

Diffusion-based Image Translation with Label Guidance for Domain Adaptive Semantic Segmentation

Duo Peng¹ Ping Hu² Qihong Ke³ Jun Liu^{1,*}

¹Singapore University of Technology and Design ²Boston University

³Monash University

duo_peng@mymail.sutd.edu.sg, pinghu@bu.edu, qihong.ke@monash.edu, jun_liu@sutd.edu.sg

Abstract

Translating images from a source domain to a target domain for learning target models is one of the most common strategies in domain adaptive semantic segmentation (DASS). However, existing methods still struggle to preserve semantically-consistent local details between the original and translated images. In this work, we present an innovative approach that addresses this challenge by using source-domain labels as explicit guidance during image translation. Concretely, we formulate cross-domain image translation as a denoising diffusion process and utilize a novel Semantic Gradient Guidance (SGG) method to constrain the translation process, conditioning it on the pixel-wise source labels. Additionally, a Progressive Translation Learning (PTL) strategy is devised to enable the SGG method to work reliably across domains with large gaps. Extensive experiments demonstrate the superiority of our approach over state-of-the-art methods.

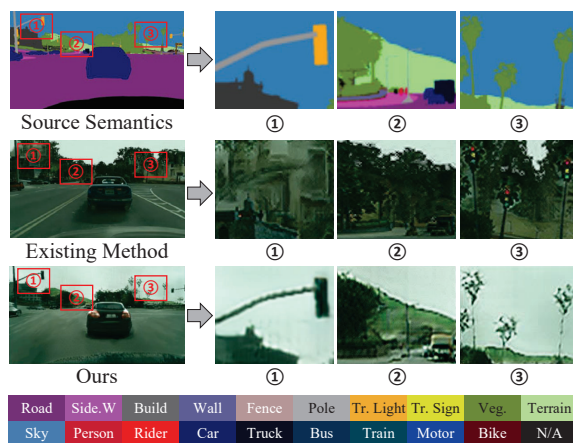
1. Introduction

Semantic segmentation is a fundamental computer vision task that aims to assign a semantic category to each pixel in an image. In the past decade, fully supervised methods [47, 59, 29, 84] have achieved remarkable progress in semantic segmentation. However, the progress is at the cost of a large number of dense pixel-level annotations, which are significantly labor-intensive and time-consuming to obtain [3, 47]. Facing this problem, one alternative is to leverage synthetic datasets, e.g., GTA5 [57] or SYNTHIA [60], where the simulated images and corresponding pixel-level labels can be generated by the computer engine freely. Although the annotation cost is much cheaper than manual labeling, segmentation models trained with such synthetic datasets often do not work well on real images due to the large domain discrepancy.

*Corresponding Author



(a) Architecture Comparison



(b) Results Comparison

Figure 1. (a) Architecture comparison between different image translation methods for DASS. Existing methods mainly translate images based on the input image. Our method further introduces the source semantic label to guide the translation process, making the translated content semantically consistent with the source label (as well as the source image). (b) Results comparison in detail. First row: the source semantic label. Second row: the translated image of an existing method [16], which is a state-of-the-art image translation method based on GAN. Third row: the translated image of our method. The translated image of the existing method shows unsatisfactory local details (e.g., in case ① of the second row, the traffic light and its surrounding sky are translated into building). In contrast, our method preserves details in a fine-grained manner.

The task of domain adaptive semantic segmentation (DASS) aims to address the domain discrepancy problem by transferring knowledge learned from a source domain to a target one with labels provided for the source domain

only. It is a very challenging problem as there can be a large discrepancy (e.g., between synthetic and real domains), resulting in great difficulties in knowledge transfer. In DASS, one type of methods [24, 7, 61, 5, 83, 12, 48, 68, 49, 66, 72] focuses on aligning the data distributions from different domains in the feature space, which obtains promising performance, while another type of approaches [86, 87, 40, 81, 38, 42] aims to generate pseudo labels for unlabeled target images, also making great achievements. Besides above approaches, image translation [85] is also an effective and important way to handle the DASS problem. It focuses on translating the labeled source images into the target domain, then using the translated images and source labels to train a target-domain segmentation model. In the past few years, lots of DASS studies [54, 23, 6, 44, 16] have been conducted based on image translation. This is because image translation allows for easy visual inspection of the adaptation method’s results, and the translated images can be saved as a dataset for the future training of any model, which is convenient. Existing image translation methods mainly adopt generative adversarial networks (GANs) [19] to translate images across a large domain gap. However, due to the intractable adversarial training, GAN models can be notoriously hard to train and prone to show unstable performance [62, 53]. Furthermore, many studies [51, 11, 55] have also shown that GAN-based methods struggle with preserving local details, leading to semantic inconsistencies between the translated images and corresponding labels, as shown in the second row of Fig. 1 (b). Training the segmentation model on such mismatched image and label pairs will cause sub-optimal domain adaptation performance. Facing these inherent limitations, the development of GAN-based image translation methods is increasingly encountering a bottleneck. Yet, there is currently little research exploring alternative approaches to improve image translation for DASS beyond GANs.

Denosing Diffusion Probabilistic Models (DDPMs), also known as diffusion models, have recently emerged as a promising alternative to GANs for various tasks, such as image generation [11, 69], image restoration [73, 79], and image editing [1, 52], etc. Diffusion models have been shown to have a more stable training process [11, 55, 22]. Inspired by its strong capability, we propose to exploit the diffusion model for image translation in DASS. Specifically, based on the diffusion technique, we propose a label-guided image translation framework to preserve local details. We observe that existing image translation methods [23, 54, 6, 44, 41] mainly focus on training a translation model with image data alone, while neglecting the utilization of source-domain labels during translation, as shown in Fig. 1 (a). Since source labels can explicitly indicate the semantic category of each pixel, introducing them to guide the translation process should improve the capability

of the translation model to preserve details. A straightforward idea of incorporating the source labels is to directly train a conditional image translation model, where the translated results are semantically conditioned on the pixel-wise source labels. However, it is non-trivial since the source labels and target images are not paired, which cannot support the standard conditional training. Recent advances [11, 55, 58] have shown that a pre-trained unconditional diffusion model can become conditional with the help of the gradient guidance [11]. The gradient guidance method can guide the pre-trained unconditional diffusion model to provide desired results by directly affecting the inference process (without any training). In light of this, we propose to first train an unconditional diffusion-based image translation model, and then apply gradient guidance to the image translation process, making it conditioned on source labels. However, we still face two challenges: (1) Traditional gradient guidance methods generally focus on guiding the diffusion model based on image-level labels (i.e., classification labels), while the DASS problem requires the image translation to be conditioned on pixel-level labels (i.e., segmentation labels). (2) The gradient guidance methods typically work within a single domain, whereas the DASS task requires it to guide the image translation across different domains. To address the first challenge, we propose a novel Semantic Gradient Guidance method (SGG), which can precisely guide the image translation based on fine-grained segmentation labels. To tackle the second challenge, we carefully design a Progressive Translation Learning strategy (PTL), which is proposed to drive the SGG method to reliably work across a large domain gap. With the help of SGG and PTL, our framework effectively handles the image translation for DASS in a fine-grained manner, which is shown in the third row of Fig. 1 (b).

