DynaST: Dynamic Sparse Transformer for Exemplar-Guided Image Generation

Songhua Liu, Jingwen Ye, Sucheng Ren, and Xinchao Wang*

National University of Singapore {songhua.liu,suchengren}@u.nus.edu, {jingweny,xinchao}@nus.edu.sg

Abstract. One key challenge of exemplar-guided image generation lies in establishing fine-grained correspondences between input and guided images. Prior approaches, despite the promising results, have relied on either estimating dense attention to compute per-point matching, which is limited to only coarse scales due to the quadratic memory cost, or fixing the number of correspondences to achieve linear complexity, which lacks flexibility. In this paper, we propose a dynamic sparse attention based Transformer model, termed Dynamic Sparse Transformer (DynaST), to achieve fine-level matching with favorable efficiency. The heart of our approach is a novel dynamic-attention unit, dedicated to covering the variation on the optimal number of tokens one position should focus on. Specifically, DynaST leverages the multi-layer nature of Transformer structure, and performs the dynamic attention scheme in a cascaded manner to refine matching results and synthesize visually-pleasing outputs. In addition, we introduce a unified training objective for DynaST, making it a versatile reference-based image translation framework for both supervised and unsupervised scenarios. Extensive experiments on three applications, pose-guided person image generation, edge-based face synthesis, and undistorted image style transfer, demonstrate that DynaST achieves superior performance in local details, outperforming the state of the art while reducing the computational cost significantly. Our code is available here.

Keywords: Dynamic Sparse Attention, Transformer, Exemplar-Guided Image Generation

1 Introduction

Semantic-based conditional image generation refers to synthesizing a photorealistic image with aligned semantic information, and finds its application across a wide spectrum of scenarios including label-to-scene [16,54,37,70,48,55,10,44], sketch-to-photo [24,11,25], and landmark-to-face [42,61,47]. Exemplar-guided image generation, as a mainstream approach along this route, provides users with the flexibility to specify an image as reference to control the appearance, style, or identity for the output image, and has recently received wide attentions across academic and industrial communities.

^{*} Corresponding author.



Fig. 1. Left: Different numbers of matching are required for different query locations in the exemplar-guided image generation task. This fact has been largely overlooked by prior methods that impose only static number of matching. The proposed DynaST, by contrast, is dedicated to handling such variations. Right: Details in faces and edges can be refined with the propagation in the multi-layer structure of Transformer. Warping results of the exemplar images using attention maps in each layer are shown here.

The core problem of exemplar-guided image generation lies in guiding input semantics to focus on appropriate context in exemplars. Early methods have largely relied on holistic convolution, normalization, and non-linear transformation [36,37,52,1,71,67,46,44]. Despite the promising results on global-style migration and the favorable efficiency, these methods overlook the fine-grained local details and thus lead to coarse results. To account for local context in the exemplar-guided image generation, recent works [45,64,62] employ per-point attention mechanism to model spacial correspondence between input and reference images. Nevertheless, hindered by the quadratic time and memory complexity, such dense matching operations, unfortunately, again limit themselves to coarse scales, making them difficult to capture fine-grained details in reference images.

To alleviate this issue, the works of [69,63] propose to fix the number of feature points that each query position focuses on, thereby achieving a linearcomplexity model. The reason behind such a design lies in that, each query position in target images is, in reality, independent to and dissimilar to most points in the reference. This rationale, in turn, implies that the matching between the target and reference image is intrinsically sparse.

Unfortunately, such the static number of correspondences, in many cases, fail to capture the dynamic nature of matching. As shown in Fig. 1, different queries may end up having different numbers of necessary matching: in Fig. 1 Left (a) and (b), due to the scale variations, the highlighted query location in the target image corresponds to different numbers of locations in the reference; in Fig. 1 Left (c), however, the query has no correspondence at all, meaning that all points in the exemplar would be negative samples.

These facts motivate us to explore a more sophisticated technique to account for the dynamic matching inherent to the exemplar-guided image generation, while maintaining the sparsity to ensure computational efficiency. To this end, we propose a novel Transformer-based model, termed Dynamic SparseTransformer (DynaST). The heart in DynaST is the dynamic sparse attention module in contrast to previous dense and static ones. Specifically, since the adopted attention strategy is sparse, a large number of potential matching candidates are dismissed. To alleviate this issue, we are inspired by the architecture of Transformer [51] and back up the proposed attention strategy with a multi-layer structure, which enables DynaST to explore and evolve matching results in a cascade manner. As shown in Fig. 1 Right, we visualize the matching results by warping exemplar images using attention maps in each layer and observe that with the feature propagation in the multi-layer structure of Transformer, the model produces finer matching results, especially in local details such as faces and edges. To refine matching results progressively, each current matching candidate would pass through a differentiable and learnable attention link pruning unit in each layer, to predict whether it is an irrelevant correspondence.

Consequently, (1) such dynamic pruning manner encourages more precise and cleaner matching results; (2) due to the effective higher-order dependency modeling capability of Transformer, DynaST is competent in aggregating relevant features and synthesizing high-quality outputs; and (3) sparse attention in DynaST guarantees the efficiency of even full-resolution matching construction. Moreover, we introduce a unified training objective, so that DynaST is readily applicable for universal exemplar-guided image generation under both supervised and unsupervised settings.

We conduct extensive evaluations on three challenging tasks: pose-guided person image generation, edge-based face synthesis, and undistorted image style transfer. In all experiments, DynaST outperforms state-of-the-art exemplar-guided image generation models significantly, by up to 36.7% in quantitative metrics, and achieves near-real-time inference efficiency, with more than $2 \times$ speedup compared with previous state-of-the-art full-resolution matching solutions.

