to its large model size and complexity. This motivates us to consider the **Model-Agnostic Zero-Shot Classification** (MA-ZSC) problem. In the MA-ZSC setting, we wish to use any non-specific architecture and methodology to perform the zero-shot classification task. In other words, any classification model and methodology can be used for $f^u(\cdot)$.

To address the MA-ZSC problem, we utilise a **Latent Diffusion Model** [27] (LDM) which can generate synthetic training images for each unseen class from its textual description. Once the synthetic training dataset, D^{syn} , are generated, we can then train any downstream classification model. Note that we are not claiming that our work is the first work to employ a synthetic image generation strategy to address the MA-ZSC problem. Rather, we show the feasibility of the strategy by combining the LDM with our proposed *bag of tricks* to generate more diverse synthetic images.

4. Generating Training Images

In our work, we use a Latent Diffusion Model [27] loaded with Stable Diffusion V1.4 weights to generate synthetic images. Stable Diffusion was trained on a subset of the LAION-5B dataset [29] and generates images of

512×512 pixels, which we resize to the native image size of each dataset. Stable Diffusion uses a frozen CLIP ViT-L/14 text encoder to provide conditioning from text prompts, similar to [27, 28]. Although LDMs can generate images from a number of different conditionings (image, text, semantic map), in our work we only generate images via text prompt. In the following sections, we will first describe text-to-image generation in more detail before investigating the potential domain gap between real and synthetic images. Finally, we hypothesise image diversity is more important for improving zero-shot classification and provide a *bag of tricks* for improving the synthetic image diversity.

4.1. Text-to-Image

For text-to-image generation, a prompt that describes the desired contents of the image is required to guide the diffusion process. The prompt is projected to an intermediate representation via CLIP’s text encoder and then mapped to the intermediate layers of the LDMs denoising UNet via cross-attention [27]. The LDM uses this guidance to diffuse a latent representation of an image starting from Gaussian noise. The resulting latent representation is then decoded back into the pixel domain to produce the final image. Figure 2 shows examples of generated images using different text prompts.

There are three main hyperparameters that control the generation process. **DDIM Steps** controls the number of steps taken by the **Denoising Diffusion Implicit Model** [31] in the denoising process. More steps generally result in more realistic and coherent images, while fewer result in more disjointed surreal images. The images in Figure 2, and all synthetic images used in this work, were generated with 40 DDIM steps. **Unconditional Guidance Scale (UGC)** controls the scale between the precision of the generated image matching the provided prompt and generation diversity. This is done by scaling between the jointly trained conditional and unconditional diffusion models [11]. A lower UGC value means less guidance and therefore more diversity and vice versa. Lastly, there is the **Seed** from which the initial Gaussian is generated, and serves as the starting point for diffusion. The images in each row of Figure 2 were all generated from common seeds, resulting in the cars in each row sharing similar features, such as car shape, position and colour.

4.2. Improving Synthetic Data for Training

Before attempting to construct a high-quality synthetic dataset we first validate that our synthetic images contain enough semantic features for classification. In a prior work, He et al. [8] conclude there is a domain gap between real and synthetic images generated via their chosen diffusion model GLIDE [19]. They suggest that reducing this gap is necessary to improve the quality of synthetic images for training purposes. If the domain gap is significant we expect




Figure 3. T-SNE plot comparing clustering of real and synthetic images. The lack of separation in the feature space suggests that the synthetic images contain semantically meaningful features.




Figure 4. Two t-SNE plots showing the clustered classes for real and synthetic images. The relative classes are clustered together, demonstrating the ability of the classifier to classify synthetic images despite only being trained on real images.

it to noticeably impact our initial validation experiment. If it does not, however, we then need to investigate what factors do impact the quality of a synthetic dataset.