In summary, our contributions are three-fold. (i) We propose a novel diffusion-based image translation framework with the guidance of pixel-wise labels, which achieves image translation with fine-grained semantic preservation. To the best of our knowledge, this is the first study to exploit the diffusion model for DASS. (ii) We devise a Semantic Gradient Guidance (SGG) scheme accompanied with a Progressive Translation Learning (PTL) strategy to guide the translation process across a large domain gap in DASS. (iii) Our method achieves the new state-of-the-art performance on two widely used DASS settings. Remarkably, our method outperforms existing GAN-based image translation methods significantly on various backbones and settings.

2. Related Work

Domain adaptive semantic segmentation (DASS) is a task that aims to improve the semantic segmentation model’s adaptation performance in new target scenarios without access to the target labels. Current DASS meth-

ods [24, 61, 40, 8, 23] have pursued three major directions to bridge the domain gap. The first direction aims to align data distributions in the latent feature space. With the advent of adversarial learning [14], many methods [24, 7, 56, 61, 5, 37, 83, 12, 48, 68, 49, 66, 72] adopt the adversarial learning to align feature distributions by fooling a domain discriminator. However, the model based on adversarial learning can be hard to train [82, 62, 53], as it requires carefully balancing the capacity between the feature encoder (i.e., generator) and the domain discriminator, which is intractable. The second direction considers exploring the network’s potential of self-training on the target domain [81, 50, 40, 38, 71, 42, 86, 87]. Self-training methods generally use the source-trained segmentation model to predict segmentation results for target-domain images. They typically select the results with high confidence scores as pseudo labels to fine-tune the segmentation model. These works mainly focus on how to decide the threshold of confidence scores for selecting highly reliable pseudo labels. However, self-training is likely to introduce inevitable label noise [44, 30, 76], as we cannot fully guarantee that the selected pseudo labels are correct especially when the domain gap is large. The third direction is inspired by the image translation technique [85], aiming to transfer the style of source labeled images to that of the target domain [8, 54, 23, 6, 44, 16]. Inspired by the advances of the generative adversarial networks (GANs), they mostly adopt GANs as image translation models. While remarkable progress has been achieved, many studies [51, 11, 55] have indicated that GAN-based image translation methods tend to show limited capacity for preserving semantic details. Unlike existing methods that handle image translation based on the input source image, our method enables the model to translate images further conditioned on the source labels, which enables fine-grained semantic preservation.

Denosing Diffusion Probabilistic Models (DDPMs), also namely diffusion models, are a class of generative models inspired by the nonequilibrium thermodynamics. Diffusion models [64, 22] generally learn a denoising model to gradually denoise from an original common distribution, e.g., Gaussian noise, to a specific data distribution. It is first proposed by Sohl-Dickstein et al. [64], and has recently attracted much attention in a wide range of research fields including computer vision [11, 69, 13, 1, 52, 9, 9, 17, 4, 20], nature language processing [2, 18, 43], audio processing [28, 28, 34, 67, 74] and multi-modality [33, 45, 27, 70]. Based on the diffusion technique, Dhariwal et al. proposed a gradient guidance method [11] to enable pre-trained unconditional diffusion models to generate images based on a specified class label. Although effective, the gradient guidance method generally works with image-level labels and is limited to a single domain, making it challenging to be adopted for DASS. To address these issues, in this paper,

we propose a novel Semantic Gradient Guidance (SGG) method and a Progressive Translation Learning (PTL) strategy, enabling the utilization of pixel-level labels for label-guided cross-domain image translation.

3. Revisiting Diffusion Models

As our method is proposed based on the Denosing Diffusion Probabilistic Model (DDPM) [64, 22], here, we briefly revisit DDPM. DDPM is a generative model that converts data from noise to clean images by gradually denoising. A standard DDPM generally contains two processes: a diffusion process and a reverse process. The diffusion process contains no learnable parameters and it just used to generate the training data for diffusion model’s learning. Specifically, in the diffusion process, a slight amount of Gaussian noise is repeatedly added to the clean image x_0 , making it gradually corrupted into a noisy image x_T after T -step noise-adding operations. A single step of diffusion process (from x_{t-1} to x_t) is represented as follows:

$$q(x_t|x_{t-1}) = \mathcal{N}(\sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I}), \quad (1)$$

where β_t is the predefined variance schedule. Given a clean image x_0 , we can obtain x_T by repeating the above diffusion step from $t = 1$ to $t = T$. Particularly, an arbitrary image x_t in the diffusion process can also be directly calculated as:

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad (2)$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ and $\epsilon \sim \mathcal{N}(0, I)$. Based on this equation, we can generate images from x_0 to x_T for training the diffusion model.

The reverse process is opposite to the diffusion process, which aims to train the diffusion model to gradually convert data from the noise x_T to the clean image x_0 . One step of the reverse process (from x_t to x_{t-1}) can be formulated as:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(\mu_\theta(x_t, t), \Sigma_\theta(x_t, t)). \quad (3)$$

Given the noisy image x_T , we can denoise it to the clean image x_0 by repeating the above equation from $t = T$ to $t = 1$. In Eq. 3, the covariance Σ_θ is generally fixed as a predefined value, while the mean μ_θ is estimated by the diffusion model. Specifically, the diffusion model estimates μ_θ via a parameterized noise estimator $\epsilon_\theta(x_t, t)$ which is usually a U-Net [59]. Given x_t as the input, the diffusion model will output the value of μ_θ as follows:

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{1 - \beta_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right), \quad (4)$$

For training the diffusion model, the images generated in the diffusion process are utilized to train the noise estimator




Figure 2. Illustration of the Diffusion-based Image Translation Baseline. Starting from a source image x_0 , the diffusion process disturbs it into a noisy image x_n , which is then denoised into the target image x_0^ϕ using the reverse process performed by a target-pre-trained diffusion model.

ϵ_θ to predict the noise added to the image, which can be formulated as:

$$L_{DM}(x) = \|\epsilon_\theta(x_t, t) - \epsilon\|^2, \quad (5)$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ is the noise added to the image x_t .

4. Method

We address the problem of domain adaptive semantic segmentation (DASS), where we are provided the source images \mathcal{X}^S with corresponding source labels \mathcal{Y}^S and the target images \mathcal{X}^T (without target labels). Our goal is to train a label-guided image translation model to translate source images into target-like ones. Then we can use the translated data to train a target-domain segmentation model. To achieve this goal, we first train a basic image translation model based on the diffusion technique, which is the baseline of our approach (Sec. 4.1). Then, we impose our SGG method to the baseline model, making it conditioned on the pixel-wise source labels (Sec. 4.2). We further propose a PTL training strategy to enable our SGG to work across a large domain gap (Sec. 4.3). Finally, we detail the training and inference procedures (Sec. 4.4).