2 Related Works

2.1 Exemplar-Guided Image Generation

Exemplar-guided image generation has recently emerged as a popular task in the computer vision community. Park *et al.* [37] propose spatially-adaptive normalization (SPADE) module that generates normalization parameters based on given semantic information. In their work, exemplar images are fed to a VAE to encode the overall style and appearance, which guide the following generation procedure. Nevertheless, it is difficult to migrate local details in exemplar images due to such global transformation. Similar drawback also exists in works like [36,52,1,71,67,46,20]. To enhance the generation of local textures, Ren *et al.* [40] introduce a neural texture extraction and distribution module. Recent attempts have also been made to introduce point-wise attention to exemplar-guided image generation and achieved superior results. For instance,

Zhang et al. [64] propose CoCosNet to learn the correspondence between input semantics and reference images. Zhan et al. [62] use unbalanced optimal transport to achieve the same goal. However, without any sparse mechanism, the quadratically increasing memory cost prevents these methods from learning fine-grained correspondence, which is important for synthesizing high-quality images. Although Zhou et al. [69] propose CoCosNet-v2 that is capable to learn full-resolution correspondence, the iterative global searching process by GRU makes a negative influence on the efficiency. DynaST in this paper brings the best of two worlds: it establishes matching at fine scales based on the dynamic sparse mechanism, which can generate images with high-quality local details while maintaining high efficiency simultaneously.

$\mathbf{2.2}$ **Image-wise Matching**

Given a pair of images, image matching such as [19,27,31,41,49,17,43] aims to find pixel-wise correspondence leveraging local features, which is a fundamental problem in computer vision and is one related field to exemplar-guided image generation in this paper. The key difference is that cross-domain matching establishment, semantic map to the exemplar image, is required in the image synthesis problem unlike matching between two highly correlated images. This is also one major difference between the general cross-domain exemplar-guided image generation and reference-based image super-resolution [34,66,18,56].

$\mathbf{2.3}$ **Efficient Transformer**

The full token-wise attention operation in standard Transformer [51,9] poses high requirements on memory and significantly increases computational cost. Thus, a lot of works are devoted to designing efficient attention mechanisms for Transformer or by extension, graph-based methods [58,57]. On the one hand, some works rely on heuristic strategies to lead a current token only focus on those in a certain local context [4,6,26,3]. Recently, more strategies based on properties of images to boost the efficiency in vision Transformer are explored [60,39]. On the other hand, random sampling based Informer [68], locality sensitive hashing based Reformer [22], and approximated Softmax based Performer [5] achieve lower complexity with a fine theoretical guarantee. Wang et al. [53] only involve tokens with top K attention scores for feature aggregation. Similar strategy is also adopted in [69,63,44]. Although effective, it is not flexible enough to fix the number of attentive tokens, which fails to model the complex and changeable matching patterns in practice. Different from all these methods, the sparse mechanism in the attention module of this paper is based on prior knowledge in image matching, targeting at exemplar-guided image generation.

3 **Dynamic Sparse Transformer**

In this section, we illustrate details of the proposed DynaST model for exemplarguided image generation. The overview of DynaST is shown in Fig. 2. DynaST

4



Fig. 2. DynaST overview. Solid arrows denote flows of features and dotted ones denote inheritance of attention maps. Yellow, green, and blue blocks in the middle adopt dense attention, inter-scale sparse attention, and inner-scale sparse attention respectively.

takes three images as input: a reference image I_{ref} , a corresponding semantic map S_{ref} of I_{ref} (e.g., a pose image or an edge map), and a target semantic map S_{tgt} . It aims at synthesizing the image \hat{I}_{tgt} with the target semantic information specified in S_{tgt} and the appearance as well as the style in I_{ref} .

The proposed DynaST consists of three parts. The first is an *embedding* module (Sec. 3.1), which is established by a set of multi-scale layers to extract and aggregate features at different levels. The second is a *Transformer module* (Sec. 3.2), which restores features of target images with features of the semantic map as target and features of reference information as memory. The last one is a lightweight *decoder module* to synthesize final images, where the multiscale features generated by the Transformer module are the input. The training objectives and supervised signals for the pipeline are described in Sec. 3.3.

3.1 Embedding Module

Given an input semantic image S_{tgt} and a reference image I_{ref} along with its corresponding semantics S_{ref} , the embedding module produces a set of feature embedding, F_{tgt} and F_{ref} . DynaST adopts a hierarchical patch embedding module as a multi-scale generative model, to enable the scale-wise cascaded matching process. The proposed embedding module is utilized to obtain rich features and contextual representation. Besides, the position embedding is also included to make the network aware of the positional information in the following matching process. Specifically, we use two independent sets of linear transformations: E_{tgt}^{i} and E_{ref}^{i} , to obtain multi-scale patch embedding for target semantic map S_{tgt} as well as reference information I_{ref} and S_{ref} , where *i* denotes the embedding of the *i*-th scale with patch size $2^{i} \times 2^{i}$. The features at the *i*-th scale, F_{tgt}^{i} and F_{ref}^{i} , can then be written as:

$$F_{tgt}^{i} = X([E_{tgt}^{j}(S_{tgt})'|0 < j < M]),$$

$$F_{ref}^{i} = Y([E_{ref}^{j}([S_{ref}, I_{ref}])'|0 < j < M]),$$
(1)



Fig. 3. Left: Intuition of inter-scale (up) and inner scale (bottom) sparse attention mechanism. Here k = 2 is used for illustration. The yellow tokens represent queries in the target feature map and the blue ones represent tokens in the reference map. The dotted lines denote links masked by dynamic attention pruning. Right: Illustration of dynamic pruning and feature aggregation in DynaST blocks.

where $0 \leq i < M$, M is the number of scales, notation ' denotes the bilinear interpolation to unify the spacial dimension for scale i, $[\cdot]$ stands for the channelwise concatenation, and X and Y are two MLPs consisting of two convolutional blocks for non-linear transformation. In other words, features of one specific scale aggregates multi-level patch embedding information, which provides rich features and contextual representations for the following matching and transformation steps. We also concatenate a learnable positional embedding to F_{tgt} and F_{ref} before computing their attention scores, denoted as $F_{tgt}^{i,pos}$ and $F_{ref}^{i,pos}$ at the *i*-th scale. The details can be found in the supplement.

