

Dual-Domain Image Synthesis using Segmentation-Guided GAN

Dena Bazazian Andrew Calway Dima Damen

Department of Computer Science, University of Bristol

{dena.bazazian, andrew.calway, dima.damen}@bristol.ac.uk

Abstract

We introduce a segmentation-guided approach to synthesise images that integrate features from two distinct domains. Images synthesised by our dual-domain model belong to one domain within the semantic-mask, and to another in the rest of the image - smoothly integrated. We build on the successes of few-shot StyleGAN and single-shot semantic segmentation to minimise the amount of training required in utilising two domains.

The method combines few-shot cross-domain StyleGAN with a latent optimiser to achieve images containing features of two distinct domains. We use a segmentation-guided perceptual loss, which compares both pixel-level and activations between domain-specific and dual-domain synthetic images. Results demonstrate qualitatively and quantitatively that our model is capable of synthesising dual-domain images on a variety of objects (faces, horses, cats, cars), domains (natural, caricature, sketches) and part-based masks (eyes, nose, mouth, hair, car bonnet). The code is publicly available¹.

1. Introduction

Automated image manipulation, particularly portrait editing, is of significant interest to researchers and practitioners alike with applications in animation, gaming and social media entertainment. Advances in Generative Adversarial Networks (GANs) [11] have led to tremendous progress and innovation, enabling the synthesis of realistic images [18, 19], and leading to techniques for image-to-image translation [24] and domain translation [25], able to synthesise a corresponding image of one domain using a reference image from another, transferring the style.

In this paper, we focus on segmentation-guided dual-domain feature integration – combining segmentation parts from one domain with an image from another in a way that is visually convincing. For example, replacing the eyes and mouth of a natural face image with those from a caricature,




Figure 1. Dual-Domain Synthesis combines segmentation-guided features from the source domain (left), and the target domain (centre) - corresponding segmentation masks are incorporated. Synthesised images (right) – top&bottom: caricature face with natural eyes/nose/mouth, middle: natural face with caricature features. The first two rows are from paired examples, i.e. the same latent code, while the third row shows an unpaired example.

so that the result is a dual-domain image. Examples are shown in Fig. 1. Note the distinction from previous mask-guided transfer methods which combine styles and/or features from the same domain [12, 26, 40].

We utilise the few-shotGAN [25] technique to generate related images of two domains from a common latent space. To create the segmentation masks efficiently, we employ the one-shot segmentation approach of RepurposingGAN [31]. We leverage these two StyleGAN-based approaches [25, 31] to obtain input images and masks, and propose a segmentation-guided weighted combination with perceptual and pixel-wise losses to synthesise the dual-domain image.

¹ <https://github.com/denabazazian/Dual-Domain-Synthesis>

To the best of our knowledge, this is the first attempt to synthesise images based on the features from two domains. The contributions can be summarised as follows. First, we introduce Dual-Domain Image Synthesis, a technique for generating images by integrating the features of two domains, guided by a semantic segmentation model. Second, we demonstrate that the approach works for images of faces, horses, cats and cars, synthesising dual-domain images from natural, caricature and sketch images. We test a variety of part-based segmentations for guidance, such as hair in faces, eyes in cats, heads of horses and car bonnets.

2. Related work

Generative Models: Deep generative models have demonstrated robust results in generating realistic images for various tasks such as image segmentation [8, 10, 28, 31], image inpainting [23], inverse graphics [39], conditional image synthesising [26, 40] and domain translations [25]. State-of-the-art GANs, such as StyleGAN [18] and StyleGAN2 [19], are capable of generating incredibly realistic images at a high resolution, making these popular models in works such as [6, 7, 25, 27, 31, 39]. The manifolds of changes in spatially localised regions for the images generated by StyleGAN are explored in [35]. Cross-domain correspondence techniques based on StyleGAN are proposed in [25, 27] to generate images from two domains using a handful of examples. We make use of [25] since it demonstrates better performance on far-domain adaptation. In this work we instead generate images which integrate the two domains into a single image.

Latent Optimisation: One of the powerful aspects of StyleGAN-based methods is the ability to generate realistic images from random latent codes and hence the option of seeking optimal codes w.r.t chosen loss criteria. A number of recent papers have adopted this approach to generate images based on perceptual loss, by comparing activation layers from convolutional networks [6, 7, 17, 38]. Optimising for the perceptual loss in [17] allows transferring the style from one image to another. Furthermore, perceptual loss is used in [6, 7] to fit a projected synthetic image into the distribution of images of a particular domain.

Image Completion: We here group works that attempt mask-guided image synthesis and inpainting. Mask-guided synthesise techniques combine source and reference images based on a given mask to generate a realistic image [12, 26, 40], while image inpainting techniques fill an arbitrary hole of a given image to generate a natural image [22, 23]. Part-based editing in synthetic images was proposed in [9] to transfer the appearance of a specific object part from a reference image to a target image. K-means is applied to the hidden layer activations of the StyleGAN generator to reveal a decomposition of the generated output into semantic objects and object-parts. However, note that all these previous part-based methods combine images

from the same domain, contrasting with our approach which seeks to integrate features from distinct domains.

Image Blending: Image blending aims to seamlessly blend an object from a source image onto a target image with lightly mask adjustment [37]. [34] combines synthesis with Gaussian-Poisson blending, and is trained using blended ground-truth. [37] combines a Poisson gradient loss, style loss, and content loss, without the need for training data. Similarly, image Harmonisation adjusts the illumination, colour, and texture of the foreground mask. All blending approaches aim to general in order to apply to any pair of images. In contrast, our proposed approach integrates semantic parts from one domain, importantly modifying the full image to ensure it belongs to a specific target domain.

3. Proposed Approach

We aim to generate a visually convincing image that integrates features of two domains, guided by part-based segmentations. Our pipeline utilises two generative models for the source and target domains along with a segmentation model for localising the corresponding mask to integrate the features of two domains. The masks indicate the semantic parts of the image that should be integrated from one domain into the other. We start by revisiting works we build on, then detail our approach.

