


# HyperNST: Hyper-Networks for Neural Style Transfer

Dan Ruta<sup>1</sup>, Andrew Gilbert<sup>1</sup>, Saeid Motiian<sup>2</sup>, Baldo Faieta<sup>2</sup>, Zhe Lin<sup>2</sup>, and John Collomosse<sup>1,2</sup>

<sup>1</sup> University of Surrey

<sup>2</sup> Adobe Research



**Abstract.** We present HyperNST; a neural style transfer (NST) technique for the artistic stylization of images, based on Hyper-networks and the StyleGAN2 architecture. Our contribution is a novel method for inducing style transfer parameterized by a metric space, pre-trained for style-based visual search (SBVS). We show for the first time that such space may be used to drive NST, enabling the application and interpolation of styles from an SBVS system. The technical contribution is a hyper-network that predicts weight updates to a StyleGAN2 pre-trained over a diverse gamut of artistic content (portraits), tailoring the style parameterization on a per-region basis using a semantic map of the facial regions. We show HyperNST to exceed state of the art in content preservation for our stylized content while retaining good style transfer performance.

## 1 Introduction

Neural style transfer (NST) methods seek to transform an image to emulate a given appearance or ‘style’ while holding the content or structure unchanged. Efficient, stylistically diverse NST remains an open challenge. Feed-forward NST methods are fast yet typically fail to represent a rich gamut of styles. At the same time, optimization based approaches can take several seconds or minutes at inference time, lacking the speed for practical, creative use. Moreover, NST algorithms are often driven by one or more exemplar ‘style’ images, rather than an intuitive parameter space, impacting their controllability as a creative tool.

In this work, we propose a fast feed-forward method for driving neural stylization (NST) parameterized by a metric embedding for style representation. Our approach is based upon a hyper-network trained to emit weight updates to a StyleGAN2 [13] model, trained on a large dataset of artistic portraits (e.g. AAHQ [16]) in order to specialize it to the depiction of the given target style. Our work is inspired by the recent StyleGAN-NADA [5], in which a CLIP [20] objective is optimized by fine-tuning a pre-trained

StyleGAN2 model to induce NST. We extend this concept to a feed-forward framework using a hyper-network to generate the weight updates. Furthermore, we introduce the use of a metric parameter space (ALADIN [22]) originally proposed for style driven visual search to condition the hyper-network prediction (vs. CLIP in [5]) and adaptively drive this parameterization using a semantic map derived from the source and target image. Without loss of generality, our experiments focus on the challenging domain of facial portraits driven using a semantic segmentation algorithm for this content class. We show our method improves target image content retention versus the state of the art while retaining comparable accuracy for diverse style transfer from a single model – and despite using a hyper-network exhibiting comparable inference speed to leading feed-forward NST [19]. Moreover, our method exhibits good controllability; using a metric space for our style code enables intuitive interpolation between diverse styles and region-level controllability of those parameters. We adopt the recent ALADIN style code for this purpose, raising the novel direction of unifying style based visual search and stylization from a single representation.

## 2 Related Work

**Neural Style Transfer (NST)** The seminal work of Gatys et al. [6] enabled artistic style transfer through neural models. This work demonstrates the correlation between artistic style and features extracted from specific layers in a pre-trained vision model. The AdaIN work [8] introduced parameterized style transfer through first and second moment matching, via mean and standard deviation values extracted from random target style images. MUNIT [9] explores domain translation in images through the de-construction of images into semantic content maps and global appearance codes. ALADIN [22] explored the creation of a metric space modeling artistic style, across numerous areas of style. The embedding space was trained in conjunction with AdaIN and a modified version of the MUNIT model and weakly supervised contrastive learning over the BAM-FG dataset. A follow-up work [21] studying multi-modal artistic style representation expanded upon ALADIN, pushing the representation quality further through a vision transformer model. CycleGAN also explored domain transfer in images, but through learned model weight spaces, encoding pairwise image translation functions into separate generators for each image domain. Enforced by cyclic consistency, the translation quality between a pair of image domains was high, at the cost of requiring bespoke models for each domain translation to be trained. Using StyleGAN as a generator model, Swapping Autoencoders [19] directly learn the embedding of images into a StyleGAN generator’s weight space while simultaneously encoding a vector representation of the visual appearance of an image externally. These models separately focus on landscapes, buildings, bedrooms, or faces.

**StyleGAN Inversion** The evolution of the StyleGAN models [12,13,10] explore generation of extremely realistic portrait images. They use weight modulation based editing of visual appearance in the generated images. They also undertake preliminary investigations into the inversion of existing images into the GANs’ weight spaces for reconstruction. The work in e4e [26] includes an undertaking of a deeper analysis of real image

embedding into the StyleGAN weight space, including the quality/editability tradeoff this imposes. Their work enables multiple vectors of editability for images generated through StyleGAN across several domains. Restyle [1] improves the reconstruction quality of images embedded into the weight space by executing three fine-tuning optimization steps at run-time. Also, using StyleGAN as target generators, HyperStyle [2] embeds real portrait images into the weight space with high fidelity. Similar to our approach, they use a hyper-network to generate weight updates for StyleGAN, trained to infer weights to bridge the gap between quick rough inversions and the fully detailed reference portrait images. Strengths of this approach are the high reconstruction fidelity and strong photorealistic editing control for portrait photos. They further undertake some early explorations at domain adaptation for images by changing StyleGAN checkpoints. However, the HyperStyle work focuses on photorealism and does not enable region-based control or style space parameterization as we propose in HyperNST.

