

What's in a Decade? Transforming Faces Through Time

Eric Ming Chen¹, Jin Sun², Apoorv Khandelwal³, Dani Lischinski⁴, Noah Snaveley¹, Hadar Averbuch-Elor⁵

¹Cornell University ²University of Georgia ³Brown University ⁴The Hebrew University of Jerusalem ⁵Tel Aviv University

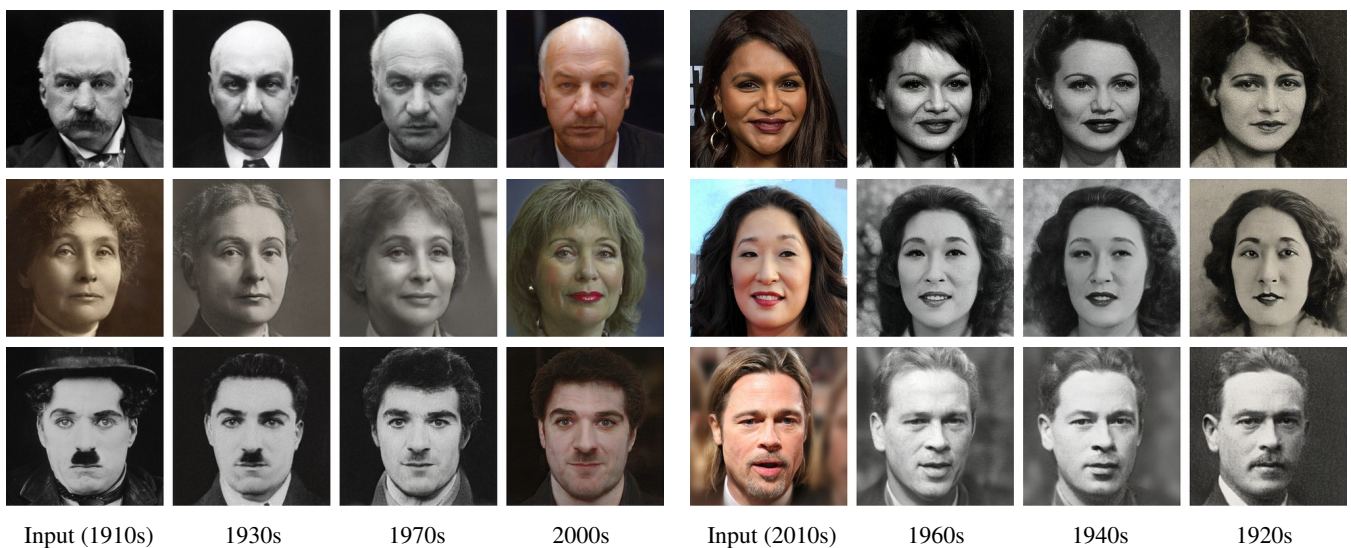


Figure 1: Given an input portrait image, our method generates plausible renditions of what that portrait might have looked like had it been taken in different decades spanning over a century. Our framework captures characteristic styles across decades while maintaining the person's identity. From top left to bottom right: J.P. Morgan, Mindy Kaling, Emmeline Pankhurst, Sandra Oh, Charlie Chaplin, Brad Pitt.

Abstract

How can one visually characterize photographs of people over time? In this work, we describe the Faces Through Time dataset, which contains over a thousand portrait images per decade from the 1880s to the present day. Using our new dataset, we devise a framework for resynthesizing portrait images across time, imagining how a portrait taken during a particular decade might have looked like had it been taken in other decades. Our framework optimizes a family of per-decade generators that reveal subtle changes that differentiate decades—such as different hairstyles or makeup—while maintaining the identity of the input portrait. Experiments show that our method can more effectively resynthesizing portraits across time compared to state-of-the-art image-to-image translation methods, as well as attribute-based and language-guided portrait editing models. Our code and data will be available at [facesthroughtime.github.io](https://github.com/facesthroughtime).