3.1. Revisiting Few-shotGAN

The work of few-shotGAN [25] considers a main domain, for which a large-scale dataset is readily available, and trains a StyleGAN model on this data. Given few-shot examples (e.g. ten samples) from a second domain, few-shotGAN [25] utilises the StyleGAN trained on the large dataset for pretraining and transfers the diversity information from this model to a second model. The training preserves the relative similarities and differences between instances via a cross-domain distance consistency loss.

We use [25] to train two generative models, one with many-shot examples and the other adapted from the first with few-shot examples. Importantly, the two generative models \mathcal{G}_1 and \mathcal{G}_2 share the same latent space \mathcal{Z} during adaptation. This is of particular importance to our method, as we base the ability to integrate features from both domains on sampling from this common latent space. Given the same random vector $z^* \in \mathcal{Z}$, two images are synthesised, one from each generator, and thus from two domains. We refer to these as *paired* images as they share the same latent code, albeit from two domains.

3.2. Revisiting RepurposingGAN

We train a part-based semantic segmentation model, using the method from RepurposingGAN [31]. Each synthesised image is passed through a series of spatial convolutions, and a binary pixel-level segmentation is computed

from a unique generative computation which can be traced back through each convolutional layer down to the initial latent code. The strength of [31] is based on the capability to train this segmentation model with one-shot, i.e. by manually segmenting a single synthetic image.

We use RepurposingGAN [31] to train our segmentation model from a randomly synthesised image of one domain. The same segmentation model is used for semantically segmenting synthetic images from both domains. This is only plausible thanks to the adaptation method in Sec. 3.1. We demonstrate experimentally that this method is robust across domains for a variety of semantic masks.

3.3. Dual-Domain Synthesis (DDS)

We introduce Dual-Domain Synthesis (DDS) framework in this section. We consider as input to our network two generative models from two domains (Sec. 3.1). We select one of these two domains to be the ‘source’ domain (\mathcal{D}_s), and the other as the ‘target’ domain (\mathcal{D}_t)², and accordingly label the two trained generative models as \mathcal{G}_s and \mathcal{G}_t . Given one latent vector $z^* \in \mathcal{N}(\mathbf{0}, \mathbf{1})$, we generate two synthetic images such that $x_s \in \mathbb{R}^{n \times n \times 3}$, where x_s is an image from the source domain (\mathcal{D}_s), and $x_t \in \mathbb{R}^{n \times n \times 3}$, where x_t is an image from the target domain (\mathcal{D}_t). We pass both images to the trained semantic segmentation model (Sec. 3.2). Corresponding spatial masks denoting the semantic part in the source and target images are defined as y_s and y_t . Our goal is to synthesise one image, that is perceptually close to the *target* domain, but contains features from the *source* domain. The spatial masks are used for guiding the regions that should be transferred from source to target.

A naive approach to achieve dual-domain synthesis is simply to use additive ‘cut and paste’ to combine the relevant masked regions. We denote this as the crossover image (x_c). The source image x_s is masked by the corresponding regions in y_s indicating the parts we wish to integrate from the source domain. Additionally, we wish to remove/subtract the corresponding segmentation parts from the target image. We thus use the binary complement of the target mask as \bar{y}_t . Fig. 2 shows examples of the generated source and target images and their corresponding masks. The naive crossover image x_c is thus computed as:

$$x_c = (x_s \otimes y_s) + (x_t \otimes \bar{y}_t), \quad (1)$$

where \otimes is pixel-wise multiplication/masking.

Importantly, while this image contains the features from both domains, it is not a visually convincing dual-domain image. We thus aim to search the latent space \mathcal{Z} for the optimal latent vector \hat{z} that can generate our dual-domain image such that $\hat{x}_t \leftarrow \mathcal{G}_t(\hat{z})$. We use a gradient-based optimisation that iteratively updates a latent code $\hat{z} \in \mathcal{N}(\mathbf{0}, \mathbf{1})$

²Note that the source domain may or may not correspond to the few-shot domain in the adaptation, and similarly for the target domain.

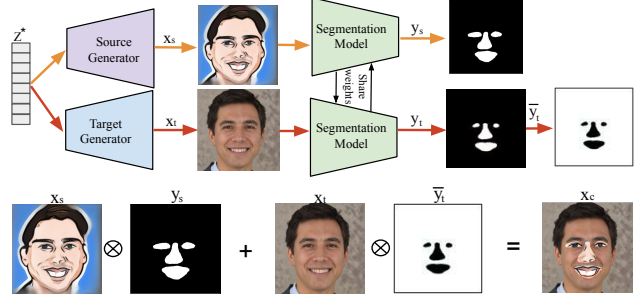


Figure 2. **Top:** A random latent code z^* is fed into two different generators to create two corresponding images from two different domains, source image x_s and target image x_t . Masks of the generated images, y_s and y_t , are then computed from a one-shot segmentation model. **Bottom:** Naive crossover image (x_c), where \bar{y}_t denotes the binary complement of y_t .

using a loss function that is based on a combination of perceptual loss [17] and pixel-level Mean Square Error (MSE).

We initialise \hat{z} randomly³. We base our perceptual loss on that proposed in [6], which compares the activations of a pre-trained convolutional network from two images, a synthesised image and a reference image. In DDS, we do not have a reference image in the dual-domain, and thus our perceptual loss considers both the source and target domains.