Our approach is inspired by the StyleGAN-NADA [5] work, which explores style transfer in the StyleGAN weight space through CLIP-based optimization. Though effective, this incurs long-running optimization passes, which are impractical for wide use. Moreover, this method has no built-in methods to effectively embed real portrait images into the weight space for reconstruction and editing. Recently, FaRL [28] undertakes representation learning of facial images, through multi-modal vision-text data, with face segmentation models that cope well with various visual appearances for portraits.

### 3 Methodology

Domain adaptation models like CycleGAN perform well at image translation, where style features from a style image are correctly mapped to the matching semantic features in a content image. A limitation of CycleGAN is that transfer only between a single pair of image domains (styles) is possible. Hypernetwork [7] models are used to predict the weights of a target model. Such a hypernetwork can be conditioned on some input before inferring the target model weights. Early experiments with a CycleGAN backbone yielded promising results, especially when we re-framed the training process into a more manageable task by learning weight offsets to a pre-trained checkpoint rather than a from-scratch model. Motivated by these findings, we use a more extensive, modern generator such as StyleGAN2 with weight updates predicted using a hyper-network. This echoes the recent optimization based approach StyleGAN-NADA which updates weights via fine-tuning for stylization rather than hyper-network prediction. HyperStyle [2] uses a channel-wise mean shift in the target generator, which significantly reduces the target number of trainable weights into a practical range.

#### 3.1 HyperNST Architecture

We compose the stylization pipeline of our model with a hyper-network set up to generate weight deltas for a target frozen, pre-trained StyleGAN2 model. Fig 1 shows an architecture diagram of HyperNST, showing the losses and the conditioning process of the hyper-network upon the content image and semantically arranged ALADIN style codes of the style image.

We direct the training to find a set of weight deltas for StyleGAN2, which can generate the same image as the content image for reconstruction. We begin with the GAN inversion, using  $e_c$  to reconstruct a given content portrait image roughly. An encoder encodes this image into a tensor used by the hyper-network weight generator modules along with style information to predict weight deltas to apply to the frozen, pre-trained StyleGAN2. Stylization then occurs by changing the style embedding in the hypernetwork conditioning.

### 3.2 Conditioning on style, and facial semantic regions

We further introduce conditioning on a target style image by injecting a style representation embedding. We project the initial  $16 \times 16 \times 512$  encoding into  $16 \times 16 \times 256$ , using ALADIN [22] codes to compose the other half of the  $16 \times 16 \times 512$  tensor. Each ALADIN code is a  $1 \times 256$  vector representing the artistic style of an image. We project semantically arranged ALADIN embeddings into a second  $16 \times 16 \times 256$ , which we then concatenate together with the first to form the  $16 \times 16 \times 512$  tensor upon which the hypernetwork is conditioned.

Facial portraits are very heavily grounded in their semantic structure. We therefore condition the stylization process on these semantic regions to ensure that the content is maintained. We execute this using segmentation masks extracted via FaRL [28].

To condition on semantic regions, we aim to use the ALADIN style code representing only the style contained within a given semantic region. Given that ALADIN codes can only be extracted from square or rectangular inputs, we use the average ALADIN code extracted across several patches per semantic region.

We use FaRL to compute the semantic regions, extract patches from the image randomly, and use mean intersection over union (IoU) to ensure that patches mainly cover pixels attributed to these respective semantic classes for each region. The ALADIN codes from these patches are averaged to form the average ALADIN code for a given semantic class. These average style codes are tiled and re-arranged on a per-pixel basis to match the semantic segmentation maps of the *content* image before they are projected by an encoder into the second  $16 \times 16 \times 256$  tensor, as above.