3.2 Transformer Module

The Transformer module is built for exemplar-guided image generation, by a set of dynamic sparse Transformer blocks (DynaST Block) and dense attention blocks (Dense Attn Block). The construction of these blocks consists of four steps, which are attention map computation, dynamic attention pruning, feature aggregation, and non-linear transformation. In the attention map computation step, Dense Attn Blocks are used at the coarsest scale to construct matching at a low resolution, while DynaST Blocks are used at higher levels to infer high-resolution matching based on previous attention results. At each higher scale, the first DynaST block adopts inter-scale sparse attention, and is therefore named as Inter-Scale DynaST Block; the remaining blocks adopt inner-scale sparse attention and are thus named as Inner-Scale DynaST Blocks. We will give details in the following sections.

Attention Map Computation. Attention map computation is crucial to guide each target semantic feature point to focus on correct positions on reference feature maps and further restore features of target images. In DynaST, at scale *i*, the *j*-th DynaST block takes feature map produced by the j - 1-th block $F_{tgt,j-1}^i$, the reference feature map F_{ref}^i , and attention map of previous DynaST block as input to compute attention scores. Note that $F_{tgt,0}^i$ is defined as F_{tgt}^i from the multi-scale patch embedding layer. At the coarsest level i = M - 1, Dense Attn Blocks with vanilla attention are used to derive attention scores:

$$A_j^{M-1} = \operatorname{softmax}(\tau \alpha (\overline{F_{tgt,j-1}^{M-1,pos}}) \beta (\overline{F_{ref}^{M-1,pos}})^{\top}), \qquad (2)$$

where α and β are implemented as two 1×1 convolutional kernels, τ is a hyperparameter controlling the smoothness of attention distribution, and \overline{x} denotes channel-independent instance normalization [50].

Then, at finer level i < M - 1, for the first DynaST block (j = 1), *inter-scale sparse attention* is proposed to compute the attention map at this layer:

$$A_{1}^{i} = \operatorname{softmax}(\tau \alpha (\overline{F_{tgt,0}^{i,pos}}) \beta (\operatorname{TopK}_{A_{-1}^{i+1}} (\overline{F_{ref}^{i,pos}}))^{\top}),$$
(3)

where $\operatorname{TopK}_{A_{-1}^{i+1}}$ denotes that matching candidates for a point in $F_{tgt,0}^{i,pos}$ come from those with top k large scores in the attention map of the last block in the previous scale. Note that a point in the previous scale would be divided into four in the current one. Therefore, there are $k \times 4$ matching candidates for attention map computation at this layer.

For the following blocks (j > 1), inner-scale sparse attention is performed to refine the attention matching at the current scale, based on the prior knowledge that the matching offset in a local area is likely to be the same [2]:

$$A_{j}^{i} = \operatorname{softmax}(\tau \alpha(\overline{F_{tgt,j-1}^{i,pos}})\beta(\mathcal{N}(\operatorname{TopK}_{A_{j-1}^{i}}(\overline{F_{ref}^{i,pos}})))^{\top}),$$
(4)

where notation \mathcal{N} is the operation to derive the points with the same matching offsets as the neighbors for one target point. For example, for one target point at the current layer, it firstly finds matching results of its right neighbor and then takes the left neighbors of these points as candidates. We define the up, bottom, left, and right points of one position plus the current point itself as the neighboring points in this paper. In this way, the number of matching candidates for each target point in the inner-scale sparse attention layer is $k \times 5$. The intuitions of inter/inner-scale sparse attention layer are illustrated in Fig. 3 Left. **Dynamic Attention Pruning.** Considering that not all the matching candidates in attention modules are necessary for feature aggregation, we propose dynamic attention pruning to decide whether an attention link between a point in the target map and that in the reference map is useful. To this end, we use two MLPs Ω and Φ to transform F_{tgt} and F_{ref} into a common feature space. Then, a sign function is applied on the inner product of the transformed results to obtain the decision D for each attention link:

$$P_j^i = \Omega(F_{tgt,j-1}^i) \Phi(F_{ref}^i)^\top$$

$$D_j^i = \begin{cases} 1, \quad P_j^i > 0, \\ 0, \quad otherwise \end{cases}$$
(5)

Note that in the above function, the sign function introduces obstacles for gradient-based optimization in training. To tackle this issue, we alternatively take gradients from *sigmoid* function during the backward propagation:

$$\frac{\mathrm{d}D_{j}^{i}}{\mathrm{d}P_{i}^{i}} = \frac{\exp(-P_{j}^{i})}{(1 + \exp(-P_{j}^{i}))^{2}}.$$
(6)

Thus, attention maps after the dynamic pruning operation are derived by:

$$\tilde{A}^i_j = D^i_j \odot A^i_j, \tag{7}$$

where \odot represents the element-wise multiplication.