We use activations from the layers $conv_{1,1}$, $conv_{1,2}$, $conv_{2,2}$ and $conv_{3,3}$ of VGG-16 [29] for the perceptual loss as proposed in [6]. We downsample the segmentation binary masks at each layer to match the resolutions of the corresponding activation layer, and duplicate the mask across the channel dimension, to match the number of channels in the layer. To measure the similarity between the generated image and the expected image during optimisation, we apply a loss function that is a weighted combination of perceptual loss and the pixel-wise MSE loss:

$$\begin{aligned} \mathcal{L}_s &= \mathcal{L}_p(\mathcal{G}_s(\hat{z}), x_s, y_s) + \|\mathcal{G}_s(\hat{z}) \otimes y_s - x_s \otimes y_s\|^2 \\ \mathcal{L}_t &= \mathcal{L}_p(\mathcal{G}_t(\hat{z}), x_t, \bar{y}_t) + \|\mathcal{G}_t(\hat{z}) \otimes \bar{y}_t - x_t \otimes \bar{y}_t\|^2 \end{aligned} \quad (2)$$

where x_s, x_t are source and target images, \bar{y}_t is a binary complement of the target mask y_t . $\mathcal{G}_s(\cdot), \mathcal{G}_t(\cdot)$ are the generators for source and target domains respectively. The perceptual loss $\mathcal{L}_p(\cdot)$ in Eq. 2 is computed as :

$$\mathcal{L}_p(I_1, I_2, \lambda) = \sum_{j=1}^J \frac{1}{N_j} \|\mathcal{F}_j(I_1) \times \lambda_j - \mathcal{F}_j(I_2) \times \lambda_j\| \quad (3)$$

where I_1, I_2 are two images, λ is the mask, λ_j is the downsampled segmentation mask according to the size of each *conv* layer j . \mathcal{F}_j is the activations of the *conv* layers and N_j is the number of scalars in the j^{th} layer output, and J is

³Initial experimental evidence showed that starting from z^* produces inferior results with the optimisation stuck in local minima.

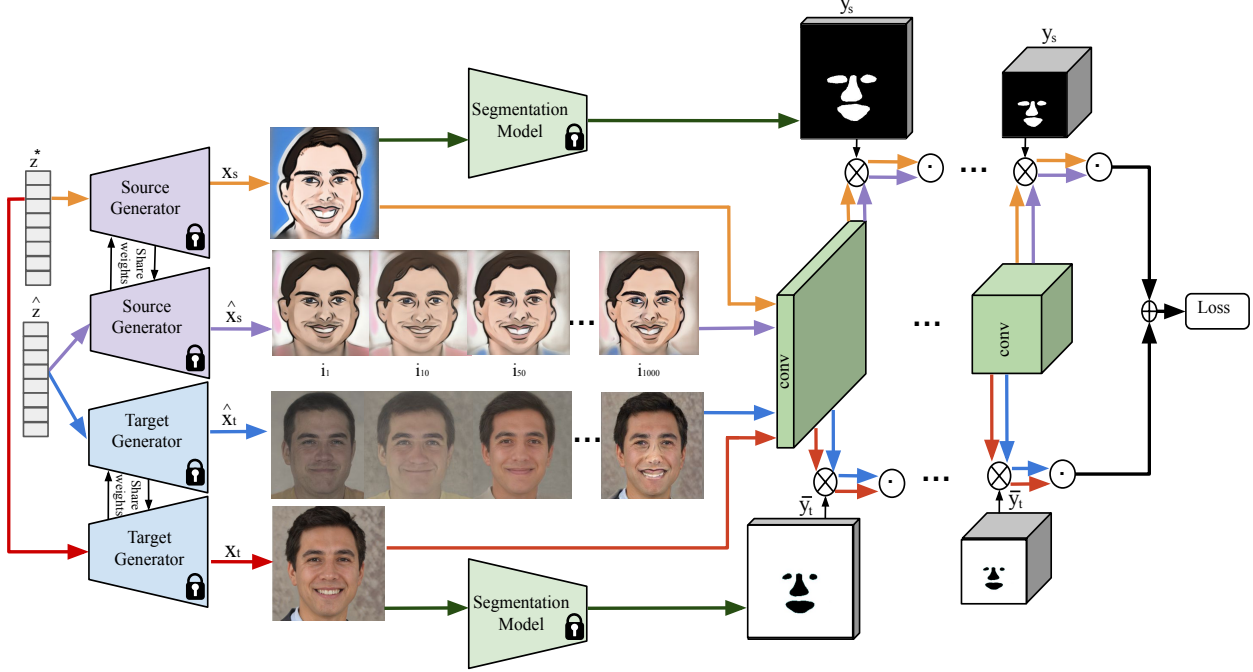


Figure 3. **Perceptual latent loss optimisation.** \odot , \otimes and \oplus indicate MSE loss, multiplication and summation, respectively. The colours of each arrow depict distinct data paths. x_s, y_s, x_t and \bar{y}_t match those illustrated in Fig. 2. z^* is fixed while \hat{z} is a random latent vector that is to be optimised. \hat{x}_s and \hat{x}_t are generated iteratively based on the optimisation of \hat{z} . Examples of \hat{x}_s and \hat{x}_t are shown for different iterations during the optimisation ($i_1, i_{10}, i_{50}, i_{1000}$).

the number of convolutional layers. Fig. 3 shows an illustration of the process of the perceptual loss.

The perceptual loss is combined with the pixel-wise MSE loss in Eq. 2 for integrating the features from the two domains. In order to increase the quality of the generated image, and ensure its segmentation-guided similarity to the source and target images, we use an additional MSE loss based on the naive cross over image as:

$$\mathcal{L}_c = \|s_c - x_c\| \quad (4)$$

where x_c is the naive crossover image from Eq. 1, and s_c is a naive crossover synthetic image computed each iteration i as:

$$s_c = (\mathcal{G}_s(\hat{z}) \otimes y_s) + (\mathcal{G}_t(\hat{z}) \otimes \bar{y}_t). \quad (5)$$

The overall loss is then defined as:

$$\mathcal{Loss} = \alpha \mathcal{L}_s + \beta \mathcal{L}_t + \gamma \mathcal{L}_c, \quad (6)$$

Note that we only optimise \hat{z} and do not update the weights of the convolutional network, generative networks or the domain-specific synthetic images x_s, x_t . We then use the target domain generator to produce the dual-domain image, such as: $\hat{x} \leftarrow \mathcal{G}_t(\hat{z})$. The complete process is summarised in Algorithm 1.