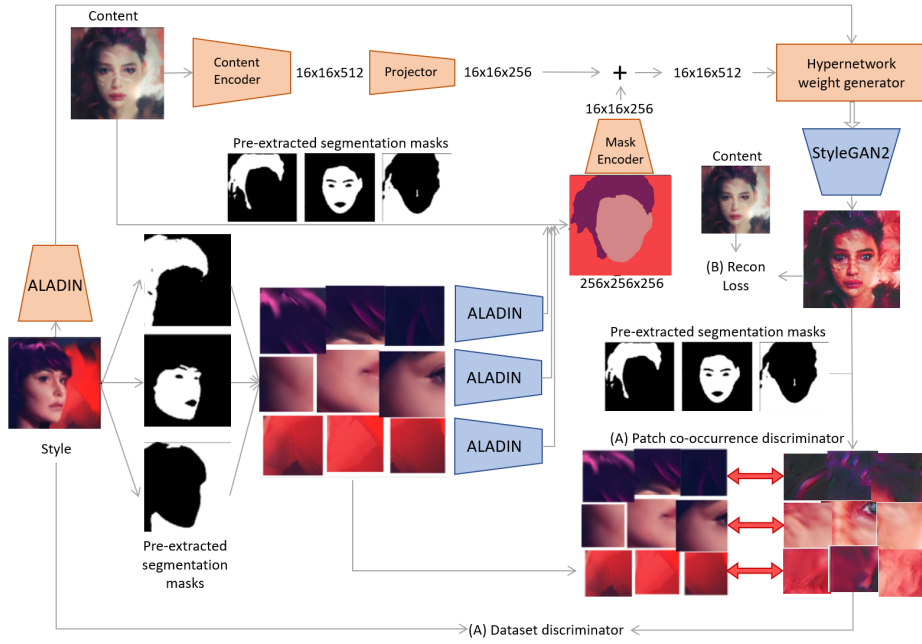
This projected  $16 \times 16 \times 256$  semantically arranged ALADIN tensor and the original  $16 \times 16 \times 256$  tensor encoded from the content image, are concatenated in the channel dimension to  $16 \times 16 \times 512$ , making up the final tensor which the hypernetwork weight delta generating modules operate over.

### 3.3 HyperNST training process

We use three iterations per batch refining methodology during training. We include a patch co-occurrence discriminator  $D_{patch}$  to introduce a style loss and an image-level discriminator  $D$ . The hypernetwork weight deltas generator  $H$  is conditioned on the encoded content images  $E_c(c)$ , and ALADIN style codes arranged by the semantic segmentation mask of the content image  $M(x, E_{mask})$ , seen in Eq 1.

$$M(x, E_{mask}) = E_{mask} \left( \text{rearrange} \left( \frac{\sum_{p=1}^P A(\text{crop}(x, SM_x)_p)}{P}, SM_c \right) \right) \quad (1)$$





**Fig. 1.** Architecture diagram of our approach. Facial semantic segmentation regions used in conditioning via ALADIN and guiding via patch co-occurrence discriminator a HyperStyle model into embedding a content portrait image into an AAHQ+FFHQ trained StyleGAN2 model, and using ALADIN style codes to stylize it towards the style of a style image. Blue modules represent frozen modules, and orange modules represent modules included in the training. + represents concatenation in the channels dimension. (A) represents the stylization pass losses, and (B) represents the reconstruction pass losses. Not pictured for clarity: the reconstruction pass uses the same semantic regions as the content image for the 256x256x256 semantically arranged ALADIN conditioning.

$$y(c, s) = G(\theta + H(M(s) \frown E_c(c), A(s))) \quad (2)$$

During training, two forward passes are executed, one for stylization (A) (Eq 2), and one for reconstruction (B). The reconstruction pass (B) loads content images and their respective ALADIN style code arranged by the content image semantic map and performs the original reconstruction code (Eq. 3). The stylization pass (A) performs style transfer by loading the content images  $c$ , and ALADIN style codes of other style images  $s$ . During stylization, the generation is performed with the mixed features combined to create stylized images, with the patch co-occurrence discriminator providing a learning signal to train the stylization (Eq. 4). In this pass, the discriminator is also trained, with the stylized images as *fake*, and the target style images as *real* (Eq 5).

$$\mathcal{L}_{rec}(c) = (c, y(c, c)) \quad (3)$$

$$\mathcal{L}_{pcd}(D_{patch}) = \mathbb{E}_{s \sim S, c \sim C} [-\log(D_{patch}(crops(s), crops(y(c, s))))] \quad (4)$$

$$\mathcal{L}_{disc}(c, s, D) = \mathbb{E}_{s \sim S, c \sim C} [-\log(D(s, y(c, s)))] \quad (5)$$

where:  $c$  is content image,  $s$  is style image,  $C$  are all content images,  $S$  are all style images,  $\theta$  are the original StyleGAN2 weights,  $H$  is the hypernetwork weight delta generator,  $G$  is the StyleGAN2 generator,  $A$  is ALADIN,  $P$  is number of patches per semantic region,  $SM$  are pre-extracted semantic segmentation masks, and  $\wedge$  represents concatenation.

$$\mathcal{L} = \lambda_1 \mathcal{L}_{rec} + \lambda_2 (\mathcal{L}_{disc} + \mathcal{L}_{pcd}) \quad (6)$$

The final loss is shown in Eq. 6, with Sec.4.3 describing ablations for the  $\lambda$  values.

### 3.4 Stylized target generator

The hypernetwork model is trained to infer weight delta values for weights in a StyleGAN2 model, which acts as the generator. Given that we work in the artistic style domain, we thus need the target generator to be able to model the weight space of a generator already able to produce high quality highly stylized images. We thus first fine-tune an FFHQ [11] StyleGAN2 model on the Artstation Artistic Face (AAHQ) dataset [16]. Largely popular in research centered around the facial images domain, the FFHQ dataset has been used by the majority of reference papers targeting this domain. Given our exploration into artistically manipulating the visual features of portrait images, we further use this AAHQ dataset, which encompasses portraits from across a large and varied corpus of artistic media. We continue to include FFHQ images, to ensure we keep high quality modeling capabilities for features more often found in photographic images than artistic renderings (like glasses and beards) to ensure real world images can still be encoded well.