CCS Concepts

• **Computing methodologies** → **Image manipulation**; **Computer graphics**; **Computer vision**;

1. Introduction

What would photographs of ourselves look like if we were born fifty, sixty, or a hundred years ago? What would Charlie Chaplin look like if he were active in the 2020s instead of the 1920s? Such is the conceit of popular diversions like [old-time photography](#), where

we imagine ourselves as we might have looked in an anachronistic time period like the Roaring Twenties. However, while many methods for editing portraits have been devised recently on the basis of powerful generative techniques like StyleGAN [KLA*20; KAH*20; SYTZ20; AZMW21; WLS21; APC21; OSF*20], little attention has been paid to the problem of automatically translating

portrait imagery in time, while preserving other aspects of the person portrayed. This paper addresses this problem, producing results like those shown in Figure 1.

To simulate such a “time travel” effect, we must be able to model and apply the characteristic features of a certain era. Such features may include stylistic trends in clothing, hair, and makeup, as well as the imaging characteristics of the cameras and film of the day. In this way, translating imagery across time differs from standard portrait editing effects that typically manipulate well-defined semantic attributes (e.g., adding or removing a smile or modifying the subject’s age). Further, while large amounts of data capturing these attributes are readily available via datasets like Flickr-Faces-HQ (FFHQ) [KLA19], diverse and high-quality imagery spanning the history of photography is comparatively scarce.

In this work, we take a step towards transforming images of people across time, focusing on portraits. We introduce *Faces Through Time* (FTT), a dataset containing thousands of images spanning fourteen decades from the 1880s to the present day. *Faces Through Time* is derived from the massive public catalog of freely-licensed images and annotations available through the [Wikimedia Commons](#) project. The extensive biographic depth available on Wikimedia Commons, as well as its organization into time-based categories, enables associating images with accurate time labels. In comparison to previous time-stamped portrait datasets, FTT is sourced from a wider assortment of images capturing notable identities varying in age, nationality, pose, etc. In contrast, a well-known prior dataset called The Yearbook Dataset [GRS*15] only contains US yearbook photos. In our work, we demonstrate FTT’s applicability for synthesis tasks. More broadly, the dataset can allow for exploration of a variety of analysis tasks, such as estimating time from images, understanding photographic styles across time, and discovering fashion trends (see Section 2). Figure 2 shows random samples from five different decades from FTT.

To transform portraits across time, we build on the success of Generative Adversarial Networks (GANs) for synthesizing high-quality facial images. In particular, we finetune the popular StyleGAN2 [KLA*20; KAH*20] generator network (trained on FFHQ) on our dataset. However, rather than modeling the entire image distribution of our dataset using a single StyleGAN2 model, we train a separate model for each decade. We introduce a method to align and map a person’s image across the latent generative spaces of the fourteen different decades (see Figure 4). Furthermore, we discover a remarkable linearity in each model’s generator weights, allowing us to fine-tune images with vector arithmetic on the model weights. This sets our approach apart from the many prior works that search for editing directions within a single StyleGAN2 model, [SYTZ20; SZ21; WLS21; AZMW21; APC21; PWS*21]. We find that by using multiple StyleGAN2 models in this way, our method is more expressive than these existing approaches. In addition, our classes are separated in a useful way for style transfer.

We demonstrate results for a variety of individuals from different backgrounds and captured during different decades. We show that prior methods struggle on our problem setting, even when trained directly on our dataset. We also perform a quantitative evaluation, demonstrating that our transformed images are of high quality and resemble the target decades, while preserving the input identity.

Our approach enables an *Analysis by Synthesis* scheme that can reveal and visualize differences between portraits across time, enabling a study of fashion trends or other visual cultural elements (related to, for instance, a [New York Times article](#) that discusses how artists use digital tools to imagine what George Washington would have looked like had he lived today).

In summary, our key contributions are:

- *Faces Through Time*, a large, diverse and high-quality dataset that can serve a variety of computer vision and graphics tasks,
- a new task—transforming faces across time—and a method for performing this task that uses unique vectorial transformations to modify generator weights across different models,
- and quantitative and qualitative results that demonstrate that our method can successfully transform faces across time.