4. Implementation Details

We used the StyleGAN2 and few-shotGAN implementations provided in [3] and [4], and 10 images for few-shot

Algorithm 1: Dual-Domain Synthesis algorithm

Input: x_s, x_t, y_s, \bar{y}_t
Model : $\mathcal{G}_s, \mathcal{G}_t$
Parameter: α, β, γ
Result: \hat{x}

$x_c \leftarrow x_s \otimes y_s + x_t \otimes \bar{y}_t$;
 $\hat{z} \leftarrow \mathcal{N}(\mathbf{0}, \mathbf{1})$;
 $\nabla \leftarrow \text{Optimiser}$;
while $i < \text{maxIterations}$ **do**
 $\mathcal{L}_s, \mathcal{L}_t \leftarrow \text{Eq. 2}$;
 $\mathcal{L}_c \leftarrow \text{Eq. 4}$;
 $\mathcal{Loss} \leftarrow \alpha \mathcal{L}_s + \beta \mathcal{L}_t + \gamma \mathcal{L}_c$;
 $\hat{z} \leftarrow \nabla(\hat{z})$;
end
 $\hat{x} \leftarrow \mathcal{G}_t(\hat{z})$;

far-domain adaptation in each case. Segmentation models are based on the implementation of RepurposingGAN provided in [5]. We manually label semantic masks in each case as follows. We first synthesise an image by feeding a random latent code to a StyleGAN model [19]. We then manually label the mask of the part(s) of interest, and train the segmentation model as in [5].

In all experiments, we optimise DDS for 1000 iterations, using a learning rate of 0.01 with the Adam optimiser [20].

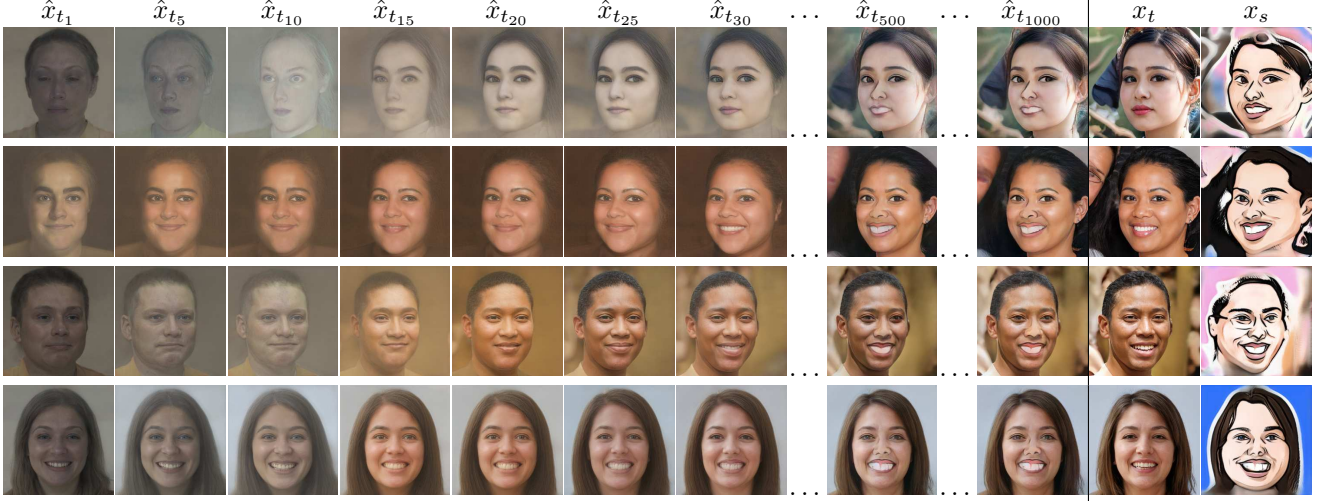


Figure 4. Intermediate DDS results for $DDS_{car \rightarrow nat}$ with a semantic mask of the eye/nose/mouth regions. Source and target domains are caricature and natural faces respectively. In the first row, note the distinct caricature eyes, nose and mouth in $i \geq 500$ iterations.

J in Eq. 3 is set to 4, given the layers chosen from *VGG-16*. Based on empirical experiments we set the α, β, γ in Eq. 6 as 0.9, 1, 0.5, respectively. Robustness to these hyperparameters was evaluated using an ablation study in Sec. 5.1.

5. Experiments

In this section, we present and discuss quantitative and qualitative experimental results of DDS on the domains of face, horse, car and cat images.

5.1. DDS on Face Domains

We perform experiments on human face images based on a StyleGAN model [3] trained on Flickr-Faces-HQ (FFHQ) [19] to generate natural human face images. Then, we use few-shotGAN model trained on Face-Caricature [25] and Face-Sketch [32] images (models from [4]). We experiment with the domain pairs natural \leftrightarrow caricature and natural \leftrightarrow sketch and two part-based segmentation masks, one indicating the regions of eyes, nose and mouth⁴ and the other indicating the hair.

Qualitative results: Fig. 4 first shows the intermediate results of DDS on the experiments for integrating the features of caricature domain into the natural face domain based on the eye/nose/mouth segmentation model. We illustrate the evolution of \hat{x}_{t_i} while the latent code \hat{z} is being iteratively optimised to generate the dual-domain image. We compare these to the target and source domains. Note that the optimisation seems to first focus on the identity and pose in the target domain, with the last iterations focusing on the dual-domain image synthesis.

⁴We combine three face parts into one mask (nose, eyes and mouth) to show more experimental results in one image.

Our qualitative results on the face domains are shown in Fig. 5. We demonstrate natural face images and their corresponding images from either caricature domain or sketch domain along with the semantic masks. We present our images in pairs, where we switch the source and target domains. For example, in the first pair, DDS incorporated the eye/nose/mouth from the caricature into the natural face. Notice how the eye colour and mouth were adapted but are distinct from that of the natural image. The nose has also been copied from the caricature image, with the skin colour integrated. In the second column, we swap the source and target domains integrating the eyes, nose and mouth from the natural image into the caricature. The natural nose is unmissable but properly integrated.