4.1. Diffusion-based Image Translation Baseline

We start with the baseline of our framework, which is a basic image translation model based on diffusion models. Before translation, we use the (unlabeled) target domain images \mathcal{X}^T to pre-train a standard diffusion model via Eq. 5. As the diffusion model is essentially a denoising model, we use the noise as an intermediary to link the two large-gap domains. As shown in Fig. 2, we first leverage the diffusion process to add n -step noise onto the source image x_0 ($x_0 \in \mathcal{X}^S$), making it corrupted into x_n , and then use the pre-trained diffusion model to reverse (denoise) it for n steps, obtaining x_0^ϕ . Since the diffusion model is pre-trained on the target domain, the reverse (denoising) process will convert the noise into target-domain content, making the denoised image x_0^ϕ shown an appearance closer to the target-domain images than the original source-domain image x_0 . Note that the diffusion model

solely trained on the target domain is able to denoise the noisy image from a different source domain towards the target domain, because the diffusion model’s training objective is derived from the variational bound on the negative log-likelihood $\mathbb{E}[-\log p_\theta(\mathcal{X}^T)]$, indicating it learns to generate target-domain data from various distributions. This property of diffusion model has been verified in [65] and has been widely used in many studies [52, 15, 9].

4.2. Semantic Gradient Guidance (SGG)

The baseline model presented above handles image translation in an unconditional manner, where denoising from much noise however will freely generate new content that may not be strictly consistent with the original image. Inspired by the gradient guidance [11], we here propose a novel SGG method to enable the unconditional baseline to be conditioned on source labels.

Our SGG consists of two guidance modules, namely the Local Class-regional Guidance (LCG) and the Global Scene-harmonious Guidance (GSG). They are proposed to affect the reverse process of the baseline model, aiming to guide the denoising process to generate content semantically consistent with source labels. Specifically, the LCG is proposed to preserve local details, while the GSG is designed to enhance the global harmony of the denoised images. As shown in Fig. 3 (a), we alternately apply the LCG and GSG modules at each reverse step, aiming to provide both local-precision and global-harmony guidance.

(1) Local Class-regional Guidance (LCG)

The key idea of LCG is to use a pixel classifier to calculate loss to measure if the generated (denoised) pixel has the same semantic as the source semantic label, and then utilize the loss gradient as the guidance signal to guide the pixel to generate same-semantic content. Assuming that we already have a pre-trained segmentation model g (i.e., the pixel classifier) that can perform pixel classification well over domains (note that we will address this assumption in Sec. 4.3), based on g , LCG can carry out guidance as follows.

Here, we use x_k^ϕ to denote the input of the LCG module, and x_{k-1}^ϕ as its output, where $k \in \{n, n-2, n-4, \dots\}$. As shown in Fig. 3 (b), given a source label map $y \in \mathcal{Y}^S$, we can leverage y to obtain a binary mask m_c that indicates the region of the c -th class in the source image ($c \in \mathcal{C}$ and \mathcal{C} denotes all classes in the source dataset). Given the input x_k^ϕ , we can multiply x_k^ϕ with m_c to obtain the pixels inside the c -th class region, which can be fed into the pre-trained segmentation model g to predict the label probabilities of the pixels. The output can be used to calculate the cross-entropy loss with the source label, aiming to use the loss gradient to guide the generation of pixels in the c -th class region. As shown in Fig. 3 (b), we can use classifier g to

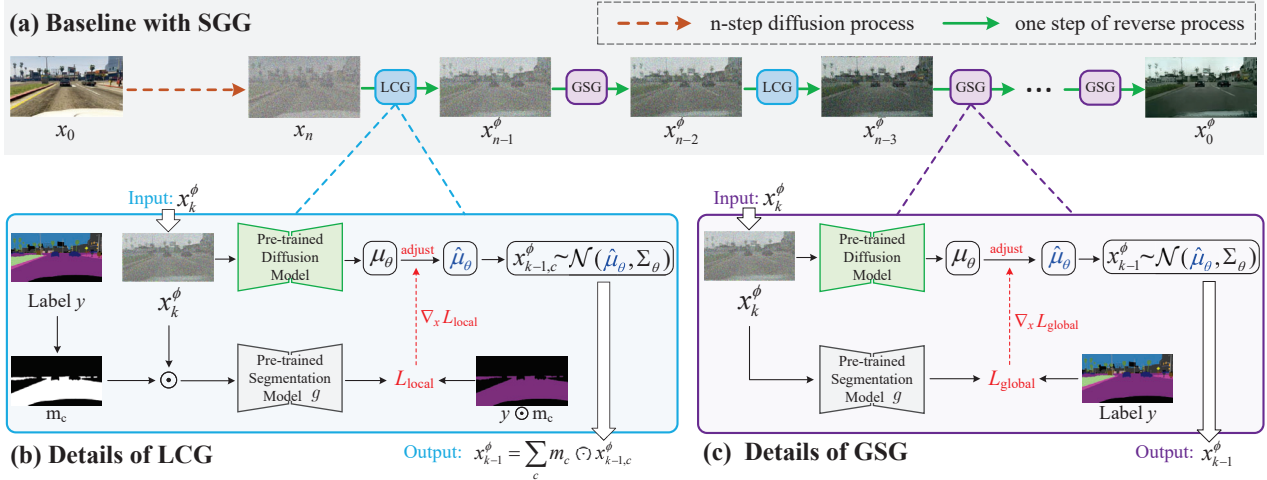


Figure 3. (a) The overall architecture of our baseline model with Semantic Gradient Guidance (SGG). SGG alternately adopts the Local Class-regional Guidance (LCG) and the Global Scene-harmonious Guidance (GSG) at each reverse step of the baseline model. (b) In LCG, we first obtain the loss gradient for class region m_c , and then use the gradient to adjust the diffusion model’s output μ_θ for generating label-related content. For clarity, we only show the workflow of one class region. We obtain the guided results for every class region and finally combine them together via $x_{k-1}^\phi = \sum_c m_c \odot x_{k-1,c}^\phi$. (c) In GSG, we calculate the gradient for the whole image and then use the gradient to adjust μ_θ for a global-harmony appearance.

compute the cross-entropy loss L_{local} for pixels in the local c -th class region, which is formulated as:

$$L_{local}(x_k^\phi, y, c) = L_{ce}(g(x_k^\phi \odot m_c), y \odot m_c), \quad (6)$$

where L_{ce} is the cross-entropy loss, \odot denotes the element-wise multiplication and $y \odot m_c$ is the source label of the class c .

In the standard reverse process, given the noisy image x_k^ϕ , the diffusion model denoises it by estimating the mean value μ_θ (Eq. 4) and then using it to obtain the denoised image x_{k-1}^ϕ via sampling (Eq. 3). Therefore, the mean μ_θ will determine the content of the denoised image. Based on this, we guide the reverse process by using the computed loss gradient $\nabla_{x_k^\phi} L_{local}$ to adjust μ_θ , as shown in Fig. 3 (b). As the loss gradient is computed based on the source label, it can enable the generated content to correspond to the label class. We follow the previous gradient guidance method [11] to formulate the adjustment operation as:

$$\hat{\mu}_\theta(x_k^\phi, k, c) = \mu_\theta(x_k^\phi, k) + \lambda \Sigma_\theta \nabla_{x_k^\phi} L_{local}(x_k^\phi, y, c), \quad (7)$$

where μ_θ is the original output of the pre-trained diffusion model, Σ_θ is a fixed value already defined in Eq. 3 and λ is a hyper-parameter which controls the guidance scale.