4.2.1 Validating the Usefulness of Synthetic Images

One approach for validating the quality of synthetic images is the Fréchet Inception Distance (FID) score [10]. Recent generative diffusion models achieve FID scores as low as 7.27 [28] on the MS-COCO dataset [15], with Stable Diffusion (V1.4) achieving a score of 16. For reference, real images overlaid with 25% Gaussian noise result in FID scores of approximately 50 [10]. With this in mind, we conjecture that the low FID scores of recent diffusion models suggest they are all capable of generating realistic images. In order to examine if synthetic images contain semantically meaningful features we take a ResNet50 model [7] with pre-trained ImageNet weights [4] and fine-tune only the classifier head on the real CIFAR10 [6] dataset. We then visualise the feature space of the real images and our generated synthetic CIFAR10 images using a t-SNE plot [36]. If the synthetic images contain vastly different semantic features compared to real images we expect the model to fail in classifying them. Additionally, we expect the real and synthetic images to occupy different areas of the feature space. However, in Figure 3 we can see the features of the real and synthetic images are intermingled. Furthermore, in Figure 4 we see the classes are closely clustered, with the model correctly classifying **76.61%** of the real images and **63.8%** of

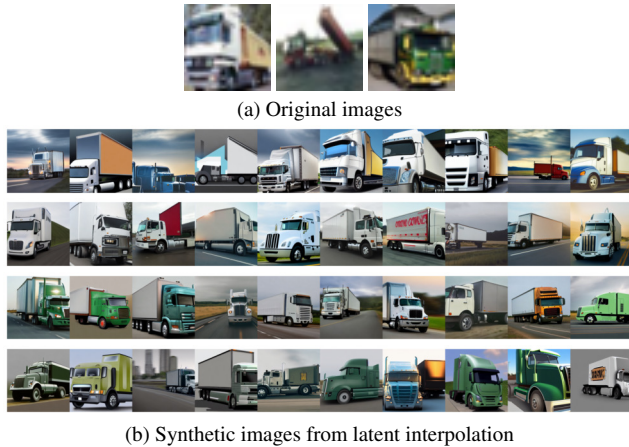


Figure 5. Examples of images generated from the combination of different initial latents. Row 1 of (b) is from the left most image of (a). Row 2 is the average of the centre and left most images. Row 3 is the average of all three images in (a). Row 4 is random combinations of all three images in (a)

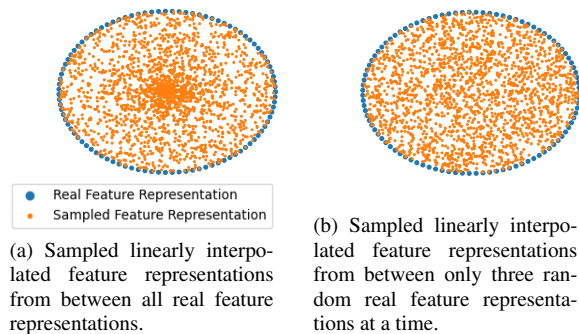


Figure 6. A visual example for demonstrating the difference in resulting diversity of sampled latent points from different sampling techniques. Sampling from a linearly interpolated combination of all real feature representations (6a) results in a bias towards the centre of the feature space; whereas sampling from three randomly selected real feature representations (6b) results in a more uniform and diverse sampling.

the synthetic images. This initial experiment validates that the synthetic images contain enough meaningful semantic features and therefore should be useful as training images.

4.2.2 Investigating Synthetic Image Diversity

Inspired by works in domain randomisation [33, 41] and data augmentation [30], which demonstrate that increasing diversity in the training data improves performance. We investigate if diversity impacts the quality of a synthetic training dataset. In order to isolate the impact of diversity, we utilise the Real Guidance (RG) technique used in [8], with a slight modification, to minimise the potential domain

gap. RG minimises the domain gap by replacing the Gaussian noise used in the diffusion process with a real image overlaid with Gaussian noise. We modify this method by generating images from the interpolated feature representations of real images instead of from the real images themselves. More explicitly, we randomly sample 1% of the images for each class in CIFAR10 [13], totalling 60 images per class. These images were then encoded using CLIP’s image encoder to obtain their feature representations. Next, we perform linear interpolation of these representations and use the interpolated feature representation as conditioning for the diffusion image generation process. Figure 5 shows an example of this where we generate synthetic images of trucks (Fig. 5b) from the feature representations of the original truck images (Fig. 5a). The first row in Figure 5b are images generated from only the feature representation of the first image in Figure 5a, the second row is from the average of the first and second images and the third is average of all three. We can qualitatively observe improved synthetic image diversity from rows one to three. In row four we generate images from random interpolations between all three original images.