We first train our model simply for ALADIN conditioned reconstruction of the AAHQ and FFHQ datasets for all layers, including the toRGB layers, which have been shown to target textures and colors [2]. We then fine-tune this pre-trained checkpoint with the goal of stylization, where we include the patch co-occurrence training for guidance and ALADIN code swapping in conditioning for style transfer.

During the stylization fine-tuning step, we freeze the content encoder and train only the hypernetwork weight delta prediction modules. This is to prevent the stylization signal from negatively affecting the model’s embedding abilities to reconstruct real images with high accuracy.

We run the stylization fine-tuning only on the part of the target StyleGAN layers. We find that the further into the StyleGAN2 model we apply stylization fine-tuning, the more the stylization affects colorization rather than textures and adjustments to the structure. We include the toRGB layers originally omitted in HyperStyle for their texture adjustments, and we train the weight deltas generation modules for layers 13 onward (out of 25). The weight delta generation for layers before this are frozen after their initial training, therefore still allowing reconstruction, but no more extended training during the stylization stage. This ensures that the overall facial structure of the images is not

too greatly affected during stylization training. Layer 13 is a *sweet middle spot* with the best balance between stylization and retaining good face structure reconstruction. A visual example of this phenomenon can be seen in Fig.3.

We further make changes to the patch co-occurrence discriminator. The Swapping Autoencoders model generates images in 1024x1024 resolution, with a discriminator operating over 8 patches of 128x128 dimensions. Our hypernetwork model generates images with a resolution of 256x256 pixels, therefore, we adjust our discriminator’s patch sizes to a lower size of 32x32.

### 3.5 Region mask driven patch discriminator

We also use facial semantic segmentation masks in the patch co-occurrence discriminator, to provide a style signal separated by semantic region.

In the original discriminator, patches are extracted from across the entire image at random. Instead, our process is repeated for each semantic class’ map. Patches are extracted at random, ensuring that the mean intersection over union (IoU) mostly covers pixels attributed to the respective semantic class. As in the original case, the losses from the patches are averaged to form the style learning signal. However, instead of comparing the stylized patches with patches randomly selected from the reference style image, we select patches bound by the respective semantic region - e.g. style hair patches are matched with stylized hair patches, background with background, face with face, etc.

**Implementation** The semantic regions predicted by FaRL contain several regions which individually cover only small areas in the pixel space. Given our use of the regions as patches we extract for ALADIN and the patch co-occurrence discriminator, we could not accurately use these small regions without the patches mainly containing pixel data from the surrounding regions. Instead, we group up the regions into 3 larger classes which typically cover larger areas of a portrait: (1) Hair, (2) Face, and (3) Background. Due to overlap that cannot be avoided, the *Face* region also contains eyes, noses, and lips. Furthermore, some images do not contain hair, for which this semantic class is not used. In the semantically driven patch co-occurrence loss, we use 3 patches for each region, totaling 9 patches.

## 4 Evaluation

Experiments generally required around 24 hours to converge. We ran our experiments on a single NVIDIA RTX 3090 with 24GB of VRAM. The HyperNST experiments were executed with a batch size of 1, which required around 22GB of VRAM, due to the high number of weights needed in the hypernetwork configuration.

### 4.1 Datasets

We create a test dataset for use in evaluation, using images from FFHQ [11], and AAHQ [16]. We extract 100 random content images from FFHQ and 100 random style images

from AAHQ. Together, these result in 10'000 evaluation images, when models are used to stylize all combinations.

We measure the content similarity using the two LPIPS variants (Alexnet [15], and VGG [25]), and we measure the style similarity using SIFID.

**LPIPS** [27] evaluates the variance between two images based on perceptual information better aligned to human perception compared to more traditional statistics based methods. We use this LPIPS variation to compute the average variation between each stylized image and its original content image. A lower average variation value would indicate a more similar semantic structure, therefore, a lower value is better.

Like Swapping Autoencoders, we employ the Single Image Fréchet Inception Distance (**SIFID**) introduced in SinGAN [23]. This metric evaluates FID, but uses only a single image at a time, which is most appropriate when evaluating style transfer computed using a single source sample. A lower value indicates better style transfer here.

Finally, we measure the time it takes to synthesize the stylized image in seconds per image. A low inference time is essential for the practicality of a model in real use cases. All timings were computed on an RTX 3090 GPU.

## 4.2 Baselines

To test the performance of our approach, we train 7 other methods on the FFHQ and AAHQ datasets: Swapping Autoencoders [19], Gatys [6], ArtFlow [3], PAMA [17], SANet [18], NNST [14], and ContraAST [4]. From these, Swapping Autoencoders is the most closely related, as it also operates over a StyleGAN2 space, thus, we separate it from the other more traditional methods which do not.