Feature Aggregation. One straightforward way to conduct feature aggregation is to use pruned attention matrix \tilde{A}_{j}^{i} directly to perform weighted summation over reference features. However, due to the pruning operation, the sum of attention weights for one target point to all reference feature points is no longer guaranteed to be 1, which would result in unbalanced magnitudes in feature aggregation. For example, in the most extreme case, a target point would be untraceable in the reference image and attention decisions would be all 0 for this point. Then the aggregated features for this point would also be all 0, which impedes the synthesis of a plausible image. To alleviate this problem, we use features restored by a SPADE block SP [37] to compensate for the masked part by dynamic attention pruning:

$$F_{out} = (1 - \sum_{i} \tilde{A}^{i}_{j}) \odot \operatorname{SP}(F^{i}_{tgt,j-1}, S_{tgt}) + \tilde{A}^{i}_{j}\eta(F^{i}_{ref}),$$

$$F^{i}_{tgt,j} = \operatorname{Norm}(F_{out} + F^{i}_{tgt,j-1}),$$
(8)

where the summation is along the dimension of reference feature points, η is another 1×1 convolutional kernel, and *Norm* denotes the layer normalization. Key steps for feature aggregation in DynaST blocks are shown in Fig. 3 Right. **Non-Linear Transformation.** Finally, following the standard Transformer architecture, a residual block is added at the end of the DynaST block for nonlinear transformation:

$$F_{tgt,j}^{i} = \operatorname{Norm}(F_{tgt,j}^{i}' + \operatorname{Conv}(\operatorname{ReLU}(\operatorname{Conv}(F_{tgt,j}^{i}')))).$$
(9)

3.3 Training Objectives

DynaST is a universal framework for exemplar-guided image generation, which is compatible for objectives of both supervised and unsupervised tasks. The overall training objective consists of two parts by default: task-specific loss \mathcal{L}_t and matching loss \mathcal{L}_m .

Task-Specific Loss. Firstly, task-specific loss \mathcal{L}_t targets at the task itself and is flexible for different forms of loss functions in our model. Typically, for supervised tasks like pose-guided person image generation, the objective is defined as the

MSE between generated images \hat{I}_{tgt} and ground truths I_{tgt} in original image space and perceptual feature space plus an adversarial loss term:

$$\mathcal{L}_{t} = \|\hat{I}_{tgt} - I_{tgt}\|_{2}^{2} + \sum_{i} \lambda_{i} \|\phi_{i}(\hat{I}_{tgt}) - \phi_{i}(I_{tgt})\|_{2}^{2} + \lambda_{adv} \max\{\log \operatorname{Dis}(S_{tgt}, I_{tgt}) + \log(1 - \operatorname{Dis}(S_{tgt}, \hat{I}_{tgt}))\},$$
(10)

where ϕ denotes a pretrained feature extractor (*e.g.*, VGG-19), subscript *i* specifies which layer features come from, λ_i controls the weight of each layer, *Dis* represents a discriminator to be trained alternatively with the generator, and λ_{adv} is the weight of the adversarial term [1]. For another example, style transfer is an unsupervised task, whose loss function can be written as:

$$\mathcal{L}_{t} = l_{c} + \lambda_{s} l_{s}, \quad l_{c} = \|\phi_{4,1}(I_{cs}) - \phi_{4,1}(I_{c})\|_{2}^{2}$$

$$l_{s} = \sum_{i=1}^{4} (\|\mu(\phi_{i,1}(I_{cs})) - \mu(\phi_{i,1}(I_{s})))\|_{2}^{2} + \|\sigma(\phi_{i,1}(I_{cs})) - \sigma(\phi_{i,1}(I_{s})))\|_{2}^{2},$$
(11)

where $\phi_{x_{-1}}$ aims to extract features of ReLU_x-1 layer of a VGG-19 network pretrained on ImageNet [7], and μ and σ denote mean and standard deviation of each feature channel respectively [21,13]. In DynaST, content images I_c and style images I_s are input to the embedding modules E_{tgt} and E_{ref} respectively, and style transfer images I_{cs} are the framework output \hat{I}_{tgt} .

Matching Loss. To provide a more direct supervision signal for matching modules and dynamic pruning modules to produce proper attention maps, we introduce a matching loss \mathcal{L}_m that uses the output attention maps to warp the reference images and measures the task-specific loss produced by the warped images. To be specific, we denote the correlation maps, results of Eq. 2, 3, and 4 before *softmax*, by C_j^i . The warp matrices W_j^i are derived taking attention decision D_j^i into consideration:

$$W_j^i = \frac{D_j^i \odot \exp(C_j^i)}{\sum \{D_j^i \odot \exp(C_j^i)\} + \epsilon},$$
(12)

where the summation is over reference feature points and ϵ is a small constant for numerical stability. Then, the warped reference images are derived by:

$$\hat{I}^i_{warp,j} = W^i_j I'_{ref},\tag{13}$$

where I'_{ref} denotes the resized version of the reference images to keep dimension scales of W_i^i and I_{ref} the same. Matching loss is defined by the MSE:

$$\mathcal{L}_m = \sum_i \sum_j \|\hat{I}_{warp,j}^i - I_{tgt}'\|_2^2.$$
(14)

Finally, the overall objective is given by a weighted sum of \mathcal{L}_t and \mathcal{L}_m :

$$\mathcal{L} = \mathcal{L}_t + \lambda_m \mathcal{L}_m, \tag{15}$$

with λ_m controlling the weight of term \mathcal{L}_m .



Fig. 4. Comparisons with state-of-the-art exemplar-guided image generation methods on *DeepFashion* and *CelebA-HQ* datasets.