We use our 60 initial feature representations to generate images via linear interpolation using two sampling methods. Firstly, we sample a starting feature representation from within the convex hull of all 60 feature representations. Secondly, we sample from only three random feature representations. We then train a ResNet-50 [7] model on these two generated datasets and validate performance on the real CIFAR10 [13] test set. The average of all latent combinations obtains **35.01%** top-1 test accuracy and the three latent combinations obtain **52.6%**. This shows data diversity is important when generating datasets and improving zero-shot accuracy is dependent on generating diverse training examples.

In Figure 6, we provide a visual example that more intuitively illustrates the differences between the two linear interpolation sampling methods. This is purely for demonstration purposes and does not represent the actual feature representations. We first generate a number of evenly spaced points (shown in blue) along the circumference of a circle, representing the feature representations of real images. We then use the two sampling methods described earlier to generate linearly interpolated feature representations (shown in orange). Figure 6a shows a significant bias towards the centre of the simulated feature space, resulting in less diverse samples. On the other hand, Figure 6b shows a more uniform and diverse set of samples.

4.3. Bag of Tricks for Improving Synthetic Diversity

Here we propose the tricks we can utilise to improve the diversity of generated images. We first describe these tricks below and later validate their effectiveness in improving zero-shot classification performance. For all synthetic