Evaluation metrics: We use standard metrics to evaluate our proposed technique. In each experiment, we consider the synthesised images as \hat{x} and images in the domain we compare to as x . Note that there is no ground-truth or examples of the target dual-domain images. We thus compare the distribution of all synthesised images to three domains: source, target and the naive crossover images. Generally, we aim for DDS images to be close to the target domain while maintaining the source features guided by the masks. We use the following metrics in our quantitative results:

- Fréchet Inception Distance (FID) [14] defined as $FID(\hat{x}, x) = \|\mu_{\hat{x}} - \mu_x\|^2 + \Gamma(\sigma_{\hat{x}} + \sigma_x - 2\sqrt{\sigma_{\hat{x}}\sigma_x})$, where $\mu_x, \sigma_x, \mu_{\hat{x}}, \sigma_{\hat{x}}$ are, respectively, the mean and covariance of x and \hat{x} , and $\Gamma(\cdot)$ is the trace operation.
- Structural Similarity Index Measure (SSIM) [33] defined as $SSIM(\hat{x}, x) = \frac{(2\mu_{\hat{x}}\mu_x + c_1)(2\sigma_{\hat{x}}\sigma_x + c_2)}{(\mu_{\hat{x}}^2 + \mu_x^2 + c_1)(\sigma_{\hat{x}}^2 + \sigma_x^2 + c_2)}$, where c_1 and c_2 stabilise the division with a weak denominator. We used $c_1 = 10^{-4}$ and $c_2 = 9 \times 10^{-4}$ in all experiments as implemented in [1].

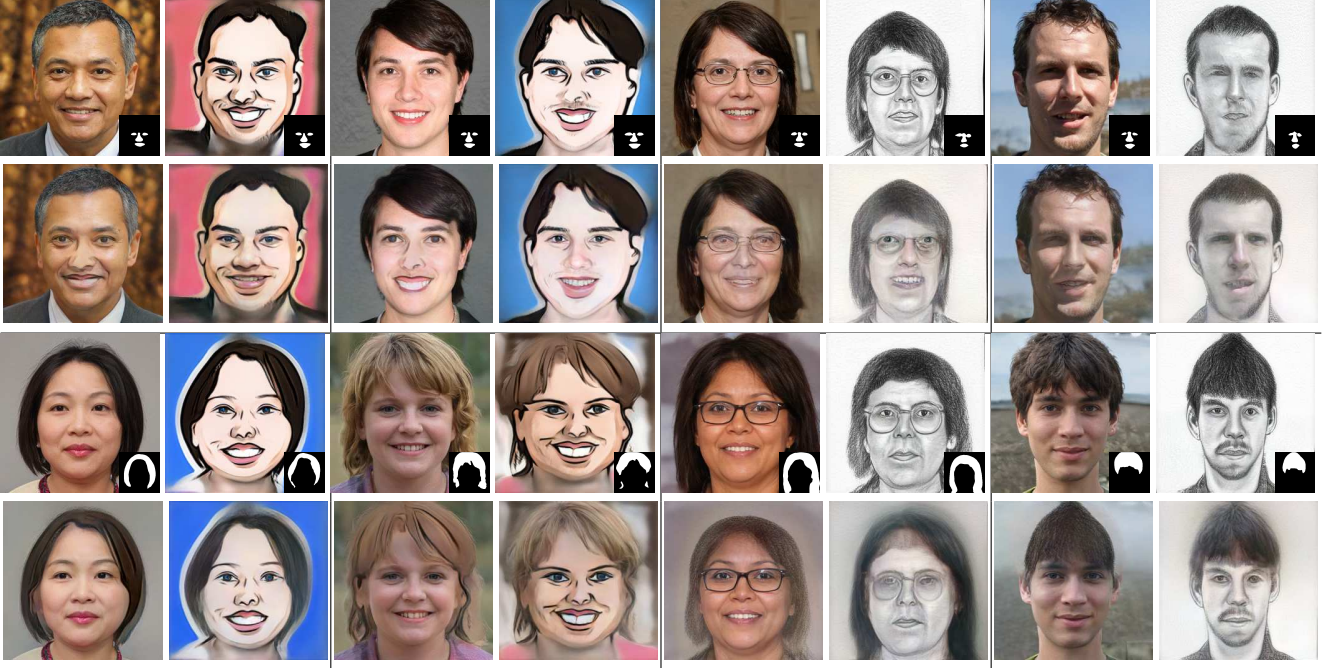


Figure 5. DDS results on human faces. Paired images are generated from the same latent code z^* . The odd columns refer to natural human face images. The second and fourth columns represent corresponding images from caricature domain. The sixth and eighth columns depict corresponding images from sketch face domain. The corresponded masks shown along with the images. The second row shows the DDS results based on the eye/nose/mouth segmentation model, while the forth row shows the DDS results based on the hair segmentation model.



Figure 6. Comparing naive crossover to $DDS_{car. \leftrightarrow nat.}$. “nat.” and “car.” denote the natural face and caricature domains respectively. We show the naive crossover $x_{ccar. \rightarrow nat.}$, as well as our DDS results $DDS_{ccar. \rightarrow nat.}$. Analogous results for $nat. \rightarrow car.$ are also shown.

- Peak Signal-to-Noise Ratio (PSNR) [15] defined as $PSNR(\hat{x}, x) = 10 \log_{10}(R^2 / MSE(\hat{x}, x))$, where $MSE(\hat{x}, x) = \frac{1}{n} \sum_{i=0}^n [\hat{x}_i - x_i]^2$, and R is the maximal in the image data.

For each DDS experiment, we use the notation $\mathcal{D}_s \rightarrow \mathcal{D}_t$ to identify the source (\mathcal{D}_s) and target (\mathcal{D}_t) domains, as we integrate features from the source domain into an image from the target domain, as illustrated in Fig. 6.