2. Related Work

2.1. Image analysis across time

While common image datasets such as ImageNet [DDS*09] and COCO [LMB*14] do not explicitly use time as an attribute, those that do show unique characteristics. Here we focus on image datasets that feature people.

Datasets. The Yearbook Dataset [GRS*15] is a collection of 37,921 front-facing portraits of American high school students from the 1900s to the 2010s. The authors design models to predict when a portrait is taken. They also analyze the prevalence of smiles, glasses, and hairstyles across different eras. The IMAGO Dataset [SAL*20] contains over 80,000 family photos from 1845 to 2009. Images are labeled in “socio-historical classes” such as *free-time* or *fashion*. Hsiao and Grauman [HG21] collect news articles and vintage photos and build a dataset that feature fashion trends in the 20th century. Our dataset contains portraits from a much wider age range (the Yearbook Dataset focuses on high school students), from diverse geographic areas (the IMAGO Dataset focuses on Italian families), and exhibiting rich variation in occupation and styles (not just fashion images).

Analysis. A standard task applicable to images with temporal information is to predict when they were taken, i.e., the date estimation problem. Müller-Budack et al. [MSE17] train two GoogLeNet models, one for classification and one for regression to predict the date of a photo. Salem et al. [SWZJ16] train different CNNs for date estimation using the face, torso, and patch of a portrait image. Other rich information can also be learned from temporal image collections. In StreetStyle and GeoStyle [MBS17; MMH*19], a worldwide set of images taken between 2013–2016 were analyzed to discover spatio-temporal fashion trends and events. In [HG21], topic models and clusterings are used to discover trends in news and vintage photos. Additionally, in [LEH13], a weakly-supervised approach is used to discover stylistic trends in cars over the decades. [LMC*15] learn trends over two centuries of architecture using Street View images. [MS14; LGZ*20] also disentangle pictures of buildings in a spatio-temporal way to see how structures have changed over time. Unlike prior works that focus on *analyzing* temporal characteristics in the data, we work on the more challenging task of *modifying* such characteristics.



Figure 2: Random samples from five decades in the Faces Through Time Dataset.

2.2. Portrait editing

Editing of face attributes has been extensively studied since before the deep learning era. A classic approach is to fit a 3D morphable model [BV99; EST*20] to a face image and edit attributes in the morphable model space. Other methods that draw on classic vision approaches includes Transfiguring Portraits [Kem16], which can render portraits in different styles via image search, selection, alignment, and blending. Given the recent success of StyleGAN (v1 [KLA19], v2 [KLA*20], and v3 [KAL*21]) in high quality face synthesis and editing, many works focus on editing portrait images using pre-trained StyleGAN models. In these frameworks, a photo is mapped into a code in one of StyleGANs latent spaces [XZY*21]. Feeding the StyleGAN generator with a modified latent code yields a modified portrait [SYTZ20; WLS21]. To find the latent code of an input image, one can directly optimize the code so that the StyleGAN can reconstruct the inputs [AQW19; WT20] or train a feed-forward network such as pSp [RAP*21] and e4e [TAN*21] that directly predicts the latent code. We adopt an optimization-based procedure to obtain better facial details.

Once a latent code of an image is obtained, portrait image editing can be done in the latent space with a pre-trained StyleGAN generator. Directions for change of viewpoint, aging, and lighting of faces can be found by PCA in the latent space [HHL20], or from facial attributes supervision [SGTZ20]. Shen and Zhou [SZ21] find that editing directions are encoded in the generators' weights and can be obtained by eigen decomposition. Collins et al. [CBPS20] perform local editing of portraits by mixing layers from reference and target images. Alternatively, the StyleGAN generator can also be modified for portrait editing. StyleGAN-nada [GPM*21] and StyleCLIP [PWS*21] use CLIP [RKH*21] to guide editing on images with target attributes. Toonify [PA20] uses layer-swapping to obtain a new generator from models in different domains. Similar to StyleAlign [WNSL21], we obtain a family of generators by finetuning a common parent model on different decades. The style change of a face is achieved by obtaining the latent code of the input image and feeding it into a modified target StyleGAN generator using PTI

[RMBC21]. Our method is conceptually simple and doesn't require exploring the latent space.