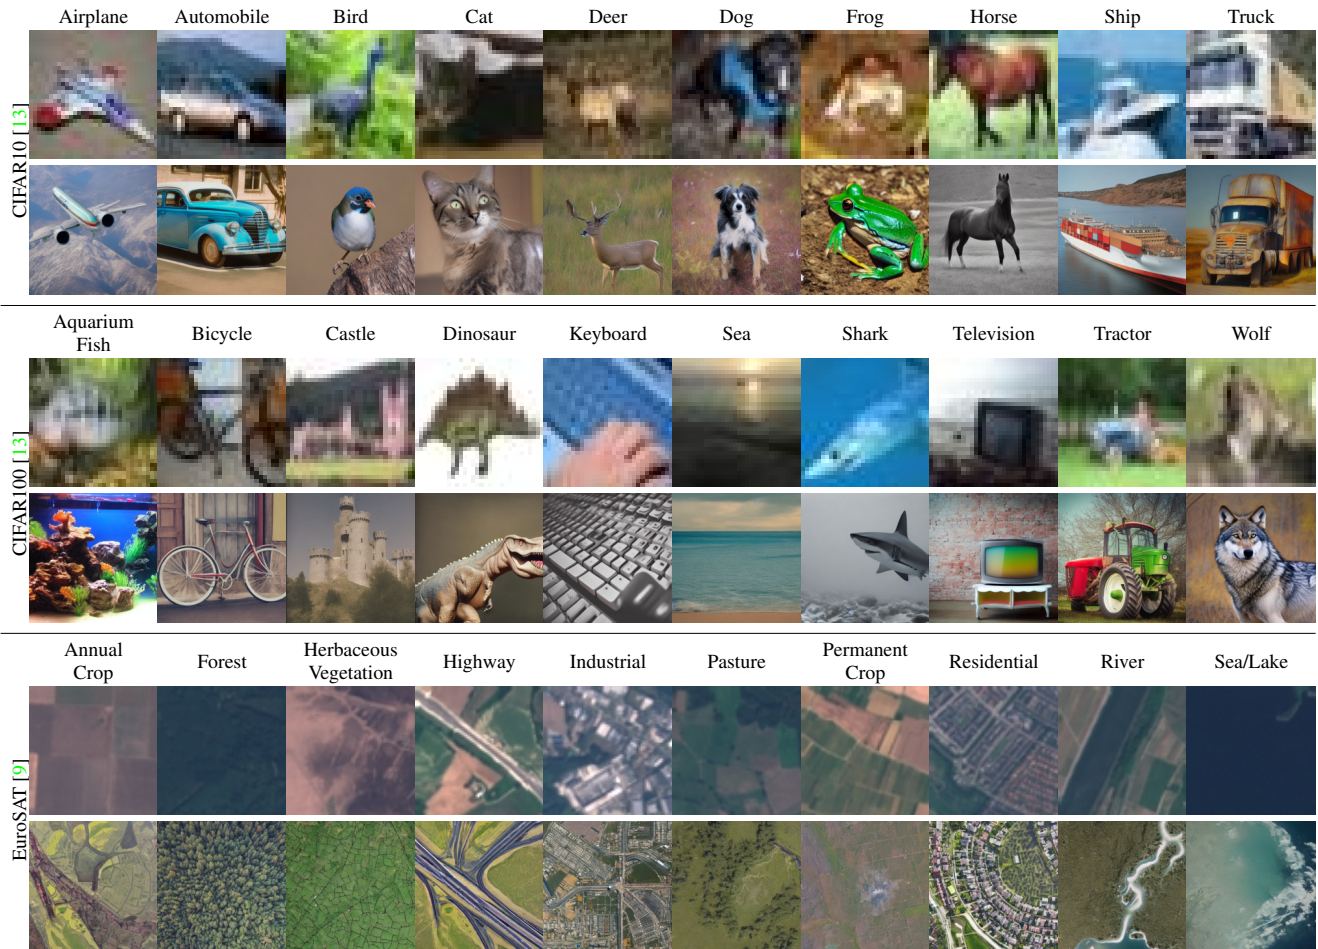


Table 1. Examples of real (top row per dataset) and generated (bottom row per dataset) images for the CIFAR10 [13], CIFAR100 [13] and EuroSAT [9], the generated images are taken from the base class generated dataset. The real images are 32x32 for CIFAR and 64x64 for EuroSAT, while the synthetic images are 512x512. The generated images are downsized to match the real images during training.

datasets, we generate the same number of images per class as their real counterparts.

Base Class - This consists of images that were generated using the prompt “an image of a {class}” where {class} is replaced with a class name from the downstream dataset. This represents the naive case for generating images as this prompt is synonymous with the prompt used for zero-shot prediction in [23]. We only slightly modify the zero-shot prompt from “a photo of” to “an image of” as we use the prompt “a photo of” later in the *multi-domain* trick.

Class Prompt - We change the prompt from “an image of a {class}” to just “{class}” as images generated from “an image of a {class}” are included in the subset of images generated by the prompt “{class}”. Therefore, using just the class name may lead to more diverse outputs.

Multi-Domain - Next, we directly influence the diversity by providing a list of domains with the prompt “a {domain} of a {class}” where the domain is one of ten preset domains (photo, drawing, painting, sketch, collage, poster, digital

art image, rock drawing, stick figure, 3D rendering) for CIFAR datasets. Due to EuroSAT requiring more domain information to correctly generate a satellite image, we use the prompt “a satellite photo of a {class} in the style of a {domain}” where the domains are (realistic photo, drawing, painting, sketch, 3D rendering). The images in Figure 2 are an example of CIFAR10 *multi-domain* images.

Random Unconditional Guidance - We use the *base class* prompt and randomly set the unconditional guidance scale between values of 1 and 5. This generates images that are highly diverse, such as generated with $UCG = 1$, as well as images containing stronger features of the target class, such as with $UCG = 5$. $UCG = 5$ was chosen as the upper bound from qualitative inspection of generated images where we found little difference in features of synthetic images with values greater than 5.

All Combined - Lastly, we combine all previous tricks into one final dataset. This should result in a more diverse dataset than any individual dataset.