Quantitative results: Table 1 presents our results. In each reported measure, we average results from 100 randomly sampled latent codes. Note in particular that the scores confirm that our results in almost all the experiments have the better *FID*, *SSIM* and *PSNR* scores in the target domain, demonstrating the expected behaviour of the synthetic images being closer to the target domain.

Monitoring latent code optimisation: Fig. 7 plots the FID metric at each iteration of the latent optimising process. We consider $DDS_{car. \rightarrow nat.}$ and $DDS_{sketch \rightarrow nat.}$, thus using the caricature and sketch images as source and the natural faces as target with the two segmentation models \mathcal{M}_1 and \mathcal{M}_2 . We observe FID during optimisation in each case. Fig. 7 showcases how the dual-domain synthesised image gets closer to both domains but importantly significantly, and more quickly, closer to the target domain, and remains distinct from either domain at the end of the optimisation.

Robustness to hyperparameters: We next assess the robustness of our results to the parametric weights of the loss function in Eq. 6. We performed this experiment on the domains of natural human faces and caricature, and the hair segmentation model, and we consider FID over the target domain (\mathcal{D}_t) as a metric to evaluate the results of several ex-

Table 1. Metric comparisons on the face domains. $DDS_{\mathcal{D}_s \rightarrow \mathcal{D}_t}$ integrating features from the source domain into the target domain in every case. $\{x_c\}$ denotes the set of all naive crossover images given by $x_{c_{\mathcal{D}_s \rightarrow \mathcal{D}_t}}$. \mathcal{M}_1 denotes the mask of eyes/nose/mouth while \mathcal{M}_2 denotes hair mask. FID↓, SSIM↑ and PSNR↑ are the metrics, where ↓ indicates lower is better and ↑ indicates higher is better.

		$DDS_{caricature \rightarrow natural}$			$DDS_{natural \rightarrow caricature}$			$DDS_{sketch \rightarrow natural}$			$DDS_{natural \rightarrow sketch}$		
		FID↓	SSIM↑	PSNR↑	FID↓	SSIM↑	PSNR↑	FID↓	SSIM↑	PSNR↑	FID↓	SSIM↑	PSNR↑
\mathcal{M}_1	\mathcal{D}_s	281.61	0.39	27.98	259.12	0.38	27.95	195.80	0.31	27.87	225.31	0.33	27.81
	$\{x_c\}$	114.03	0.69	29.45	121.21	0.75	28.75	200.35	0.71	29.81	207.41	0.65	29.28
	\mathcal{D}_t	73.51	0.70	29.51	74.51	0.76	28.76	79.87	0.73	29.84	97.65	0.66	29.29
\mathcal{M}_2	\mathcal{D}_s	296.61	0.47	27.99	290.63	0.42	27.92	245.61	0.38	27.87	262.90	0.36	27.83
	$\{x_c\}$	155.91	0.73	29.63	95.30	0.71	28.43	177.03	0.66	29.25	152.41	0.61	28.54
	\mathcal{D}_t	87.28	0.74	29.65	76.96	0.72	28.45	90.58	0.64	29.31	94.83	0.62	28.49

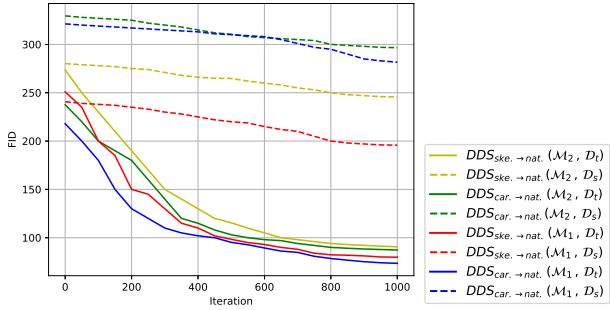


Figure 7. FID during latent code optimisation iterations, comparing the dual-domain images to source (dashed) and target (solid) domains for 4 experiments. Dual-domain images get closer to both domains during optimisation but significantly closer to target.

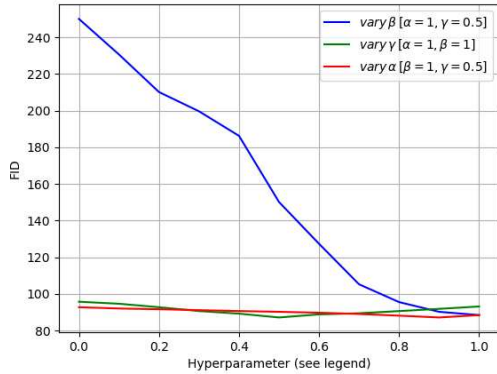


Figure 8. Impact of α , β and γ in loss function (Eq. 6) on natural face (target) and caricature domains (source) with the hair mask. FID values are computed over the target domain (\mathcal{D}_t).

periments. Fig. 8 demonstrates that by considering α, β, γ as 0.9, 1, 0.5 respectively we can have the best FID in this experiment. It can be perceived from Fig. 8 that β is the most sensitive hyperparameter in Eq. 6. The reason for this sensitivity is that β weights the target loss (\mathcal{L}_t) and these FID values are computed over the target domain. Therefore, when β is equal to zero, we do not consider the target loss (\mathcal{L}_t) in our optimisation.

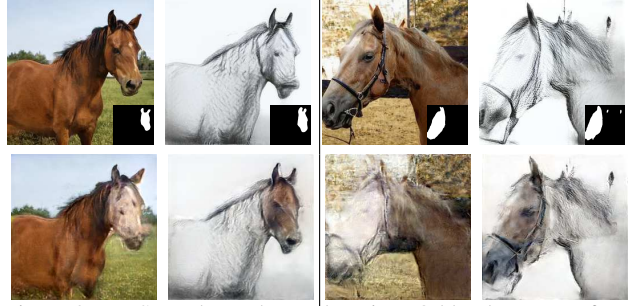


Figure 9. DDS results on horse domains. Odd columns are from natural horse domain, and even columns are from sketch horse domain. Segmentation masks are incorporated. Second row shows the corresponding dual-domain images.