Table 1 shows the LPIPS, SIFID, and timing results computed for our own HyperNST method, and the other methods we evaluate against. The results show the superiority of HyperNST in retaining the most semantic structure in the facial reconstructions, indicated by the lowest LPIPS value. Meanwhile, we achieve comparable SIFID values (2.279) to Swapping AutoEncoders (1.948), indicating a small gap in stylization quality. Only our HyperNST model can embed a portrait into a model with semantic editing capabilities. Finally, the inference time of our model is similar to other methods and low enough for practical use of the model in real applications.

One disadvantage of the hypernetwork based approach is the current limitation with how strongly the StyleGAN2 model can be fine-tuned. As per our findings and the HyperStyle authors', hyper-learning weight deltas for every individual weight would be infeasible on current hardware (requiring over 3 billion parameters for a StyleGAN2 model) - we are limited by the technology of our time. So at best, we can aim to at most match a method where the model is fully trained, the closest comparison being Swapping AutoEncoders. The Swapping AutoEncoders model also fine-tunes a StyleGAN2 model. However, the more straightforward training approach affords the luxury to fine-tune individual weights in the generator, compared to the channel-wise mean shifts, currently possible with a hyper-network on today's hardware.

Nevertheless, hyper-network based NST is interesting to study, as such limitations are temporary. Both model and hardware improvements with subsequent works will reduce such limitations in time. We can compare hyper-network approaches directly with fully-trained counterparts for the current state of the art.

Model	LPIPS (Alexnet)	LPIPS (VGG)	SIFID	Time (s/img)	Interpolation	Editing
SAE [19]	0.334500	0.4857	1.948	0.10	✓	✗
HyperNST (ours)	0.000042	0.0017	2.279	0.35	✓	✓
Gatys [6]	0.000164	0.0030	1.369	14.43	✗	✗
ArtFlow [3]	0.000080	0.0022	1.347	0.32	✗	✗
PAMA [17]	0.000109	0.0029	0.522	0.14	✓	✗
SANet [18]	0.000180	0.0068	0.486	0.11	✓	✗
NNST [14]	0.000149	0.0030	0.871	55.40	✗	✗
ContraAST [4]	0.000133	0.0035	0.666	0.10	✗	✗

**Table 1.** Overall results, comparing metrics for our model’s content, style, and timing and the closest most similar model, Swapping AutoEncoders (top), and some further, more traditional methods (bottom). All methods were trained using AAHQ as the style dataset and FFHQ as the content dataset.

Style loss strength	LPIPS (Alexnet)	LPIPS (VGG)	SIFID
0.5	0.001480	0.002657	3.312
1.0	0.001310	0.002626	2.590
2.0	0.000042	0.001688	2.279
5.0	0.000097	0.000191	2.463
10.0	0.000091	0.000222	2.571

**Table 2.** Ablation results for varying the style loss strength

Figure 4 visualizes some style transfer results obtained with a Swapping Autoencoders (SAE) model, re-trained with the AAHQ dataset. Figure 2 shows the same visualization, obtained with HyperStyle. It is worth noting that the SAE model generates 1024x1024 images, whereas the target StyleGAN2 model in the hypernetwork setting uses a smaller, 256x256 version. Quality of stylization is somewhat equivocal for these approaches; in most cases, the SAE results are slightly more faithful in their style transfer accuracy but are less faithful at retaining the semantic structure of the portraits vs. HyperNST. The SAE model also has difficulty stylizing backgrounds.

Layer	LPIPS (Alexnet)	LPIPS (VGG)	SIFID
3	0.000140	0.00399	3.123
7	0.000190	0.00302	2.792
13	0.000042	0.00169	2.279
17	0.000153	0.00268	3.337
22	0.000042	0.00104	3.506

**Table 3.** Ablation results for varying the starting StyleGAN2 layer past which layers are included in the stylization fine-tuning.

### 4.3 Ablations

We perform the stylization fine-tuning on only part of the StyleGAN2 layers’ hyper-weight generators. Out of 25 layers, we find layer 13 to be a good balance between




Fig. 2. Qualitative HyperNST visualization of portrait style transfer.