	CIFAR10 [13]	CIFAR100 [13]	EuroSAT [9]
He et al. (ResNet50) [8]	-	28.74	-
CLIP-ResNet50 [23]	75.6	41.6	41.1
Base Class	60.5	29.72	36.18
Anti-aliasing Rescale	63.84 \uparrow 3.34	33.61 \uparrow 3.89	34.4 \downarrow 1.78
Class Prompt	62.32 \uparrow 1.82	26.4 \downarrow 3.32	-
Multi-Domain	67.97 \uparrow 7.47	32.55 \uparrow 1.96	35.68 \downarrow 0.5
Random Guidance	72.93 \uparrow 12.43	31.19 \uparrow 1.47	40.18 \uparrow 4
All Combined	81 \uparrow 20.5	45.63 \uparrow 15.91	39.92 \uparrow 3.74

Table 2. Zero-shot classification top-1 test accuracy on the CIFAR10, CIFAR100 and EuroSAT datasets from training on different permutations of synthetic datasets with a ResNet50 model. The change (\downarrow , \uparrow) in top-1 accuracy is measured with respect to the *base class*.

5. EXPERIMENTS

In this section, we first discuss our training setup and then present our baseline zero-shot results with ResNet50 [7]. We then apply the *bag of tricks* when generating our synthetic datasets of CIFAR10 [13], CIFAR100 [13] and EuroSAT [9]. We have chosen the CIFAR datasets due to their widespread use and low resolution, allowing for ease of training. We have also included EuroSAT because of its low resolution and challenging domain for synthetic images. Additionally, it has been shown to be a challenging dataset for zero-shot classification [23]. We show examples of our generated datasets in Table 1.

Following the *bag of tricks*, we have identified the the best tricks for each dataset and tested them on four additional classification architectures; ResNet101 [7], MobileNetV3 [12], ViT [5] and ConvNeXt [16].

5.1. Training Setup

All models were trained from random initialisation for 200 epochs, with a batch size of 128, the AdamW optimiser [17] and cosine annealing learning rate decay. All training used an initial learning rate of $2e^{-4}$. MobileNetV3 models used a weight decay of 0.1 for all training; whereas all other models used a weight decay of 0.9 when training on the CIFAR datasets and 0.3 when training on EuroSAT. We found the weight decay hyperparameter of the AdamW optimiser to be important, as it helps reduce overfitting to the synthetic images.

5.2. Baseline Results

In Table 2, we gather baseline zero-shot classification results of CLIP-ResNet50 from [23] and ResNet50 from [8]. For our own baseline, we use the *base class* synthetic dataset as described in section 4.3. Our *base class* dataset achieves zero-shot accuracy of 60.5%, 29.72% and 36.18% for CIFAR10, CIFAR100 and EuroSAT respectively. In comparison, CLIP-ResNet50 from [23] achieves zero-shot accuracies of 75.6%, 41.6% and 41.1% respectively. The 15.1% and 11.88% difference between CLIP’s ResNet zero-shot and our ResNet zero-shot results show that our gen-

erated dataset does not currently capture the full diversity of CLIP’s knowledge of each class. In theory, we expect that with infinite training examples, we should achieve CLIP’s zero-shot accuracy. The baseline results from [8] on CIFAR100 are the most directly comparable to our *base class* results as these results are obtained from training a ResNet50 model from scratch on a synthetic version of CIFAR100 generated using GLIDE [8]. An important point to note however is that the results in [8] are obtained after improving the prompt quality for generating synthetic datasets, whereas our *base class* results are already higher without any improvements. This is most likely due to the difference in diffusion models used for the generation of the synthetic datasets. As mentioned, [8] use GLIDE [19] which was trained on a dataset of 250 million image and caption pairs. In contrast, we use Stable Diffusion, which was trained on 2.3 billion image caption pairs from the LAION-5B dataset [29]. This $9.2\times$ increase in training data appears to result in inherently better generative abilities.