Figure 10. DDS results on car domains. Odd columns are standard car images, and even columns are abandoned cars. Segmentation masks are incorporated. Second row shows the DDS results.

5.2. DDS on Horse and Car Domains

We performed further experiments on horse and car images. We used the StyleGAN2 [19] model to generate natural horse images from [2] which is trained on LSUN Horses [36]. We collected 10 sketch horses for the purpose of this experiment through a web search. We scale these images to 256×256 and use these 10 images for the few-shotGAN far-domain adaptation [25]. We train our one-shot segmentation model to segment the head of the horse. Fig. 9 shows qualitative results.

For car images, we used the StyleGAN2 [19] model to generate the standard car images from [2], which is trained

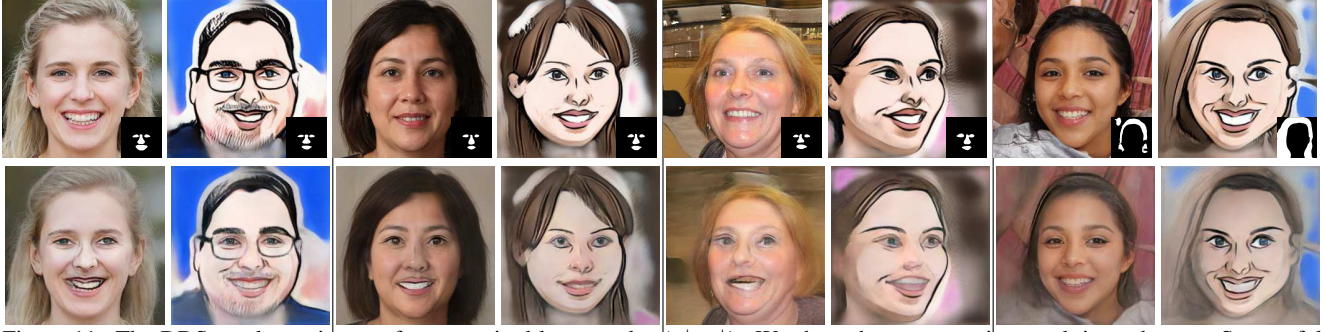


Figure 11. The DDS results on images from unpaired latent codes (z_1^*, z_2^*). We show the segmentation mask in each case. Successful examples (col 1-4) show that DDS can bridge small pose differences. Failures (col 5-8) are due to larger pose variations.

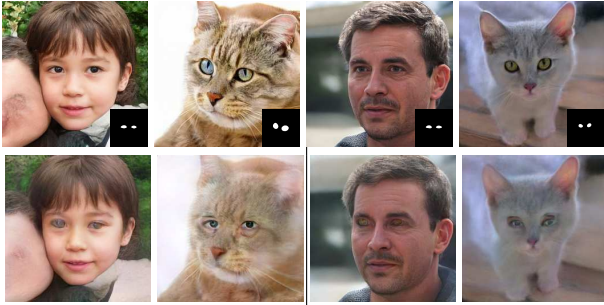


Figure 12. DDS results on unpaired human and cat images based on the segmentation model on eyes. First row shows the unpaired images with their segmentation mask. Second row shows human faces with the cat eyes and cats with the human eyes.

on LSUN Cars [36], and used 10 abandoned car images from the few-shotGAN implementation [4] for few-shot adaptation. Fig. 10 shows the qualitative results. The one-shot segmentation model is trained to segment the bonnet and lights of a car. Note that for both cars and horses, we use the *conv* layers of AlexNet [21] for the perceptual loss. This performs better than VGG [29] (see appendix).

5.3. Unpaired latents, different objects, limitations

Failures in the domain adaptation (Sec. 3.1) or in the segmentation masks (Sec. 3.2) will negatively impact DDS. In all previous results, we pass the same latent code z^* to the two generative models. Thanks to the adaptation algorithm, this guarantees the same pose in the corresponding images.

In Fig. 11, we show that DDS can work with unpaired latent codes with small pose variations. In all the domains and examples, DDS assumes the corresponding source and target images x_s, x_t are roughly aligned in pose. Fig. 11 also shows that DDS cannot recover from large pose variations in the unpaired images where the segmentation masks are not aligned.

We performed further unpaired experiments on cat and natural face images. We used the StyleGAN2 [19] model to generate cat images from [2] which is trained on LSUN Cats [36]. We used a segmentation model on eyes with the

purpose of generating dual-domain images of human face with the cat eyes and vice versa. Fig. 12 shows successful dual-domain synthesis. This demonstrates that DDS is not limited to examples when one object types, and can be used with two different many-shot styleGANs and are not adapted from one another, when masks are aligned. We present further results on this in appendix where both the source and target domains are from the few-shotGAN. The results in Fig. 12 demonstrate the potential of DDS on two unpaired examples from unrelated GANs and distinct segmentation models. However, successful examples are challenging to find due to the variation in pose when randomly sampling latent variables. This forms a limitation that we wish to address in future work, by utilising segmentation alignment approaches.

We also noted that in contrast with prior assumption [38], the choice of backbone for the perceptual loss impacts DDS results and is domain specific. We present further results on this in appendix. Further exploration is needed to ensure robustness to backbone choice in the perceptual loss.

6. Conclusion

We presented Dual-Domain Synthesis framework, an approach to generate visually convincing images that contain features from different domains. Our generated images are obtained by combining two correspondence images guided by part-based segmentation from a one-shot segmentation model. We demonstrated the approach on several image classes (face, car, horse, cat), multiple domains (natural, sketch, caricature, abandoned) and various part-based segmentation models (eye, nose, mouth, hair, car bonnet).

Extensive experiments illustrate the potential of the approach, both in terms of qualitative and quantitative assessment. While the majority of the results focus on paired latent codes and object (e.g. faces), the approach can be extended to distinct GANs trained independently and unpaired latent codes, assuming part-based alignment.