Next, we adopt the adjusted mean value $\hat{\mu}_\theta$ to sample the denoised image:

$$x_{k-1,c}^\phi \sim \mathcal{N}(\hat{\mu}_\theta(x_k^\phi, k, c), \Sigma_\theta). \quad (8)$$

With the above process, the guided denoised image $x_{k-1,c}^\phi$ can preserve the semantics for pixels in the c -th class

region. We parallelly perform the above operations for different c , and combine different guided regions into a complete image x_{k-1}^ϕ :

$$x_{k-1}^\phi = \sum_c m_c \odot x_{k-1,c}^\phi, \quad (9)$$

where x_{k-1}^ϕ is the output of the LCG module.

(2) Global Scene-harmonious Guidance (GSG)

The LCG is performed locally, achieving a per-category guidance for image translation. As a complement, we also propose a global guidance module to enhance the global harmony of the translated image, which further boosts the translation quality.

As shown in Fig. 3 (c), our proposed GSG also follows the similar operation flow as LCG, i.e., “computing loss gradient \rightarrow adjusting mean value \rightarrow sampling new image”. Given x_k^ϕ ($k \in \{n-1, n-3, n-5, \dots\}$) as the input of the GSG module, the loss function is computed based on the whole image to achieve the global guidance:

$$L_{global}(x_k^\phi, y) = L_{ce}(g(x_k^\phi), y). \quad (10)$$

In SGG, we apply the LCG or GSG at each reverse (denoising) step of the baseline model. By doing so, we can gradually apply the guidance in a step-wise manner, which is suitable for the step-wise translation process of the baseline model. To achieve both local-precise and global-harmony guidance, we incorporate an equal number of LCG and GSG modules into our approach. Specifically, since the baseline model has n reverse steps, we use $\frac{n}{2}$ LCG modules

and $\frac{n}{2}$ GSG modules. As for the arrangement of LCG and GSG modules, we design several plausible options: (1) first using $\frac{n}{2}$ LCG modules and then $\frac{n}{2}$ GSG modules, (2) first using $\frac{n}{2}$ GSG modules and then $\frac{n}{2}$ LCG modules, (3) alternating between LCG and GSG, and (4) randomly mixing the two types of modules. We observe that all options perform at the same level. Since the alternating option shows a slightly better performance than others, we adopt the alternating scheme for SGG.

4.3. Progressive Translation Learning (PTL)

The proposed SGG method is based on the assumption that we already have a segmentation model trained over domains. That is, if we want to guide the translation towards the target domain, we need to have a segmentation model that is pre-trained on the target domain. However, in DASS, we only have labeled images in the source domain for training a source-domain segmentation network. To enable the SGG to reliably work for translation towards the target domain with the source-domain segmentation model, we propose a Progressive Translation Learning (PTL) strategy.

As mentioned before, our baseline model can directly translate images by first diffusing for n steps and then reversing for n steps (see Fig. 2). We formulate this process as: $\mathcal{X}^n = \phi^n(\mathcal{X}^S)$, where \mathcal{X}^S denotes all source-domain images, ϕ^n denotes the baseline model that handles n -step diffusion and reverse process, \mathcal{X}^n denotes the translated output images which is related to the parameter n . As the diffusion model can denoise the image towards the target domain, the parameter n essentially controls the amount of target-domain information added to the source image. Inspired by [15], we can flexibly change the domain of translated images by varying the value of n . When n is small, we can obtain translated images \mathcal{X}^n with minor differences to source-domain images \mathcal{X}^S . Conversely, when n is large enough, the translated images \mathcal{X}^n can completely approach the target domain. By walking through n from 1 to N , our baseline model can generate a set of image domains $\{\mathcal{X}^n | n = 1, 2, \dots, N\}$, where \mathcal{X}^N approximately follows the target-domain distribution. Since each step of diffusion/reverse operation only adds/removes a very small amount of noise, the translated images from neighboring domains, i.e., \mathcal{X}^n and \mathcal{X}^{n+1} , show very slight differences.

Specifically, we start with the source-domain segmentation model g_0 , which is well-trained on source images \mathcal{X}^S with source labels \mathcal{Y}^S . As the domain gap between two adjacent domains is very small, the segmentation network trained on source domain \mathcal{X}^S can also work well for the adjacent domain \mathcal{X}^1 . Based on this property, we can use g_0 to carry out SGG guidance for image translation of the next domain \mathcal{X}^1 . We formulate this operation in a general representation, i.e., given g_n which is well-trained on \mathcal{X}_n , we

can guide the translation of \mathcal{X}_{n+1} as follows:

$$\mathcal{X}^{n+1} = \text{SGG}(\phi^{n+1}(\mathcal{X}^S), g_n), \quad (11)$$

where $\text{SGG}(\phi^{n+1}, g_n)$ denotes the SGG-guided baseline model with $(n + 1)$ -step diffusion and reverse process, which is equipped with g_n .

Under the g_0 -based SGG guidance, the translated images \mathcal{X}^1 have fine-grained semantic consistency with source labels \mathcal{Y}^S . Then, we can use the guided images \mathcal{X}^1 to fine-tune the segmentation model g_0 , making it further fit with this domain. This operation is also formulated in a general representation, i.e., given the guided images \mathcal{X}^{n+1} , we can fine-tune the segmentation model g_n as follows:

$$L_{ft} = L_{ce}(g_n(\mathcal{X}^{n+1}), \mathcal{Y}^S) + L_{ce}(g_n(\mathcal{X}^S), \mathcal{Y}^S). \quad (12)$$

We name the fine-tuned segmentation model as g_{n+1} . The former loss component aims to use the guided images \mathcal{X}^{n+1} to adapt the segmentation model to the new domain. The latter component is to train the segmentation model on the source domain images \mathcal{X}^S . We empirically observe combining both that augments the model’s learning scope leads to slightly better results than using the former loss only.

After fine-tuning the segmentation model g_0 on \mathcal{X}^1 , we can perform SGG-guided image translation again using the fine-tuned model (namely g_1), obtaining guided images \mathcal{X}^2 . Then we can use \mathcal{X}^2 to further fine-tune the segmentation model g_1 for adapting to a new domain. In this manner, we can perform Eq. 11 and Eq. 12 iteratively with n increasing from 1 to N , finally obtaining the translated target-domain images \mathcal{X}^N and target-fine-tuned segmentation model g_N . We use the segmentation model g_N for test-time inference. Throughout the whole process of PTL, the segmentation model always performs SGG guidance at the adjacent domain with little domain discrepancy, which is the key to enabling the SGG to work across a large domain gap.