5.3. Implementing the Bag of Tricks

Here we iterate over the *bag of tricks*, as described in Section 4.3, in an effort to improve diversity and zero-shot classification. Results are shown in Table 2 comparing the accuracy obtained using each *trick*, with a ResNet50 model. Although not a trick, we test the impact of using anti-aliasing during the rescaling of images from 512×512 pixels to 32×32 for CIFAR datasets and 64×64 for EuroSAT. We find anti-aliasing significantly benefits CIFAR datasets but not EuroSAT. Thus we do not apply anti-aliasing on EuroSAT.

Class Prompt - Using only the class name as the prompt we see a small improvement in CIFAR10 accuracy while CIFAR100 accuracy reduces. We suspect the reduction in CIFAR100 performance is due to the generation of incorrect images for classes which can have multiple meanings. Such as: ‘Apple’ generating images of the fruit and the software company Apple logo. ‘Beetle’ generating images of the insect and the Volkswagen car. ‘Orange’ generating images of orange items of clothing mainly instead of the fruit. ‘Ray’ generating Sun rays, sting rays and men. This is not the case for CIFAR10 where the class names are unambiguous. We do not test this trick on the EuroSAT dataset as this dataset requires some context in the prompt relating to satellite images.

Multi-Domain - Despite approximately 90% of the images being generated under this setting not being realistic; we see the most significant improvement for the CIFAR datasets, with a 7.47% and 1.96% improvement with CIFAR10 and CIFAR100 respectively. These images, especially the posters, paintings and drawings, are not within the real CIFAR domain. For EuroSAT, we see a slight reduction of 0.5% compared to the *base class*. Both the CIFAR and EuroSAT results show that out-of-domain training images

		Base Class	Best Tricks
Dataset	Model		
CIFAR10 [13]	CLIP-ResNet50* [23]	75.6	-
	ResNet50 [7]	60.5	81 \uparrow 20.5
	ResNet101 [7]	60.89	81.84 \uparrow 20.95
	ViT-B [5]	42.34	75.72 \uparrow 33.38
	MobileNetV3-S [12]	51.05	74.38 \uparrow 23.33
	ConvNeXt-S [16]	49.15	80.1 \uparrow 30.95
CIFAR100 [13]	CLIP-ResNet50* [23]	41.6	-
	ResNet50* [8]	28.74	-
	ResNet50	29.72	45.63 \uparrow 15.91
	ResNet101	27.66	46.63 \uparrow 18.97
	ViT-B	16.38	32.38 \uparrow 16
	MobileNetV3-S	17.78	39.64 \uparrow 21.86
	ConvNeXt-S	20.93	45.14 \uparrow 24.21
EuroSAT [9]	CLIP-ResNet50* [23]	41.1	-
	ResNet50	36.18	42.59 \uparrow 6.41
	ResNet101	34.73	37.31 \uparrow 2.58
	ViT-B	19.53	21.71 \uparrow 2.18
	MobileNetV3-S	34.08	39.13 \uparrow 5.05
	ConvNeXt-S	18.57	20.22 \uparrow 1.65

Table 3. Top-1 zero-shot accuracy of various classification models on the *Base Class* and *Best Trick* synthetic CIFAR10, CIFAR100 and EuroSAT datasets. Models marked with * are baseline results (not *base class* results) from the cited papers, similar to Table 2. The change (\uparrow) in *Best Trick* top-1 accuracy is relative to the *Base Class* top-1 accuracy.

are not the main constraint for improving the zero-shot potential of synthetic datasets.

Random Unconditional Guidance - When we directly enforce diversity over precision by setting a random unconditional guidance scale we see an improvement in zero-shot classification across all datasets. Interestingly, we see the most significant improvements in CIFAR10 and EuroSAT. We conjecture this is due to CIFAR10 and EuroSAT containing more training examples per class than CIFAR100. This supports the finding in [8] that synthetic images are less data efficient than real images. Random unconditional guidance resulting in an accuracy improvement further supports our hypothesis that increasing diversity is more important than reducing a domain gap when generating synthetic training images.