4 Experiments

Implementation Details. For all the experiments, DynaST are trained under 256×256 resolution. 4 different scales are set to be 32, 64, 128 and 256, where dimensions of the corresponding feature channel are 512, 256, 128 and 64 respectively. Each level of Transformation module is built by 2 blocks, where on the coarsest scale there are 2 *Dense Attn Blocks* and on each upper level there is 1 *Inter-Scale DynaST Block* and 1 *Inner-Scale DynaST Block*. For supervised tasks, hyper-parameters λ_m and λ_{adv} are set as 100 and 10, respectively. For style transfer, λ_m and λ_s are set as 1 and 3, respectively. Matching loss defined in Eq. 14 is not adopted here. The smoothness parameter τ is set as 100 and k = 4 is used by default when selecting attention candidates. Training is done on 8 Tesla V100 GPUs with batch size 32.

Datasets. For the pose-guided person image generation and the edge-guided face generation tasks, *DeepFashion* [32] and *CelebA-HQ* [33] datasets are used. Splits of training and validation sets and policies on retrieval of input-exemplar image pairs are consistent with those in [64]. For style transfer, following the common settings, *MS-COCO* [29] and *WikiArt* [38] are adopted as content and style image sets for training respectively. During training, all the images are resized to 512×512 and then randomly cropped to 256×256 . Inference results under 512×512 resolution are reported in this paper.

4.1 Supervised Tasks

Comparison with Other Methods. On the pose-guided person image generation and the edge-guided face generation problems, we mainly compare our method with three state-of-the-art attention-based exemplar-guided image generation methods, including UNITE [62], CoCosNet [64], and CoCosNet-v2 [69]. The attention matching in UNITE and CoCosNet is limited on a relatively coarse scale (64×64) by the quadratic memory cost of the dense attention operation. Thus, as shown in Fig. 4, some detailed textures are not good enough, *e.g.*, cloth and face details in the 1st row of comparisons on *DeepFashion* and the beard in the 2nd row of comparisons on *CelebA-HQ*. CoCosNet-v2 leverages the



Fig. 5. Challenging scenarios involving (a) scale-dissimilar matching, (b) scale-similar matching, and (c) partially no matching, where the static matching scheme inevitably fails. The proposed DynaST, thanks to its dedicated scheme for handling dynamic numbers of matching, may well handle all such cases and yield visually plausible results.

Conv-GRU module to predict correspondence at fine scales, where noisy correspondences would be added easily due to the large search space under high resolutions. And there is no pruning to mask irrelevant matching in CoCosNetv2, which may lead to some artifacts, *e.g.*, straps in the 2nd row of comparisons on *DeepFashion*. As shown in the last column of each example, our method addresses these problems successfully with the proposed dynamic sparse attention based Transformer model, which generates more high-quality results.

Notably, one major difference between our DynaST and CoCosNet-v2 is that DynaST uses a dynamic number of matching points for feature aggregation, while CoCosNet-v2 only considers a fixed number of candidates, which fails to account for the dynamic property of matching in different cases. One illustrative example is shown in Fig. 5, where results by CoCosNet-v2 are less robust to cases when scales of input and exemplar are different, since one location must select a fixed number of matching points. When there are less informative correspondences, noises are inevitably introduced. By comparison, dynamic pruning involved in DynaST is more competent to handle such scale variation robustly.

To further illustrate the advantage of our method on matching construction, we use the attention matrix derived by each method to warp the exemplar image and report the warped results. As shown in Fig. 6, results by UNITE and Co-CosNet are blurry due to low-resolution matching, and results by CoCosNet-v2 contain too much noise. Compared with the above methods, our method can generate matching under the full resolution with the best quality.

Quantitatively, leveraging the paired samples in *DeepFashion* dataset, we compare warped results with ground-truth images and show the average L1 loss, PSNR, and SSIM scores in Tab. 1 Left, where our method performs the best. The

Method	DeepFashion				DeepFashion		CelebA-HQ	
momou	$L1\downarrow$	$\mathrm{PSNR}\uparrow$	$\mathrm{SSIM}\uparrow$	Time↓	FID↓	SWD↓	FID↓	SWD↓
Pix2PixHD [54]	-	-	-	-	25.2	16.4	62.7	43.3
SPADE [37]	-	-	-	-	36.2	27.8	31.5	26.9
MUNIT [14]	-	-	-	-	74.0	46.2	56.8	40.8
EGSC-IT [35]	-	-	-	-	29.0	39.1	29.5	23.8
UNITE [62]	13.1	16.7	13.2	14.9	13.1	16.7	13.2	14.9
CoCosNet [64]	0.067	18.48	0.80	11.5	14.4	17.2	14.3	15.2
CoCosNet-v2 [69]	$\underline{0.064}$	18.24	0.80	21.7	<u>13.0</u>	16.7	13.2	<u>14.0</u>
Ours-32	0.077	18.12	0.73	5.53	8.55	15.4	16.0	17.8
Ours-64	0.063	18.22	0.78	7.24	8.50	12.8	13.1	13.1
Ours-128	0.061	19.13	0.82	8.43	8.41	12.9	12.3	12.7
Ours wo Inner	0.064	18.30	0.82	9.45	8.88	12.0	14.7	17.2
Ours wo Dyna	0.063	19.04	0.81	9.26	9.32	21.8	15.3	19.0
Ours	0.054	19.25	0.83	9.63	8.36	11.8	12.0	12.4

Table 1. Left: Quantitative metrics on the quality and efficiency of matching establishment. Results are measured by comparing warped results with the ground truth on DeepFashion dataset. Running time for one sample ($\times 10^{-2}$ sec.) is shown here. Right: Quantitative metrics of image quality on two datasets.