Age editing on portraits is related to our task since both involve modifying the temporal aspects of an image. In the work of Or-EI, et al. [OSF*20], age is represented as a latent code and applied to the decoder network of the face generator. An identity loss is used to preserve identities across ages. Alaluf et al. [APC21] design an age encoder that takes a portrait and a target age, produces a style code modification on pSp-coded styles and generates a new portrait with the target age. In contrast, our editing aims to change the decade a photo was taken without altering the subject's age.

GANs can also be used to recolor historic photographs [ZZI*17; Ant19; LZY*21]. In particular, Time-Travel Rephotography [LZY*21] uses a StyleGAN2 model to project historic photos into a latent space of modern high-resolution color photos with modern imaging characteristics. Rather than focusing solely on low-level characteristics like color, our method alters a diverse collection of visual attributes, such as facial hair and make-up styles. Moreover, our method can transform images across a wide range of decades, instead of learning binary transformations between "old" and "modern" as in [LZY*21].

3. The Faces Through Time Dataset

Our *Faces Through Time* (FTT) dataset features 26,247 images of notable people from the 19th to 21st centuries, with roughly 1,900 images per decade on average. It is sourced from Wikimedia Commons (WC), a crowdsourced and open-licensed collection of 50M images.

We automatically curate data from WC to construct FTT (Figure 2) as follows: (1) The "People by name" category on WC contains 407K distinct people identities. We query each identity's hierarchy of people-centric subcategories (similar to [CKA*21]) and organize retrieved images by identity. (2) We use a Faster R-CNN model [RHGS15; JL17; RR17] trained on the WIDER Face dataset

[YLLT16] as a face detector. For each detected face, 68 facial landmarks are found using the Deep Alignment Network [KNT17], and alignments are applied as in the FFHQ dataset [KLA19] given these landmarks. (3) We devise a clustering method based on clique-finding in face similarity graphs to group faces by identities (see appendix). This resolves ambiguities in photos that feature multiple people. (4) We gather additional samples without biographic information from the “19th Century Portrait Photographs of Women” and “20th Century Portrait Photographs of Women” categories. These make up about 15% of the dataset.

We leverage image metadata, identity labels, and biographic information available in WC to further assist in balancing and filtering our data. We discard any photos without time labels or taken before 1880, and sample a subset of 3,000 faces each for the 2000s and 2010s decades (which tend to feature many more images than other decades) to maintain a roughly balanced distribution of images across decades. For images where the identity is known, we further filter by only keeping images where the identity is between 18 and 80 years old (comparing the image timestamp and identity’s birth year). We also estimate face pose using Hopenet [RCR18] and remove images with yaw or pitch greater than 30 degrees. After these automated collection and filtering steps, we manually inspected the entire dataset and removed images with clearly incorrect dates, images that were not cropped properly, images that were duplicates of other identities, and images featuring objectionable content. This resulted in a removal of 6% of the assembled data.

The total number of samples from each decade in our curated dataset, demographic and other biographic distributions, and further implementation details can be found in our supplementary material. We create train and test splits by randomly selecting 100 images per decade as a test set, with the remaining images used for training. Samples from the dataset are shown in Figure 2.

As detailed in Section 2, the Yearbook dataset [GRS*15] is another dataset that spans multiple decades and could potentially allow for transforming faces across time. However, it is grayscale only, contains one age group, and critically, its images are low resolution (186×171). In the supplementary material, we show that high-quality synthesized images cannot be obtained using this dataset, further highlighting the benefit of our FTT dataset.

4. Transforming Faces Across Time

Given a portrait image from a particular decade, our goal is to predict what the same person might look like across various decades ranging from 1880 to 2010. The key challenges are: (1) maintaining the identity of the person across time, while (2) ensuring the result fits the natural distribution of images of the target decade in terms of style and other characteristics. We present a novel two-stage approach that addresses these challenges. Figure 3 shows an overview of our approach.