All Combined - Finally, we combine all previously generated datasets into one large dataset, in order to test if combining all the tricks, and further increasing diversity, gives more zero-shot classification improvements. In doing so we obtain our most significant improvements, further supporting our hypothesis. Surprisingly, both CIFAR zero-shot results have now surpassed the CLIP-ResNet50 zero-shot results [23], showing that our *bag of tricks* may distil the important signals or features in CLIPs understanding of a concept.

5.4. Model Agnostic Zero-shot Classification

Using our *bag of tricks* we can now endow any model with zero-shot classification capabilities. To demonstrate this we test the best tricks for each dataset on four additional

classification architectures, results are shown in Table 3. For CIFAR10, the *all combined* dataset is used for the *best tricks*. CIFAR100 uses all tricks except *class prompt*. For EuroSAT, only the *random unconditional guidance* trick improved performance, therefore to further increase diversity, we generate an additional 2700 images per class, doubling the size of the dataset. We use the ResNet101 [7] architecture in order to test if simply a deeper ResNet is able to obtain higher zero-shot performance and we see only slight improvements over ResNet50. When training with the ViT-B model [5] we see reduced performance compared to the ResNet models across all datasets. We conjecture this is due to training from scratch, as ViTs are known to benefit greatly from ImageNet pre-training [5]. Despite this, we still see an improvement in zero-shot performance when training using the *best tricks*. Lastly, we use MobileNetV3-small [12] and ConvNeXt-small [16] as examples of architectures that previously have not been used for zero-shot classification, demonstrating our approach applies to any existing model. Again we see improvements in zero-shot classification across all datasets when applying the *best tricks*.

6. Conclusion

In conclusion, we investigate the problem of Model-Agnostic Zero-Shot Classification (MA-ZSC). Where MA-ZSC aims to train any downstream classification architecture to classify real images without training on any real images. We investigated how to improve the quality of a synthetic dataset for the purpose of training and found diversity in the synthetic images to be an important factor. From this, we then proposed a set of modifications to the text-to-image generation process via diffusion models, named our *bag of tricks*. This *bag of tricks* is designed only to improve the diversity of synthetic images, with no mitigation of the potential domain gap, as reported by previous works. Applying the *bag of tricks* achieves notable improvements across five classification architectures on the CIFAR10, CIFAR100 and EuroSAT datasets. Some architectures even achieve zero-shot classification accuracies comparable to state-of-the-art zero-shot models, such as CLIP. Our findings provide initial insights into the problem of MA-ZSC using diffusion models and opens up new avenues for research in this area.

Acknowledgement

This work has been supported by the SmartSat CRC, whose activities are funded by the Australian Government’s CRC Program; and partly supported by Sentient Vision Systems. Sentient Vision Systems is one of the leading Australian developers of computer vision and artificial intelligence software solutions for defence and civilian applications.