Fig. 6. Ablation studies on dynamic pruning and inner-scale sparse attention. Intermediate warping results using attention maps are visualized to demonstrate the evolving of input-exemplar matching. Zoom-in for better visualization.

time needed for generating matching and synthesizing results for one sample is also included, measured on a single Nvidia 3090 GPU by averaging 1000 samples. Dense attention mechanism and iterative solving of optimal transport problem in CoCosNet and UNITE respectively leave a high computational burden. Recurrent prediction under full resolution in CoCosNet-v2 increases the latency further. Different from previous approaches, the efficient dynamic sparse attention operation in DynaST makes it achieve the most satisfactory computational speed while generating the best matching results impressively.

We report the widely-used FID [12] and SWD [23] metrics to reflect the distance of feature distributions between generated samples and real images, following Zhang *et al.* [64] in Tab. 1 Right. Our method outperforms previous ones significantly under both datasets, suggesting the best quality of results. On the other hand, we show the semantic, color, and texture consistency in Tab. 2, also under the same setting as [64]. Results in Tab. 2 demonstrate that our method achieves competent semantic restoration and style migration performance.

Method	1	DeepFashio	1	CelebA-HQ			
	$\mathrm{Sem.}\uparrow$	$\operatorname{Col.}\uparrow$	$\text{Tex.}\uparrow$	Sem.↑	$\operatorname{Col.}\uparrow$	$Tex.\uparrow$	
Pix2PixHD	0.943	NA	NA	0.914	NA	NA	
SPADE	0.936	0.943	0.904	0.922	0.955	0.927	
MUNIT	0.910	0.893	0.861	0.848	0.939	0.884	
EGSC-IT	0.942	0.945	0.916	0.915	0.965	0.942	
UNITE	0.957	0.973	0.930	0.952	0.966	0.950	
CoCosNet	0.968	0.982	0.958	0.949	0.977	0.958	
CoCosNet-v2	0.959	0.974	0.925	0.948	0.975	0.954	
Ours	0.975	0.974	0.937	0.952	0.980	0.959	

Table 2. Quantitative metrics of semantic (Sem.), color (Col.), and texture (Tex.) consistency on two datasets compared with state-of-the-art image synthesis methods.



Fig. 7. Comparisons with state-of-the-art undistorted style transfer methods.

Ablation Study. We conduct ablation studies on the two core ideas in this paper: dynamic pruning and sparse attention. Based on the full model, we (1) remove the dynamic pruning mechanism, with the corresponding result denoted as *Ours wo Dyna*; (2) replace all inner-scale DynaST blocks with inter-scale ones, denoted as *Ours wo Inner*; and (3) remove inter-scale sparse attention layers at different scales and only use attention up to a coarse scale instead of full resolution, denoted as *Ours-x*, where $x \in \{32, 64, 128\}$ representing the highest resolution used for matching. Evaluations on warped results and final results mentioned above are repeated using the resulting models. The quantitative results in Tab. 1 indicate that the incomplete models lead to inferior results.

Qualitatively, we visualize the intermediate warping results using attention maps in Fig. 6, to demonstrate how the matching results are evolving through multi-layer dynamic pruning and inter/inner-scale attention. For one thing, dynamic pruning is capable of suppressing noisy matching and helps generate a clearer view. The example in Fig. 5 also demonstrates the importance of dynamic pruning for the robustness to handle scale variation. For another, replacing inner-scale DynaST blocks may lead to some artifacts such as checkerboard, as shown in Fig. 6, due to the missing of local refinement. Removing them would have negative impacts on local details like areas of hair and cloth.

4.2 Undistorted Image Style Transfer

The full resolution matching mechanism makes DynaST well suitable to generate undistorted style transfer results. In order to demonstrate such advantage, we compare our DynaST with 6 state-of-the-art style transfer methods with the same or similar goal, including LST [28], MST [65], WCT2 [59], MCCNet [8], AdaAttN [30], and MAST [15]. As shown in Fig. 7, compared with the photorealistic style transfer method WCT2, our results migrate global and



Fig. 8. Left: visualization of content SSIM and style loss of different undistorted style transfer methods; Right: distribution of user preference.

local style patterns better, while compared with other methods, our results preserve texture details of content images best, *e.g.*, hair in the 1st row. In particular, our method is capable of dealing with complicated scenes like the 2nd row without distortion, where other methods fail. We also visualize *SSIM scores* with content images and *style loss* against style images in Fig. 8 Left, using the test dataset in [59]. It is demonstrated that DynaST, as the first full-resolution matching based method in style transfer, can even achieve comparable content preservation ability with photorealistic style transfer methods, while significantly improving the stylization effects. Furthermore, we conduct a *user study* with the same test dataset to reflect user preference. There are 155 users involved and 12 content-style pairs are shown to each randomly. They are invited to select their favorite one for each pair among results by the 7 methods. We receive 1860 votes in total and the preference distribution is shown in Fig. 8 Right, where our preference score outperforms others significantly. In this way, both qualitative and quantitative comparisons demonstrate the superiority of DynaST.

5 Conclusion

In this paper, we introduce a novel multi-scale Transformer model, DynaST, to account for dynamic sparse attention and construct fine-level matching in exemplar-guided image generation tasks. DynaST comes with a unified training objective, making it a versatile model for various exemplar-guided image generation tasks under both supervised and unsupervised settings. Extensive evaluations on multiple benchmarks demonstrate that the proposed DynaST outperforms previous state-of-the-art methods, on both the matching quality and the running efficiency.