First, rather than training a single generative model that covers all decades (e.g., [OSF*20]), we train a family of StyleGAN models, one for each decade (Section 4.1, Figure 4). These are obtained by fine-tuning the same *parent* model, resulting in a set of *child* models whose latent spaces are roughly *aligned* [WNSL21], as described in Section 4.1. The alignment ensures that providing the

same latent code to different models results in portraits with similar high-level properties, such as pose. At the same time, the resulting images exhibit the unique characteristics of each decade. Given a real portrait from a particular decade, it is first inverted into the latent space of the corresponding model, and the resulting latent code can then be fed into the model of any desired target decade. Our approach makes it unnecessary to search for editing directions in the latent space (e.g., [SYTZ20; HHL20]).

Next, to better fit the identity of the input individual, we apply single-image finetuning of the family of per-decade StyleGAN generators (Section Section 4.2). Specifically, we introduce Transferable Model Tuning (TMT), a modified PTI (Pivotal Tuning Inversion) [RMBC21] procedure, to obtain an adjustment for the weights of the source decade generator, and apply the resulting adjustment to the target generator(s). This input-specific adjustment is done in the generator’s parameter space, enabling us to better preserve the input individual’s identity, while maintaining the style and characteristics of the target decade. We now describe these two stages in more detail.

4.1. Learning coarsely-aligned decade models

We are interested in learning a family of StyleGAN2-ADA [KAH*20] generators $\{G_t\}$, $t \in \{1880, 1890, \dots, 2010\}$ each of which maps a latent vector $w \in \mathcal{W}$ to an RGB image. For each decade, we finetune a separate StyleGAN model with weights initialized from an FFHQ-pretrained model. We call the FFHQ-pretrained model the *parent model* G_p , and the finetuned network for decade t the *child model* G_t . Consistent with the findings in prior work [WNSL21], we observe that the collection of generators $\{G_t\}$, $t \in \{1880, \dots, 2010\}$ exhibits semantic alignment of faces generated from the same latent code w : they share similar face poses and shapes. However, various fine facial characteristics such as eyes and noses, which are important for recognizing a person, often drift from one another (as evident in Figures 4 and 8).

To better preserve identity across decades when finetuning each child model, to the standard StyleGAN objective \mathcal{L}_{GAN} we add an identity loss. Specifically, we measure the cosine similarity between ArcFace [DGXZ19] embeddings of images generated by G_p and by G_t :

$$\mathcal{L}_{\text{ID}} = 1 - \cos_sim(\text{ArcFace}(G_p(w)), \text{ArcFace}(G_t^B(w))). \quad (1)$$

A similar loss is used in [RAP*21]. However, since the ArcFace model was only trained on modern day images, we found that this raw identity loss performed poorly on historical images, due to the domain gap. To solve this issue, we use a blended version $G_t^B(w)$ instead of the original $G_t(w)$. We create $G_t^B(w)$ using layer swapping [PA20] to mix $G_p(w)$ and $G_t(w)$ at different spatial resolutions: we combine the *coarse* layers of the child model $G_t(w)$ with the *fine* layers of the parent model $G_p(w)$. By doing so, we “condition” our input image for the identity loss by making its colors more similar to the image generated by the parent, and thus more similar to the distribution of the images used to train the ArcFace model. Figure 3 (left) shows that the blended (middle) image retains the structure of the 1900s image, but its colors better resemble those of a modern day photo. In addition, this technique restricts the identity loss to focus on layers which generally control head shape, position, and