4.4. Training and Inference

Training. First, we train a standard diffusion model with the unlabeled target-domain images \mathcal{X}^T via Eq. 5. After that, we keep the pre-trained diffusion model frozen and use it to denoise the source-diffused image, which forms our basic image translation model, namely baseline model ϕ^n (Sec. 4.1). Our baseline model ϕ^n can translate source images \mathcal{X}^S into the images \mathcal{X}^n . The model ϕ^n is flexible with a variable parameter n , which can control the domain of the translated images \mathcal{X}^n . When $n = 1$, the translated images are extremely close to the source domain. When $n = N$, the translated images approximately approach the target domain. We use the source-domain labeled data $\{\mathcal{X}^S, \mathcal{Y}^S\}$ to train a segmentation model g_0 . Then, we set $n = 1$ for the baseline model ϕ^n and use it to translate source images \mathcal{X}^S into the adjacent-domain images \mathcal{X}^1 . During the

translation process, we use the segmentation model g_0 to carry out SGG guidance (Eq.11), obtaining the guided images \mathcal{X}^1 . Then, we use the guided images and source labels $\{\mathcal{X}^1, \mathcal{Y}^S\}$ to fine-tune the segmentation model g_0 (Eq.12), obtaining fine-tuned segmentation model g_1 . We set $n = 2$ for the baseline model ϕ^n , and use g_1 to carry out SGG guidance, obtaining the guided images \mathcal{X}^2 . Next, we can use $\{\mathcal{X}^2, \mathcal{Y}^S\}$ to execute the fine-tuning again. In this way, we alternate the SGG guidance and the fine-tuning of segmentation model with n growing from 1 to N (i.e., PTL strategy). Finally, we obtain the translated target-domain images \mathcal{X}^N and the target-fine-tuned segmentation model g_N . **Inference.** After training, we directly use the segmentation model g_N for inference.

5. Experiment

5.1. Implementation Details

Datasets. As a common practice in DASS, we evaluate our framework on two standard benchmarks (GTA5 [57] \rightarrow Cityscapes [10], and SYNTHIA [60] \rightarrow Cityscapes [10]). GTA5 and SYNTHIA provide 24,996 and 9400 labeled images, respectively. Cityscapes consists of 2975 and 500 images for training and validation respectively. Following the standard protocol [24], we report the mean intersection over Union (mIoU) on 19 classes for GTA5 \rightarrow Cityscapes and 16 classes for Synthia \rightarrow Cityscapes.

Network Architecture. To make a fair and comprehensive comparison, we test three typical types of networks as the segmentation model g : (1) DeepLab-V2 [3] architecture with ResNet-101 [21] backbone; (2) FCN-8s [47] architecture with VGG16 [63] backbone; (3) DAFormer [25] architecture with SegFormer [75] backbone. All of them are initialized with the network pre-trained on ImageNet [36]. For the diffusion model, we follow the previous work [22] to adopt a U-Net [59] architecture with Wide-ResNet [80] backbone. The diffusion model is trained from scratch without loading any pre-trained parameters.

Parameter Setting. Following DDPM [22], we train the diffusion model with a batch size of 4, using the Adam [35] as the optimizer with learning rate as 2×10^{-5} and momentum as 0.9 and 0.99. Specifically, for the hyperparameters of diffusion model, we still follow [22] to set β_t from $\beta_1 = 10^{-4}$ to $\beta_T = 0.02$ linearly and $T = 1000$. The covariance Σ_θ is fixed and defined as $\Sigma_\theta = \beta_t \mathbf{I}$. We train the segmentation model g by using the SGD optimizer [32] with a batch size of 2, a learning rate of 2.5×10^{-4} and a momentum of 0.9. In SGG, we set $\lambda = 80$. In PTL, we set $N = 600$. We pre-train the segmentation model on the source domain for 20,000 iterations and fine-tune it on other intermediate domains for 300 iterations each. Since the discrepancy between adjacent domains is minor, fine-tuning for 300 iterations is enough to effectively adapt to a

Table 1. Performance comparison in terms of mIoU (%). Three backbone networks, ResNet-101 (Res), VGG16 (Vgg) and SegFormer (Seg), are used in our study. The results of the task ‘‘GTA5 \rightarrow Cityscapes’’ (‘‘SYNTHIA \rightarrow Cityscapes’’) are averaged over 19 (16) classes. * denotes previous image translation methods.

Method	Venue	GTA5 \rightarrow Cityscapes			SYN. \rightarrow Cityscapes		
		Res	Vgg	Seg	Res	Vgg	Seg
Source only	-	37.3	27.1	45.4	33.5	22.9	40.7
CyCADA [23]*	ICML’18	42.7	35.4	-	-	-	-
IIT [54]*	CVPR’18	-	35.7	-	-	-	-
CrDoCo [6]*	CVPR’19	-	38.1	-	-	38.2	-
BDL [44]*	CVPR’19	48.5	41.3	-	-	33.2	-
CLAN [49]	CVPR’19	43.2	36.6	-	-	-	-
Intra [56]	CVPR’20	45.3	34.1	-	41.7	33.8	-
SUIT [41]*	PR’20	45.3	40.6	-	40.9	38.1	-
LDR [76]	ECCV’20	49.5	43.6	-	-	41.1	-
CCM [40]	ECCV’20	49.9	-	-	45.2	-	-
ProDA [81]	CVPR’21	57.5	-	-	55.5	-	-
CTF [50]	CVPR’21	56.1	-	-	48.2	-	-
CADA [77]*	WACV’21	49.2	44.9	-	-	40.8	-
UPLR [71]	ICCV’21	52.6	-	-	48.0	-	-
DPL [8]	ICCV’21	53.3	46.5	-	47.0	43.0	-
BAPA [46]	ICCV’21	57.4	-	-	53.3	-	-
CIR [16]*	AAAI’21	49.1	-	-	43.9	-	-
CRAM [82]	TITS’22	48.6	41.8	-	45.8	38.7	-
ADAS [39]*	CVPR’22	45.8	-	-	-	-	-
CPSL [42]	CVPR’22	55.7	-	-	54.4	-	-
FDA [78]	CVPR’22	50.5	42.2	-	-	40.0	-
DAP [30]	CVPR’22	59.8	-	-	59.8	-	-
DAFormer [25]	CVPR’22	-	-	68.3	-	-	60.9
HRDA [26]	ECCV’22	-	-	73.8	-	-	65.8
ProCA [31]	ECCV’22	56.3	-	-	53.0	-	-
Bi-CL [38]	ECCV’22	57.1	-	-	55.6	-	-
Deco-Net [37]	ECCV’22	59.1	-	-	57.0	-	-
Ours	ICCV’23	61.9	48.1	75.3	61.0	44.2	66.5

new domain. As the SGG guidance needs the segmentation model g to classify for noisy and masked images, during the training of the segmentation model, we also generate noisy and masked images for data augmentation to train a robust segmentation model.

5.2. Comparative Studies

We compared the results of our method with the state-of-the-art methods on two DASS tasks, i.e., ‘‘GTA5 \rightarrow Cityscapes’’ and ‘‘SYNTHIA \rightarrow Cityscapes’’, using three different backbone networks: ResNet-101 [21], VGG16 [63] and SegFormer [75]. To ensure reliability, we run our model five times and report the average accuracy. The comparative results are reported in Tab. 1. Overall, our method consistently gains the best performance across all settings and backbones. Remarkably, when compared to existing image translation methods (marked with *), our method brings 3.2% \sim 20.1% improvement on all backbones and settings, which demonstrates the superiority of our diffusion-based approach.