Acknowledgement

This project is supported by AI Singapore (Award No.: AISG2-RP-2021-023) and NUS Faculty Research Committee Grant (WBS: A-0009440-00-00).

References

- AlBahar, B., Huang, J.B.: Guided image-to-image translation with bi-directional feature transformation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9016–9025 (2019)
- Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: Patchmatch: A randomized correspondence algorithm for structural image editing. ACM Trans. Graph. 28(3), 24 (2009)
- 3. Beltagy, I., Peters, M.E., Cohan, A.: Longformer: The long-document transformer. arXiv preprint arXiv:2004.05150 (2020)
- Child, R., Gray, S., Radford, A., Sutskever, I.: Generating long sequences with sparse transformers. arXiv preprint arXiv:1904.10509 (2019)
- Choromanski, K., Likhosherstov, V., Dohan, D., Song, X., Gane, A., Sarlos, T., Hawkins, P., Davis, J., Mohiuddin, A., Kaiser, L., Belanger, D., Colwell, L., Weller, A.: Rethinking attention with performers (2021)
- Dai, Z., Yang, Z., Yang, Y., Carbonell, J.G., Le, Q., Salakhutdinov, R.: Transformer-xl: Attentive language models beyond a fixed-length context. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 2978–2988 (2019)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A largescale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
- Deng, Y., Tang, F., Dong, W., Huang, H., Ma, C., Xu, C.: Arbitrary video style transfer via multi-channel correlation. arXiv preprint arXiv:2009.08003 (2020)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
- Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12873–12883 (2021)
- Gao, C., Liu, Q., Xu, Q., Wang, L., Liu, J., Zou, C.: Sketchycoco: Image generation from freehand scene sketches. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5174–5183 (2020)
- 12. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)
- Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1501–1510 (2017)
- Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-toimage translation. In: Proceedings of the European conference on computer vision (ECCV). pp. 172–189 (2018)
- Huo, J., Jin, S., Li, W., Wu, J., Lai, Y.K., Shi, Y., Gao, Y.: Manifold alignment for semantically aligned style transfer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14861–14869 (2021)
- Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)

- 16 Liu et al.
- Jiang, W., Trulls, E., Hosang, J., Tagliasacchi, A., Yi, K.M.: Cotr: Correspondence transformer for matching across images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6207–6217 (2021)
- Jiang, Y., Chan, K.C., Wang, X., Loy, C.C., Liu, Z.: Robust reference-based superresolution via c2-matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2103–2112 (2021)
- Jin, Y., Mishkin, D., Mishchuk, A., Matas, J., Fua, P., Yi, K.M., Trulls, E.: Image matching across wide baselines: From paper to practice. International Journal of Computer Vision 129(2), 517–547 (2021)
- Jing, Y., Liu, X., Ding, Y., Wang, X., Ding, E., Song, M., Wen, S.: Dynamic instance normalization for arbitrary style transfer. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 4369–4376 (2020)
- Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European conference on computer vision. pp. 694–711. Springer (2016)
- Kitaev, N., Kaiser, L., Levskaya, A.: Reformer: The efficient transformer. arXiv preprint arXiv:2001.04451 (2020)
- Lee, C.Y., Batra, T., Baig, M.H., Ulbricht, D.: Sliced wasserstein discrepancy for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10285–10295 (2019)
- Lee, J., Kim, E., Lee, Y., Kim, D., Chang, J., Choo, J.: Reference-based sketch image colorization using augmented-self reference and dense semantic correspondence. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5801–5810 (2020)
- Li, B., Zhao, F., Su, Z., Liang, X., Lai, Y.K., Rosin, P.L.: Example-based image colorization using locality consistent sparse representation. IEEE transactions on image processing 26(11), 5188–5202 (2017)
- Li, S., Jin, X., Xuan, Y., Zhou, X., Chen, W., Wang, Y.X., Yan, X.: Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. Advances in Neural Information Processing Systems **32**, 5243–5253 (2019)
- Li, X., Han, K., Li, S., Prisacariu, V.: Dual-resolution correspondence networks. Advances in Neural Information Processing Systems 33, 17346–17357 (2020)
- Li, X., Liu, S., Kautz, J., Yang, M.H.: Learning linear transformations for fast image and video style transfer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3809–3817 (2019)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
- Liu, S., Lin, T., He, D., Li, F., Wang, M., Li, X., Sun, Z., Li, Q., Ding, E.: Adaattn: Revisit attention mechanism in arbitrary neural style transfer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6649–6658 (2021)
- Liu, X., Zheng, Y., Killeen, B., Ishii, M., Hager, G.D., Taylor, R.H., Unberath, M.: Extremely dense point correspondences using a learned feature descriptor. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4847–4856 (2020)
- Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1096–1104 (2016)

- Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: Proceedings of the IEEE international conference on computer vision. pp. 3730– 3738 (2015)
- Lu, L., Li, W., Tao, X., Lu, J., Jia, J.: Masa-sr: Matching acceleration and spatial adaptation for reference-based image super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6368– 6377 (2021)
- Ma, L., Jia, X., Georgoulis, S., Tuytelaars, T., Van Gool, L.: Exemplar guided unsupervised image-to-image translation with semantic consistency. arXiv preprint arXiv:1805.11145 (2018)
- Ma, L., Jia, X., Sun, Q., Schiele, B., Tuytelaars, T., Van Gool, L.: Pose guided person image generation. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. pp. 405–415 (2017)
- Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2337–2346 (2019)
- Phillips, F., Mackintosh, B.: Wiki art gallery, inc.: A case for critical thinking. Issues in Accounting Education 26(3), 593–608 (2011)
- Ren, S., Zhou, D., He, S., Feng, J., Wang, X.: Shunted self-attention via multi-scale token aggregation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022)
- 40. Ren, Y., Fan, X., Li, G., Liu, S., Li, T.H.: Neural texture extraction and distribution for controllable person image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13535–13544 (2022)
- 41. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: Superglue: Learning feature matching with graph neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4938–4947 (2020)
- 42. Song, L., Lu, Z., He, R., Sun, Z., Tan, T.: Geometry guided adversarial facial expression synthesis. In: Proceedings of the 26th ACM international conference on Multimedia. pp. 627–635 (2018)
- 43. Sun, J., Shen, Z., Wang, Y., Bao, H., Zhou, X.: Loftr: Detector-free local feature matching with transformers. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8922–8931 (2021)
- 44. Tan, Z., Chen, D., Chu, Q., Chai, M., Liao, J., He, M., Yuan, L., Hua, G., Yu, N.: Efficient semantic image synthesis via class-adaptive normalization. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021)
- 45. Tang, H., Bai, S., Torr, P., Sebe, N.: Bipartite graph reasoning gans for person image generation
- Tang, H., Bai, S., Zhang, L., Torr, P.H., Sebe, N.: Xinggan for person image generation. In: European Conference on Computer Vision. pp. 717–734. Springer (2020)
- 47. Tang, H., Xu, D., Liu, G., Wang, W., Sebe, N., Yan, Y.: Cycle in cycle generative adversarial networks for keypoint-guided image generation. In: Proceedings of the 27th ACM International Conference on Multimedia. pp. 2052–2060 (2019)
- Tang, H., Xu, D., Yan, Y., Torr, P.H., Sebe, N.: Local class-specific and global image-level generative adversarial networks for semantic-guided scene generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7870–7879 (2020)
- Truong, P., Danelljan, M., Timofte, R.: Glu-net: Global-local universal network for dense flow and correspondences. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6258–6268 (2020)

- 18 Liu et al.
- 50. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Instance normalization: The missing ingredient for fast stylization. arXiv preprint arXiv:1607.08022 (2016)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017)
- 52. Wang, M., Yang, G.Y., Li, R., Liang, R.Z., Zhang, S.H., Hall, P.M., Hu, S.M.: Example-guided style-consistent image synthesis from semantic labeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1495–1504 (2019)
- Wang, P., Wang, X., Wang, F., Lin, M., Chang, S., Xie, W., Li, H., Jin, R.: Kvt: k-nn attention for boosting vision transformers. arXiv preprint arXiv:2106.00515 (2021)
- Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: Highresolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8798–8807 (2018)
- Wang, Y., Qi, L., Chen, Y.C., Zhang, X., Jia, J.: Image synthesis via semantic composition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13749–13758 (2021)
- 56. Yang, F., Yang, H., Fu, J., Lu, H., Guo, B.: Learning texture transformer network for image super-resolution. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5791–5800 (2020)
- Yang, Y., Feng, Z., Song, M., Wang, X.: Factorizable graph convolutional networks. In: Conference on Neural Information Processing Systems (2020)
- Yang, Y., Qiu, J., Song, M., Tao, D., Wang, X.: Distilling knowledge from graph convolutional networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2020)
- Yoo, J., Uh, Y., Chun, S., Kang, B., Ha, J.W.: Photorealistic style transfer via wavelet transforms. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9036–9045 (2019)
- 60. Yu, W., Luo, M., Zhou, P., Si, C., Zhou, Y., Wang, X., Feng, J., Yan, S.: Metaformer is actually what you need for vision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022)
- Zakharov, E., Shysheya, A., Burkov, E., Lempitsky, V.: Few-shot adversarial learning of realistic neural talking head models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9459–9468 (2019)
- Zhan, F., Yu, Y., Cui, K., Zhang, G., Lu, S., Pan, J., Zhang, C., Ma, F., Xie, X., Miao, C.: Unbalanced feature transport for exemplar-based image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15028–15038 (2021)
- Zhan, F., Yu, Y., Wu, R., Cui, K., Xiao, A., Lu, S., Shao, L.: Bi-level feature alignment for versatile image translation and manipulation. arXiv preprint arXiv:2107.03021 (2021)
- Zhang, P., Zhang, B., Chen, D., Yuan, L., Wen, F.: Cross-domain correspondence learning for exemplar-based image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5143–5153 (2020)
- Zhang, Y., Fang, C., Wang, Y., Wang, Z., Lin, Z., Fu, Y., Yang, J.: Multimodal style transfer via graph cuts. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5943–5951 (2019)

- Zhang, Z., Wang, Z., Lin, Z., Qi, H.: Image super-resolution by neural texture transfer. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7982–7991 (2019)
- Zheng, H., Liao, H., Chen, L., Xiong, W., Chen, T., Luo, J.: Example-guided image synthesis using masked spatial-channel attention and self-supervision. In: European Conference on Computer Vision. pp. 422–439. Springer (2020)
- Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., Zhang, W.: Informer: Beyond efficient transformer for long sequence time-series forecasting. In: Proceedings of AAAI (2021)
- Zhou, X., Zhang, B., Zhang, T., Zhang, P., Bao, J., Chen, D., Zhang, Z., Wen, F.: Cocosnet v2: Full-resolution correspondence learning for image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11465–11475 (2021)
- Zhu, P., Abdal, R., Qin, Y., Wonka, P.: Sean: Image synthesis with semantic regionadaptive normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5104–5113 (2020)
- Zhu, Z., Huang, T., Shi, B., Yu, M., Wang, B., Bai, X.: Progressive pose attention transfer for person image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2347–2356 (2019)