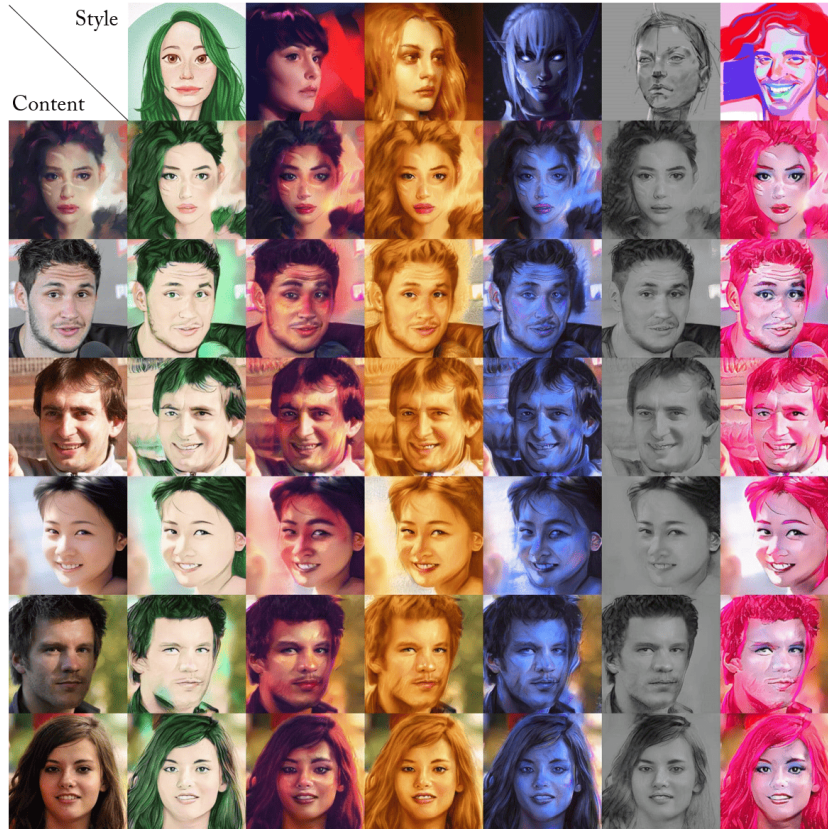
inducing style transfer and retaining semantic structure. The earlier the target StyleGAN2 layers that hyper-weight generators get fine-tuned for stylization, the stronger the emphasis on style is, therefore losing structure reconstruction quality. Conversely, the later the starting layer is, the less pronounced the stylization is, thereby focusing the training more on the semantic reconstruction quality and instead just performing a color transfer. Figure 3 shows a representative visualization of this effect, and Tbl. 3 shows quantitative metrics. The images further to the left represent a deeper fine-tuning of layers of style and losing structure. The images to the right show a more shallow fine-tuning for style, retaining more structure from the original content image. We targeted these specific layers, as these are toRGB layers in the target StyleGAN2 layer, which are known to more significantly affect color and texture.

Finally, we ablate the use of facial semantic region information in the model and training pipeline. We explore the effect of including the semantic regions as a conditioning signal in the pipeline (cond) and as a way to match region types in the patch co-occurrence discriminator (loss). We populate the results in Tbl4, showing the usefulness of these components.





**Fig. 3.** Representative visualization of the effect of choosing where in the target StyleGAN2 layers to start fine-tuning hyper-weight generators for. The left-most image is the reference content image, and the right-most is the reference style image. In the middle, from left to right, images represent visualizations for layers 3, 7, 13, 17, and 22.

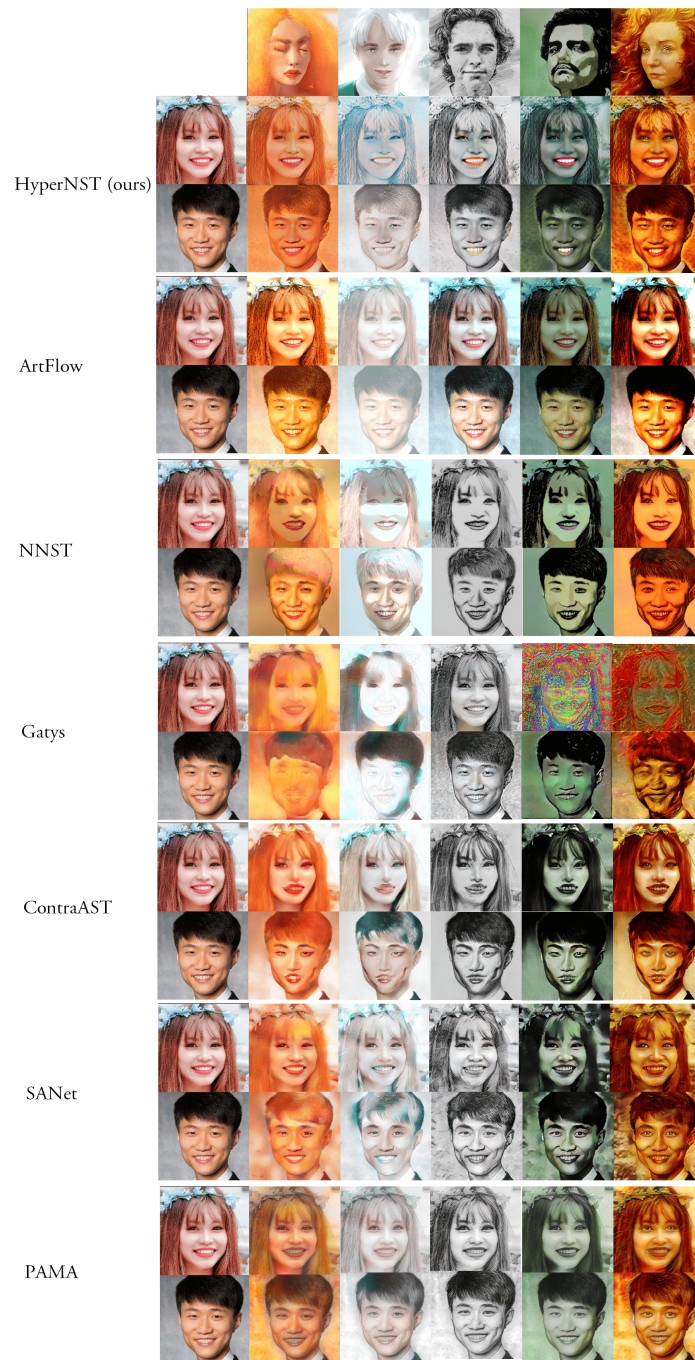


**Fig. 4.** Baseline comparison to Swapping Autoencoders (re-trained with AAHQ), the closest work - a fully per-weight-trained model.