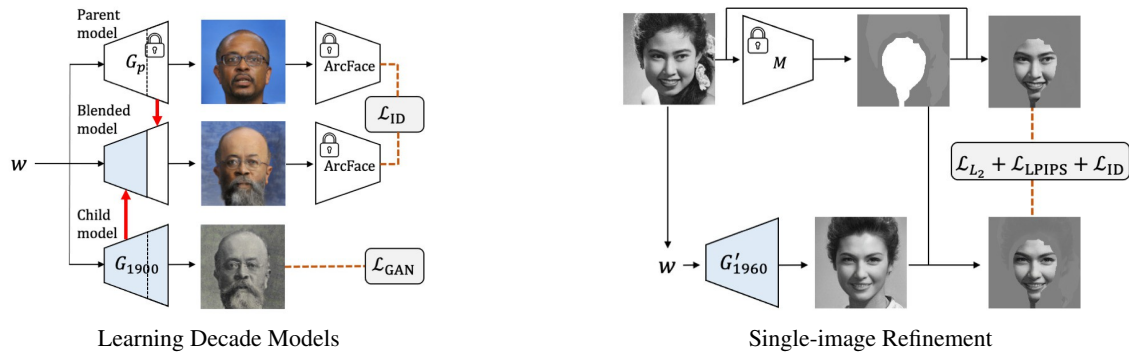


Figure 3: Overview of our method. Left: We first train a family of StyleGAN models, one for each decade, using adversarial losses and an identity loss on a blended face, which resembles the parent model in its colors. Right: Afterwards, each real image is projected onto a vector w on the decade manifold (1960 in the example above). We learn a refined generator G'_t and transfer the learned offset to all models (this process is visualized in Figure 5). To better encourage the refined model to preserve facial details, we mask the input image and apply all losses in a weighted manner (further described in the text).

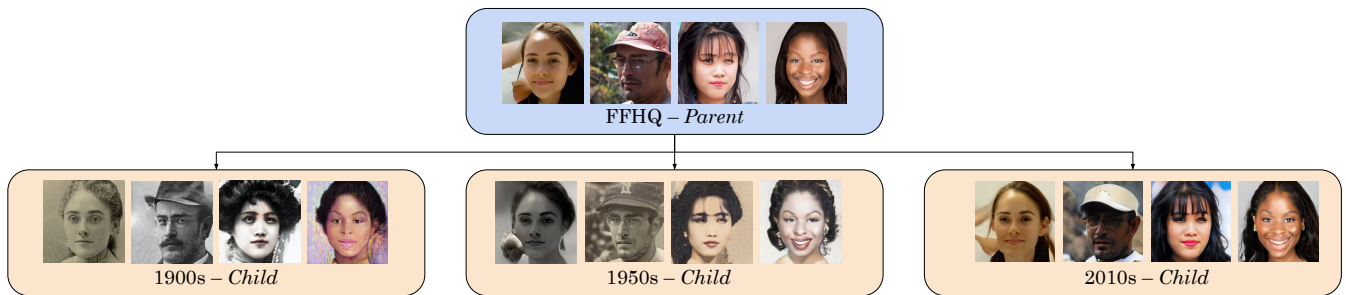


Figure 4: We finetune a family of decade generators (child models) from an FFHQ-trained parent model. While each generator captures unique styles, the generated images from the same latent code are aligned in terms of high-level properties such as pose.

identity. Note that this blended image is only used to compute the loss, and not in the transformed results.

4.2. Single-image refinement and transferable model tuning

As previously described, we first train the family of aligned StyleGAN models with randomly sampled latent codes from \mathcal{W} and our collections of real per-decade images. Given these coarsely-aligned per-decade models, we are given a single real face image as input and aim to generate a set of faces across various decades. In order to better preserve the identity of the input image across these decades, we introduce Transferable Model Tuning (TMT). TMT is inspired by PTI [RMBC21], which is a procedure for optimizing a model's parameters to better fit to an input image after GAN inversion. TMT extends PTI from a single generator model to a family of models. Our TMT procedure produces a set of face images where the identity is preserved in the presence of changing style over time (Figure 5).