5.3. Ablation Studies

Effect of Baseline model. In Sec. 4.1, we propose a Diffusion-based Image Translation Baseline, i.e., the baseline model ϕ^n . To study the effectiveness of the baseline model ϕ^n , we set $n = N$ to directly translate source images




Figure 4. The translated images of our approach (including images from intermediate domains). n denotes the serial number of intermediate domains. We can see that the semantic knowledge is transferred smoothly and precisely.

Table 2. Ablation studies of different components on the task “GTA5 \rightarrow Cityscapes” (G \rightarrow C) and “SYNTHIA \rightarrow Cityscapes” (S \rightarrow C) with backbone Res-101.

	Model	G \rightarrow C	S \rightarrow C
(a)	Baseline Model	48.3	45.1
(b)	Ours w/o SGG	55.4	52.7
(c)	Ours w/o PTL	58.3	56.5
(d)	Ours w/o LCG	59.8	59.1
(e)	Ours w/o GSG	60.5	59.4
(f)	Ours (full)	61.9	61.0

into the target domain. After translation, we use the translated images and source images, along with the source labels, to train a segmentation model for inference. As shown in Tab. 2 (a), our baseline model achieves 48.3% and 45.1% mIoU on two DASS tasks, respectively. This performance is comparable to the results of previous image translation methods (marked with * in Tab. 1), which demonstrates the effectiveness of our baseline model.

Effect of SGG. To demonstrate the effectiveness of SGG, during the PTL training, we remove the SGG guidance and directly use the unguided images of each intermediate domain to fine-tune the segmentation model. By comparing the results in Tab. 2 (b) and (f), we can see that when adopting SGG, our framework shows a significant improvement (more than 6%) on both tasks, clearly demonstrating the effectiveness of SGG in preserving semantics.

Effect of PTL. The key design of PTL is to generate intermediate domains to decompose the large gap into small ones. To investigate the effectiveness of PTL, we construct an ablation experiment by removing the generation of intermediate domains, i.e., we set $n = N$ for the baseline model ϕ^n and use it to translate source images into the target domain, i.e., $\mathcal{X}^N = \phi^N(\mathcal{X}^S)$. During the translation, we directly use the source-trained segmentation model to execute SGG guidance to guide the translation process. Finally, we use the guided target-domain images to fine-tune the segmentation model. From Tab. 2 (c) and (f), we can see that our framework without PTL strategy consistently incurs an obvious performance drop on both DASS tasks, which demonstrates the effectiveness of PTL strategy. We show in Fig. 4 some translated image examples and their corresponding intermediate domains. We can see that the

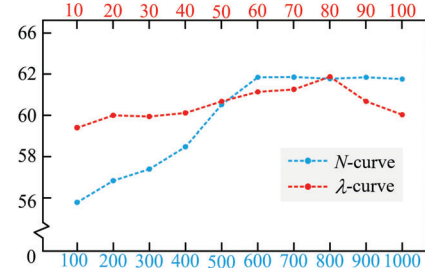


Figure 5. Parameter analysis on N and λ .

image domain is gradually transferred, and all images contain semantics that are well preserved in a fine-grained manner, which demonstrates that our PTL can enable the SGG to reliably work on each intermediate domain.

Effect of LCG and GSG. In order to evaluate the impact of the LCG and GSG modules, we conduct ablation experiments by removing each module respectively. As shown in Tab. 2 (d) and (e), we find that the model’s performance without LCG (or GSG) degrades, which demonstrates their usefulness. Furthermore, to comprehensively demonstrate the effect of LCG and GSG modules, we provide some qualitative ablation results in *Supplementary*.

Effect of the second loss component of fine-tuning loss in PTL. As mentioned in Sec. 4.3, in our PTL strategy, we propose a fine-tuning loss L_{ft} to progressively fine-tune the segmentation model towards the target domain, which is formulated as follows:

$$L_{ft} = \underbrace{L_{ce}(g_n(\mathcal{X}^{n+1}), \mathcal{Y}^S)}_{L_{ada}} + \underbrace{L_{ce}(g_n(\mathcal{X}^S), \mathcal{Y}^S)}_{L_{src}}. \quad (13)$$

The fine-tuning loss has two components. The former loss component L_{ada} aims to adapt the segmentation model to a new domain, which is closer to the target domain. The later component L_{src} is to train the segmentation model on the source domain. As shown in Tab. 3, compared to using the former loss L_{ada} only, the combination of both components can augment the model’s learning scope, leading to better results (i.e., a gain of 0.4% and 0.3%).

Parameter Analysis. For parameters β , T and Σ_θ , we follow the previous work DDPM [22] to set their values. In this paper, we only need to study the impact of N (the num-




Figure 6. Qualitative semantic segmentation results of our framework. It can be seen that the segmentation performance is gradually improved by adding our proposed components one by one. Best viewed in color.

Table 3. Ablation study on the loss component L_{src} in fine-tuning loss L_{ft} .

	Model	G→C	S→C
(a)	L_{ada}	61.5	60.7
(b)	$L_{ada} + L_{src}$	61.9	61.0

Table 4. Comparison of training and inference time.

Methods	Training Time	Inference Time	Acc.
BDL [44]	1.8 days	58.40 ms	48.5
ProCA [31]	1.9 days	57.48 ms	56.3
Bi-CL [38]	1.7 days	57.91 ms	57.1
Deco-Net [37]	1.9 days	59.67 ms	59.1
Ours	2.3 days	57.68 ms	61.9

ber of domains in PTL) and λ (the guidance scale parameter in SGG). As for N , we desire the number of domains as large as possible, as we want to smoothly bridge the domain gap. On the other hand, we need to avoid inducing too much training time, i.e., the number of domains should be limited. We present the ablation results of N in Fig. 5 (blue curve). We find the accuracy begins to level off when $N > 600$. We thus set the optimal value of N to 600 with the consideration of saving training time. As for λ , we study its impact with different values and show the results in Fig. 5 (red curve). We can see that the best option of λ is $\lambda = 80$.

Training & Inference Time. We compare our method with other methods in terms of training and inference time on the task “GTA5→Cityscapes”. Although our model is trained across N intermediate domains ($N = 600$), the training time does not increase much, as we only need to fine-tune the model for several iterations on each intermediate domain (more details in Sec. 5.1). As shown in Tab. 4, compared to recent state-of-the-art models, though our method achieves a significant performance gain, it only requires slightly more training time. Since our approach does not change the segmentation model’s structure, our method performs inference almost the same as others.

Ablation on different arrangement options of LCG and GSG modules in SGG. As shown in Fig. 3, in our framework, the LCG and GSG modules are arranged as a sequence with n steps. Here, we conduct ablation studies to investigate the impact of different arrangement options. Tab. 5 shows the ablation results, where “LCG/GSG” de-

Table 5. Ablation studies on different arrangement options of LCG and GSG modules in SGG.