#### 4.4 Downstream experiments

**Style interpolation** In Figure 6, we visualize a controllability aspect of the stylization process based on interpolations done in ALADIN style space. Basing the stylization upon a metric style embedding space, it is possible to perform smooth transitions between the style of a reference content image and a target style image. We also show in

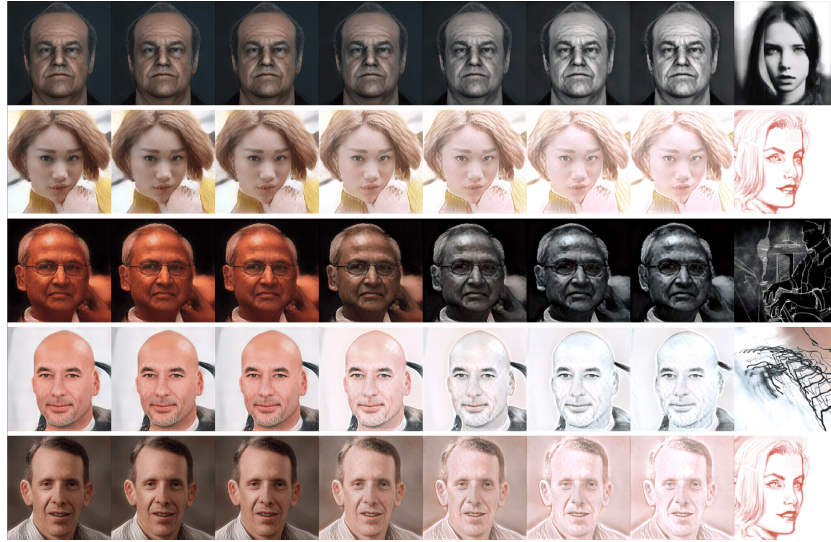




**Fig. 5.** Stylization comparison of HyperNST (1st pair), against Artflow (2nd), NNST (3rd), Gatys et.al (4th), ContraAST (4th), SANet (5th), and PAMA (6th)

Experiment	LPIPS (Alexnet)	LPIPS (VGG)	SIFID
Baseline (with both cond and loss)	0.00004	0.0017	2.279
No mask cond and no mask loss	0.00017	0.0024	2.892
No mask cond	0.00012	0.0025	2.848
No mask loss	0.00016	0.0026	3.438

**Table 4.** Ablation results for varying the facial semantic information used in the architecture.



**Fig. 6.** Visualization of style interpolation between a content image, and a target style image. Note in this example, we show that a non-portrait style image can also be used.

this example that it is possible to perform stylization using style images not containing portraits. In this case, the image-wide ALADIN code is used instead of semantically arranged per-feature ALADIN codes.

**Semantic controlability** One advantage over the SAE model is the editable method of GAN inversion that the HyperNST model uses. Coupled with InterFaceGAN [24], latent directions in the StyleGAN2 weight space can be used to expose semantic controls over portrait aspects such as age, pose, and facial expression. On top of the artistic style aspects of the style transfer capabilities of HyperNST, these semantic abilities offer a deeper level of artistic control over the editing and generation of portraits. Figure 7 visualizes some style transfer examples, where these latent directions are also used in tandem for more advanced control. We show the *smile*, *pose*, and *age* vectors.

## 5 Conclusion

We explored a novel application of a class of models, hypernetworks, to the task of style transfer. We explored the merits and drawbacks of such an approach and what is currently



**Fig. 7.** Semantic face/portrait editing capabilities. The top two example rows range across the *age* vector, the second two example rows show the *smile* vector, and the third shows the *pose* rotation vector.

possible. The hypernetwork approach is excellent at maintaining high quality semantic structure and allows for artistic controls, such as semantic editing and style-code-based stylization interpolation.

The main limitation of the current implementation of hyper-weight prediction is the very high VRAM consumption of such a model (22GB for a batch size of 1) during training, forcing weight deltas to affect the target generator only on a per-channel mean shifting basis. It is possible that the style transfer quality could be pushed further with a larger batch size. For example, the batch size used by Swapping Autoencoders is 16. This could be mitigated by exploring smaller alternatives to StyleGAN2, or by implementing GPU optimization tricks such as memory/compute tradeoff training approaches based on gradient checkpointing. Such reductions in VRAM use could also afford an increase in batch size, bringing the training more in line with previous works, primarily affecting discriminators. A more promising direction could be to more deeply analyse which layers can be omitted from training (i.e. need not be updated by the hyper network). In this paper, we focused our attention on the domain of human portraits. However, other domains would be equally important to study, further generalizing this novel approach to neural artistic stylization using more general semantic segmentation techniques or larger semantically diverse models such as StyleGAN XL, once hardware evolves to support hyper-network inference of larger models' weights.