References

- [1] SSIM Implementation. "<https://github.com/Po-Hsun-Su/pytorch-ssim/>", 2017. [Online]. 5
- [2] StyleGAN2 Models. "<https://nvlabs-fi-cdn.nvidia.com/stylegan2/networks/>", 2019. [Online]. 7, 8
- [3] PyTorch Implementation of Analyzing and Improving the Image Quality of StyleGAN. "<https://github.com/rosinality/stylegan2-pytorch/>", 2020. [Online]. 4, 5
- [4] Implementation of Few-shot Image Generation via Cross-domain Correspondence. "<https://github.com/utkarshojha/few-shot-gan-adaptation/>", 2021. [Online]. 4, 5, 8
- [5] Implementation of Repurposing GANs for One-shot Semantic Part Segmentation. "<https://github.com/bryandlee/repurpose-gan/>", 2021. [Online]. 4
- [6] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2StyleGAN: How to Embed Images into the StyleGAN Latent Space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4432–4441, 2019. 2, 3, 11
- [7] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2StyleGAN++: How to Edit the Embedded Images? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8296–8305, 2020. 2, 11
- [8] Mohammed Al-Rawi, Dena Bazazian, and Ernest Valveny. Can Generative Adversarial Networks Teach Themselves Text Segmentation? In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019. 2
- [9] Edo Collins, Raja Bala, Bob Price, and Sabine Susstrunk. Editing in Style: Uncovering the Local Semantics of GANs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5771–5780, 2020. 2
- [10] Wolfgang Fuhl, David Geisler, Wolfgang Rosenstiel, and Enkelejda Kasneci. The Applicability of Cycle GANs for Pupil and Eyelid Segmentation, Data Generation and Image Refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019. 2
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. *Advances in Neural Information Processing Systems*, 27, 2014. 1
- [12] Shuyang Gu, Jianmin Bao, Hao Yang, Dong Chen, Fang Wen, and Lu Yuan. Mask-guided Portrait Editing with Conditional GANs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3436–3445, 2019. 1, 2
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 11
- [14] 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. *Advances in Neural Information Processing Systems*, 30, 2017. 5
- [15] Quan Huynh-Thu and Mohammed Ghanbari. Scope of Validity of PSNR in Image/Video Quality Assessment. *Electronics Letters*, 44(13):800–801, 2008. 6
- [16] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. SqueezeNetX: AlexNet-level Accuracy with 50x Fewer Parameters and < 0.5MB Model Size. *arXiv preprint arXiv:1602.07360*, 2016. 11
- [17] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual Losses for Real-Time Style Transfer and Super-Resolution. In *European Conference on Computer Vision*, pages 694–711. Springer, 2016. 2, 3
- [18] Tero Karras, Samuli Laine, and Timo Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 1, 2
- [19] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and Improving the Image Quality of StyleGAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 1, 2, 4, 5, 7, 8
- [20] Diederik P Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*, 2014. 4
- [21] Alex Krizhevsky. One Weird Trick for Parallelizing Convolutional Neural Networks. *arXiv preprint arXiv:1404.5997*, 2014. 8, 11, 12
- [22] 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 *Proceedings of the European Conference on Computer Vision*, pages 85–100, 2018. 2
- [23] Hongyu Liu, Ziyu Wan, Wei Huang, Yibing Song, Xintong Han, and Jing Liao. PD-GAN: Probabilistic Diverse GAN for Image Inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9371–9381, 2021. 2
- [24] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised Image-to-Image Translation Networks. In *Advances in Neural Information Processing Systems*, pages 700–708, 2017. 1
- [25] Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. Few-shot Image Generation via Cross-domain Correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10743–10752, 2021. 1, 2, 5, 7
- [26] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic Image Synthesis with Spatially-Adaptive Normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2337–2346, 2019. 1, 2
- [27] Esther Robb, Wen-Sheng Chu, Abhishek Kumar, and Jia-Bin Huang. Few-shot Adaptation of Generative Adversarial Networks. *arXiv preprint arXiv:2010.11943*, 2020. 2

- [28] Dimitrios Sakkos, Edmond SL Ho, and Hubert PH Shum. Illumination-aware Multi-Task GANs for Foreground Segmentation. *IEEE Access*, 7:10976–10986, 2019. 2
- [29] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3, 8, 11, 12
- [30] Konstantin Sofiiuk, Polina Popenova, and Anton Konushin. Foreground-aware semantic representations for image harmonization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1620–1629, 2021. 11, 12
- [31] Nontawat Tritrong, Pitchaporn Rewatbowornwong, and Supasorn Suwajanakorn. Repurposing GANs for One-shot Semantic Part Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4475–4485, 2021. 1, 2, 3
- [32] Xiaogang Wang and Xiaoou Tang. Face Photo-Sketch Synthesis and Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(11):1955–1967, 2008. 5
- [33] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. 5
- [34] Huikai Wu, Shuai Zheng, Junge Zhang, and Kaiqi Huang. GP-GAN: Towards realistic high-resolution image blending. In *Proceedings of the 27th ACM international conference on multimedia*, pages 2487–2495, 2019. 2
- [35] Mengyu Yang, David Rokeby, and Xavier Snelgrove. Mask-Guided Discovery of Semantic Manifolds in Generative Models. *4th Workshop on Machine Learning for Creativity and Design at NeurIPS 2020*, 2020. 2
- [36] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. LSUN: Construction of a Large-scale Image Dataset using Deep Learning with Humans in the Loop. *arXiv preprint arXiv: 1506.03365*, 2015. 7, 8
- [37] Lingzhi Zhang, Tarmily Wen, and Jianbo Shi. Deep image blending. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 231–240, 2020. 2, 11, 12
- [38] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 2, 8, 11
- [39] Yuxuan Zhang, Wenzheng Chen, Huan Ling, Jun Gao, Yanan Zhang, A. Torralba, and S. Fidler. Image GANs meet Differentiable Rendering for Inverse Graphics and Interpretable 3D Neural Rendering. *arXiv preprint arXiv: abs/2010.09125*, 2021. 2
- [40] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. SEAN: Image Synthesis with Semantic Region-Adaptive Normalization. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5103–5112, 2020. 1, 2