	Model	G→C	S→C
(a)	LCG/GSG	61.4	60.5
(b)	GSG/LCG	61.3	60.5
(c)	Alternate	61.9	61.0
(d)	RandMix	61.6	60.7

notes first using $\frac{n}{2}$ LCG modules and then $\frac{n}{2}$ GSG modules, “GSG/LCG” represents first using $\frac{n}{2}$ GSG modules and then $\frac{n}{2}$ LSG modules, “Alternate” refers to alternating between LCG and GSG, and “RandMix” represents randomly arranging the two types of modules. We can see that the alternating option performs best, outperforming others by 0.3%~0.6%. Therefore, we adopt the alternating option to arrange LCG and GSG modules in the SGG scheme.

Qualitative Segmentation Results. Fig. 6 shows some qualitative segmentation results. The “Source Only” results were obtained by directly applying the segmentation model trained on the source domain to the target domain. We can see that even only using our baseline model, the segmentation results are improved obviously (e.g., the road and car). After adding the SGG scheme to our baseline, the boundary of small objects (e.g., the person) becomes more precise, which demonstrates the SGG’s effectiveness in preserving details. When further using the PTL strategy, our approach can provide more accurate segmentation results, indicating that the PTL can help SGG to preserve details better.

6. Conclusion

In this paper, we have proposed a novel label-guided image translation framework to tackle DASS. Our approach leverages the diffusion model and incorporates a Semantic Gradient Guidance (SGG) scheme to guide the image translation based on source labels. We have also proposed a Progressive Translation Learning (PTL) strategy to facilitate our SGG in working across a large domain gap. Extensive experiments have shown the effectiveness of our method.

Acknowledgement This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-PhD-2022-01-027).

References

- [1] Johannes Ackermann and Minjun Li. High-resolution image editing via multi-stage blended diffusion. *arXiv preprint arXiv:2210.12965*, 2022. 2, 3
- [2] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarrow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:17981–17993, 2021. 3
- [3] 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. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 40(4):834–848, 2017. 1, 7
- [4] Shoufa Chen, Peize Sun, Yibing Song, and Ping Luo. Diffusiondet: Diffusion model for object detection. *arXiv preprint arXiv:2211.09788*, 2022. 3
- [5] Yuhua Chen, Wen Li, and Luc Van Gool. Road: Reality oriented adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7892–7901, 2018. 2, 3
- [6] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. Crdoco: Pixel-level domain transfer with cross-domain consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1791–1800, 2019. 2, 3, 7
- [7] Yi-Hsin Chen, Wei-Yu Chen, Yu-Ting Chen, Bo-Cheng Tsai, Yu-Chiang Frank Wang, and Min Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1992–2001, 2017. 2, 3
- [8] Yiting Cheng, Fangyun Wei, Jianmin Bao, Dong Chen, Fang Wen, and Wenqiang Zhang. Dual path learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9082–9091, 2021. 3, 7
- [9] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14367–14376, October 2021. 3, 4
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3213–3223, 2016. 7
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:8780–8794, 2021. 2, 3, 4, 5
- [12] Liang Du, Jingang Tan, Hongye Yang, Jianfeng Feng, Xiangyang Xue, Qibao Zheng, Xiaoqing Ye, and Xiaolin Zhang. Ssf-dan: Separated semantic feature based domain adaptation network for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 982–991, 2019. 2, 3
- [13] Lin Geng Foo, Jia Gong, Hossein Rahmani, and Jun Liu. Distribution-aligned diffusion for human mesh recovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 3
- [14] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. 3
- [15] Jin Gao, Jialing Zhang, Xihui Liu, Trevor Darrell, Evan Shelhamer, and Dequan Wang. Back to the source: Diffusion-driven test-time adaptation. *arXiv preprint arXiv:2207.03442*, 2022. 4, 6
- [16] Li Gao, Lefei Zhang, and Qian Zhang. Addressing domain gap via content invariant representation for semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 7528–7536, 2021. 1, 2, 3, 7
- [17] Jia Gong, Lin Geng Foo, Zhipeng Fan, Qihong Ke, Hossein Rahmani, and Jun Liu. Diffpose: Toward more reliable 3d pose estimation. *arXiv preprint arXiv:2211.16940*, 2022. 3
- [18] Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and LingPeng Kong. Diffuseq: Sequence to sequence text generation with diffusion models. *arXiv preprint arXiv:2210.08933*, 2022. 3
- [19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 27. Curran Associates, Inc., 2014. 2
- [20] Tianpei Gu, Guangyi Chen, Junlong Li, Chunze Lin, Yongming Rao, Jie Zhou, and Jiwen Lu. Stochastic trajectory prediction via motion indeterminacy diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17113–17122, 2022. 3
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 7
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:6840–6851, 2020. 2, 3, 7, 8
- [23] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning (ICML)*, pages 1989–1998. Pmlr, 2018. 2, 3, 7
- [24] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016. 2, 3, 7
- [25] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In *Proceedings*

- of the *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9924–9935, 2022. 7
- [26] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Hrda: Context-aware high-resolution domain-adaptive semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 372–391. Springer, 2022. 7
- [27] Qingqing Huang, Daniel S Park, Tao Wang, Timo I Denk, Andy Ly, Nanxin Chen, Zhengdong Zhang, Zhishuai Zhang, Jiahui Yu, Christian Frank, et al. Noise2music: Text-conditioned music generation with diffusion models. *arXiv preprint arXiv:2302.03917*, 2023. 3
- [28] Rongjie Huang, Zhou Zhao, Huadai Liu, Jinglin Liu, Chenye Cui, and Yi Ren. Prodiff: Progressive fast diffusion model for high-quality text-to-speech. In *Proceedings of the 30th ACM International Conference on Multimedia (ACM MM)*, pages 2595–2605, 2022. 3
- [29] Zilong Huang, Xinggong Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 603–612, 2019. 1
- [30] Xinyue Huo, Lingxi Xie, Hengtong Hu, Wengang Zhou, Houqiang Li, and Qi Tian. Domain-agnostic prior for transfer semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7075–7085, 2022. 3, 7
- [31] Zhengkai Jiang, Yuxi Li, Ceyuan Yang, Peng Gao, Yabiao Wang, Ying Tai, and Chengjie Wang. Prototypical contrast adaptation for domain adaptive semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 36–54. Springer, 2022. 7, 9
- [32] Jack Kiefer and Jacob Wolfowitz. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, pages 462–466, 1952. 7
- [33] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2426–2435, 2022. 3
- [34] Heeseung Kim, Sungwon Kim, and Sungroh Yoon. Guided-ts: A diffusion model for text-to-speech via classifier guidance. In *International Conference on Machine Learning (ICML)*, pages 11119–11133. PMLR, 2022. 3
- [35] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [36] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. 7
- [37] Xin Lai, Zhuotao Tian, Xiaogang Xu, Yingcong Chen, Shu Liu, Hengshuang Zhao, Liwei Wang, and Jiaya Jia. Decouplenet: Decoupled network for domain adaptive semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 369–387. Springer, 2022. 3, 7, 9
- [38] Geon Lee, Chanho Eom, Wonkyung Lee, Hyekang Park, and Bumsub Ham. Bi-directional contrastive learning for domain adaptive semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 38–55. Springer, 2022. 2, 3, 7, 9
- [39] Seunghun Lee, Wonhyeok Choi, Changjae Kim, Minwoo Choi, and Sunghoon Im. Adas: A direct adaptation strategy for multi-target domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19196–19206, 2022. 7
- [40] Guangrui Li, Guoliang Kang, Wu Liu, Yunchao Wei, and Yi Yang. Content-consistent matching for domain adaptive semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 440–456. Springer, 2020. 2, 3, 7
- [41] Rui Li, Wenming Cao, Qianfen Jiao, Si Wu, and Hau-San Wong. Simplified unsupervised image translation for semantic segmentation adaptation. *Pattern Recognition (PR)*, 105:107343, 2020. 2, 7
- [42] Ruihuang Li, Shuai Li, Chenhang He, Yabin Zhang, Xu Jia, and Lei Zhang. Class-balanced pixel-level self-labeling for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11593–11603, 2022. 2, 3, 7
- [43] Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *arXiv preprint arXiv:2205.14217*, 2022. 3
- [44] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6936–6945, 2019. 2, 3, 7, 9
- [45] Ting-Hsuan Liao, Songwei Ge, Yiran Xu, Yao-Chih Lee, Badour AlBahar, and Jia-Bin Huang. Text-driven visual synthesis with latent diffusion prior. *arXiv preprint arXiv:2302.08510*, 2023. 3
- [46] Yahao Liu, Jinhong Deng, Xincheng Gao, Wen Li, and Lixin Duan. Bapa-net: Boundary adaptation and prototype alignment for cross-domain semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8801–8811, 2021. 7
- [47] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440, 2015. 1, 7
- [48] Yawei Luo, Ping Liu, Tao Guan, Junqing Yu, and Yi Yang. Significance-aware information bottleneck for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6778–6787, 2019. 2, 3
- [49] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2507–2516, 2019. 2, 3, 7
- [50] Haoyu Ma, Xiangru Lin, Zifeng Wu, and Yizhou Yu. Coarse-to-fine domain adaptive semantic segmentation with photo-