References

- [1] Hritik Bansal and Aditya Grover. Leaving Reality to Imagination: Robust Classification via Generated Datasets. *ArXiv*, abs/2302.02503, 2023. 3
- [2] Victor Besnier, Himalaya Jain, Andrei Bursuc, Matthieu Cord, and Patrick Pérez. This dataset does not exist: training models from generated images. In *ICASSP*, 2020. 3
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *NeurIPS*, 2020. 1
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 1, 2, 3, 4
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2, 7, 8
- [6] Haibo He and Edwardo A. Garcia. Learning from Imbalanced Data. *TKDE*, 2009. 4
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 4, 5, 7, 8
- [8] Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and Xiaojuan Qi. Is synthetic data from generative models ready for image recognition? In *ICLR*, 2023. 2, 3, 4, 5, 7, 8
- [9] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *JS-TARS*, 2019. 2, 6, 7, 8
- [10] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *NeurIPS*, 2017. 4
- [11] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS Workshop on Deep Generative Models and Downstream Applications*, 2021. 4
- [12] A. Howard, M. Sandler, B. Chen, W. Wang, L. Chen, M. Tan, G. Chu, V. Vasudevan, Y. Zhu, R. Pang, H. Adam, and Q. Le. Searching for mobilenetv3. In *ICCV*, 2019. 7, 8
- [13] Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. 2009. 2, 5, 6, 7, 8
- [14] Christoph H. Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *PAMI*, 36(3), 2014. 3
- [15] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 4
- [16] Z. Liu, H. Mao, C. Wu, C. Feichtenhofer, T. Darrell, and S. Xie. A convnet for the 2020s. In *CVPR*, 2022. 7, 8
- [17] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 7
- [18] Sachit Menon and Carl Vondrick. Visual Classification via Description from Large Language Models. *ArXiv*, abs/2210.07183, 2022. 1, 2
- [19] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *ICML*, 2021. 2, 3, 4, 7
- [20] Zachary Novack, Saurabh Garg, Julian McAuley, and Zachary C. Lipton. CHiLS: Zero-Shot Image Classification with Hierarchical Label Sets. *ArXiv*, abs/2302.02551, 2023. 1, 2
- [21] Genevieve Patterson and James Hays. SUN attribute database: Discovering, annotating, and recognizing scene attributes. In *CVPR*, 2012. 1, 2
- [22] Sarah Pratt, Rosanne Liu, and Ali Farhadi. What does a platypus look like? Generating customized prompts for zero-shot image classification. *ArXiv*, abs/2209.03320, 2022. 1, 2
- [23] 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 *ICML*, 2021. 1, 2, 3, 6, 7, 8
- [24] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 2019. 1
- [25] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical Text-Conditional Image Generation with CLIP Latents, 2022. 3
- [26] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021. 3
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3, 4
- [28] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, Seyedeh Sara Mahdavi, Raphael Gontijo Lopes, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. *ArXiv*, abs/2205.11487, 2022. 3, 4
- [29] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022. 1, 3, 7
- [30] Connor Shorten and Taghi M. Khoshgoftaar. A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data*, 2019. 2, 5

- [31] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 4
- [32] Gencer Sumbul, Marcela Charfuelan, Begüm Demir, and Volker Markl. Bigearthnet: A Large-Scale Benchmark Archive for Remote Sensing Image Understanding. In *IGARSS*, 2019. 1
- [33] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *IEEE IROS*, 2017. 2, 5
- [34] Brandon Trabucco, Kyle Doherty, Max Gurinas, and Ruslan Salakhutdinov. Effective Data Augmentation With Diffusion Models. *ArXiv*, abs/2302.07944, 2023. 3
- [35] Vishaal Udandarao, Ankush Gupta, and Samuel Albanie. SuS-X: Training-Free Name-Only Transfer of Vision-Language Models. *ArXiv*, abs/2211.16198, 2023. 1, 2
- [36] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *JMLR*, 2008. 4
- [37] Wei Wang, Vincent W Zheng, Han Yu, and Chunyan Miao. A survey of zero-shot learning: Settings, methods, and applications. *TIST*, 2019. 1, 2, 3
- [38] Yongqin Xian, Bernt Schiele, and Zeynep Akata. Zero-shot learning-the good, the bad and the ugly. In *CVPR*, 2017. 2
- [39] Zebin You, Yong Zhong, Fan Bao, Jiacheng Sun, Chongxuan Li, and Jun Zhu. Diffusion Models and Semi-Supervised Learners Benefit Mutually with Few Labels. *ArXiv*, abs/2302.10586, 2023. 3
- [40] Jianhao Yuan, Francesco Pinto, Adam Davies, Aarushi Gupta, and Philip Torr. Not Just Pretty Pictures: Text-to-Image Generators Enable Interpretable Interventions for Robust Representations. *ArXiv*, abs/2212.11237, 2022. 3
- [41] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain Randomization and Pyramid Consistency: Simulation-to-Real Generalization Without Accessing Target Domain Data. In *ICCV*, 2019. 2, 5