## References

1. Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. Restyle: A residual-based stylegan encoder via iterative refinement. *CoRR*, abs/2104.02699, 2021. (Cited on page 3)
2. Yuval Alaluf, Omer Tov, Ron Mokady, Rinon Gal, and Amit H. Bermano. Hyperstyle: Stylegan inversion with hypernetworks for real image editing. *CoRR*, abs/2111.15666, 2021. (Cited on page 3, 6)
3. Jie An, Siyu Huang, Yibing Song, Dejing Dou, Wei Liu, and Jiebo Luo. Artflow: Unbiased image style transfer via reversible neural flows. *CoRR*, abs/2103.16877, 2021. (Cited on page 8, 9)
4. Haibo Chen, Lei Zhao, Zhizhong Wang, Zhang Hui Ming, Zhiwen Zuo, Ailin Li, Wei Xing, and Dongming Lu. Artistic style transfer with internal-external learning and contrastive learning. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. (Cited on page 8, 9)
5. Rinon Gal, Or Patashnik, Haggai Maron, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *CoRR*, abs/2108.00946, 2021. (Cited on page 1, 2, 3)
6. L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In *Proc. CVPR*, pages 2414–2423, 2016. (Cited on page 2, 8, 9)
7. David Ha, Andrew M. Dai, and Quoc V. Le. Hypernetworks. *CoRR*, abs/1609.09106, 2016. (Cited on page 3)
8. X. Huang and S. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proc. ICCV*, 2017. (Cited on page 2)
9. Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018. (Cited on page 2)
10. Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Proc. NeurIPS*, 2021. (Cited on page 2)
11. Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *CoRR*, abs/1812.04948, 2018. (Cited on page 6, 7)
12. T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *Proc. CVPR*, 2019. (Cited on page 2)
13. Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. *CoRR*, abs/1912.04958, 2019. (Cited on page 1, 2)
14. Nicholas Kolkin, Michal Kucera, Sylvain Paris, Daniel Sykora, Eli Shechtman, and Greg Shakhnarovich. Neural neighbor style transfer, 2022. (Cited on page 8, 9)
15. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’12, page 1097–1105, Red Hook, NY, USA, 2012. Curran Associates Inc. (Cited on page 8)
16. Mingcong Liu, Qiang Li, Zekui Qin, Guoxin Zhang, Pengfei Wan, and Wen Zheng. Blendgan: Implicitly GAN blending for arbitrary stylized face generation. *CoRR*, abs/2110.11728, 2021. (Cited on page 1, 6, 7)
17. Xuan Luo, Zhen Han, Lingfang Yang, and Lingling Zhang. Consistent style transfer. *CoRR*, abs/2201.02233, 2022. (Cited on page 8, 9)
18. Dae Young Park and Kwang Hee Lee. Arbitrary style transfer with style-attentional networks. *CoRR*, abs/1812.02342, 2018. (Cited on page 8, 9)
19. Taesung Park, Jun-Yan Zhu, Oliver Wang, Jingwan Lu, Eli Shechtman, Alexei A. Efros, and Richard Zhang. Swapping autoencoder for deep image manipulation. *CoRR*, abs/2007.00653, 2020. (Cited on page 2, 8, 9)

20. A. Radford, J. Wook Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021. (Cited on page 1)
21. Dan Ruta, Andrew Gilbert, Pranav Aggarwal, Naveen Marri, Ajinkya Kale, Jo Briggs, Chris Speed, Hailin Jin, Baldo Faieta, Alex Filipkowski, Zhe Lin, and John Collomosse. Stylelabel: Artistic style tagging and captioning, 2022. (Cited on page 2)
22. Dan Ruta, S. Motiian, B. Faieta, Z. Lin, H. Jin, A. Filipkowski, A. Gilbert, and J. Collomosse. Aladin: All layer adaptive instance normalization for fine-grained style similarity. *arXiv preprint arXiv:2103.09776*, 2021. (Cited on page 2, 4)
23. Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. Singan: Learning a generative model from a single natural image. *CoRR*, abs/1905.01164, 2019. (Cited on page 8)
24. Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *CoRR*, abs/2005.09635, 2020. (Cited on page 13)
25. K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition, 2014. (Cited on page 8)
26. Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *CoRR*, abs/2102.02766, 2021. (Cited on page 2)
27. Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *CoRR*, abs/1801.03924, 2018. (Cited on page 8)
28. Yinglin Zheng, Hao Yang, Ting Zhang, Jianmin Bao, Dongdong Chen, Yangyu Huang, Lu Yuan, Dong Chen, Ming Zeng, and Fang Wen. General facial representation learning in a visual-linguistic manner. *CoRR*, abs/2112.03109, 2021. (Cited on page 3, 4)