- metric alignment and category-center regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4051–4060, 2021. 3, 7
- [51] Naoki Matsunaga, Masato Ishii, Akio Hayakawa, Kenji Suzuki, and Takuya Narihira. Fine-grained image editing by pixel-wise guidance using diffusion models. *arXiv preprint arXiv:2212.02024*, 2022. 2, 3
- [52] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2021. 2, 3, 4
- [53] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2018. 2, 3
- [54] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4500–4509, 2018. 2, 3, 7
- [55] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning (ICML)*, pages 8162–8171. PMLR, 2021. 2, 3
- [56] Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3764–3773, 2020. 3, 7
- [57] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European Conference on Computer Vision (ECCV)*, pages 102–118. Springer, 2016. 1, 7
- [58] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. 2
- [59] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241. Springer, 2015. 1, 3, 7
- [60] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3234–3243, 2016. 1, 7
- [61] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Unsupervised domain adaptation for semantic segmentation with gans. *arXiv preprint arXiv:1711.06969*, 2(2):2, 2017. 2, 3
- [62] Divya Saxena and Jiannong Cao. Generative adversarial networks (gans) challenges, solutions, and future directions. *ACM Computing Surveys (CSUR)*, 54(3):1–42, 2021. 2, 3
- [63] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 7
- [64] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning (ICML)*, pages 2256–2265. PMLR, 2015. 3
- [65] Xuan Su, Jiaming Song, Chenlin Meng, and Stefano Ermon. Dual diffusion implicit bridges for image-to-image translation. In *International Conference on Learning Representations (ICLR)*, 2022. 4
- [66] Ruoqi Sun, Xinge Zhu, Chongruo Wu, Chen Huang, Jianping Shi, and Lizhuang Ma. Not all areas are equal: Transfer learning for semantic segmentation via hierarchical region selection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4360–4369, 2019. 2, 3
- [67] Jaesung Tae, Hyeongju Kim, and Taesu Kim. Editts: Score-based editing for controllable text-to-speech. *arXiv preprint arXiv:2110.02584*, 2021. 3
- [68] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schuster, and Manmohan Chandraker. Domain adaptation for structured output via discriminative patch representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1456–1465, 2019. 2, 3
- [69] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:11287–11302, 2021. 2, 3
- [70] Andrey Voynov, Kfir Aberman, and Daniel Cohen-Or. Sketch-guided text-to-image diffusion models. *arXiv preprint arXiv:2211.13752*, 2022. 3
- [71] Yuxi Wang, Junran Peng, and ZhaoXiang Zhang. Uncertainty-aware pseudo label refinery for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9092–9101, 2021. 3, 7
- [72] Zhonghao Wang, Mo Yu, Yunchao Wei, Rogerio Feris, Jinjun Xiong, Wen-mei Hwu, Thomas S Huang, and Honghui Shi. Differential treatment for stuff and things: A simple unsupervised domain adaptation method for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12635–12644, 2020. 2, 3
- [73] Simon Welker, Henry N Chapman, and Timo Gerkmann. Driftrec: Adapting diffusion models to blind image restoration tasks. *arXiv preprint arXiv:2211.06757*, 2022. 2
- [74] Shoule Wu and Ziqiang Shi. Itôwave: Itô stochastic differential equation is all you need for wave generation. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8422–8426. IEEE, 2022. 3
- [75] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:12077–12090, 2021. 7

- [76] Jinyu Yang, Weizhi An, Sheng Wang, Xinliang Zhu, Chaochao Yan, and Junzhou Huang. Label-driven reconstruction for domain adaptation in semantic segmentation. In *European Conference on Computer Vision (ECCV)*, pages 480–498. Springer, 2020. 3, 7
- [77] Jinyu Yang, Weizhi An, Chaochao Yan, Peilin Zhao, and Junzhou Huang. Context-aware domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 514–524, 2021. 7
- [78] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4085–4095, 2020. 7
- [79] Yueqin Yin, Lianghua Huang, Yu Liu, and Kaiqi Huang. Diffgar: Model-agnostic restoration from generative artifacts using image-to-image diffusion models. *arXiv preprint arXiv:2210.08573*, 2022. 2
- [80] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. 7
- [81] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, Yong Wang, and Fang Wen. Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12414–12424, 2021. 2, 3, 7
- [82] Xiaohong Zhang, Yi Chen, Ziyi Shen, Yuming Shen, Haofeng Zhang, and Yudong Zhang. Confidence-and-refinement adaptation model for cross-domain semantic segmentation. *IEEE Transactions on Intelligent Transportation Systems (TITS)*, 23(7):9529–9542, 2022. 3, 7
- [83] Yiheng Zhang, Zhaofan Qiu, Ting Yao, Dong Liu, and Tao Mei. Fully convolutional adaptation networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6810–6818, 2018. 2, 3
- [84] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip HS Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1529–1537, 2015. 1
- [85] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2223–2232, 2017. 2, 3
- [86] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289–305, 2018. 2, 3
- [87] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5982–5991, 2019. 2, 3