Specifically, given an input image x from decade t , we first obtain its latent code $w \in \mathcal{W}$ using the projection method from StyleGAN2: $w = \arg \min_w \|x - G_t(w; \theta_t)\|$, where $\theta_t \in \Theta$ is the vector of all parameters in G_t . We only work with child models in this stage. As the first step of TMT, we fix the obtained latent code w

and optimize over θ_t to obtain a new model G'_t with parameters θ'_t (Figure 3, right).

$$\theta'_t = \arg \min_{\theta_t} \|x - G_t(w; \theta_t)\|. \quad (2)$$

Tuning in the parameter space Θ of generators, instead of only working in the latent space \mathcal{W} as in previous work [NBLC20], allows us to better fit the facial details of the input individual, such as eyes and expression. Treating θ_t and θ'_t as vectors in the parameter space Θ , this tuning can be thought of as applying an offset $\Delta\theta = \theta'_t - \theta_t$ to the original model.

We found that this $\Delta\theta$ offset is surprisingly transferable from one TMT-tuned generator to all other decade generators in the parameter space Θ . Concretely, to obtain the style-transformed face of x in any target decade d , we simply apply:

$$x_d = G_d(w; \theta_d + \Delta\theta), \quad (3)$$

where G_d is the generator for decade d and θ_d is the vector of all its parameters. We visualize the learned TMT offsets for two different input images in Figure 5. As illustrated in the figure, these offsets greatly improve the identity preservation of synthesized portraits.

Intuitively, $\Delta\theta$ ‘refines’ the parameters of the generator family to the single input face image to reconstruct better facial details.



Figure 5: Visualization of TMT offsets. We obtain offset vectors $\Delta\theta_1$ (for the image in row 1) and $\Delta\theta_2$ (row 2) for the weights of the source decade generator and apply it to every target decade generator. On the left we use PCA to visualize the convolutional parameters for all target generators in 2D. Each dot represents the weights of a single generator, colored according to decade, and with edges connecting adjacent decades. We illustrate the offset vectors optimized for the two input images (colored in gray and red) and the corresponding transformed images for three different decades. For each decade we show images before (left) and after (right) applying TMT. Adding these offsets has the effect of improving identity preservation.

We hypothesize that the found $\Delta\theta$ offsets mainly focus on improving identities. Since the coarsely-aligned family of generators share similar weights that are responsible for identity-related features, those offsets are easily transferable. While most prior works, e.g., [SYTZ20; HHLP20; WLS21] modify images using linear directions in various *latent spaces* of StyleGAN, we are the first to apply a linear offset in the generator *parameter space*, to a collection of generators. We hope our work will inspire future investigations on understanding the linear properties in GAN’s parameter spaces. An analysis of the effects of applying TMT to a family of models can be found in the supplementary material.

As demonstrated in Figure 3 (right), to focus the loss computation on facial details instead of hair and background, we apply masks to images before calculating the losses. We use a DeepLab segmentation network [CZP*18] trained on CelebAMask-HQ photos [OSF*20; LLWL20]. Empirically we determine it is best to apply a weight of 1.0 on the face, 0.1 on the hair and 0.0 elsewhere. We put a small weight on the hair to accurately reconstruct it, as it does contribute to the image’s stylization. However, we do not want to prioritize it over facial features. In addition, we find that it is best to keep StyleGAN’s ToRGB layers frozen. Otherwise, color artifacts are introduced into the generators. We follow the objectives introduced in [RMBC21] and minimize a perceptual loss (\mathcal{L}_{LPIPS}), and a reconstruction loss (\mathcal{L}_{L2}). In addition, we add another identity loss (\mathcal{L}_{id}) to further enhance identity preservation for the generated images.

Using our two-stage approach, for each portrait, we obtain a set of faces that maintain the identity as well as demonstrate diverse styles across decades.

5. Results and Evaluation

We conduct extensive experiments on the *Faces Through Time* test set. We compare our approach to several state-of-the-art techniques across a variety of metrics that quantify how well transformed images depict their associated target decades and to what extent the identity is preserved. We also present an ablation study to examine the impact of the different components of our approach. Additional

uncurated results and visualizations of the full test set (1,400 individuals) are in the supplementary material.

5.1. State-of-the-art Image Editing Alternatives

As no prior works directly address our task, we adapt commonly used image editing models to our setting and perform comprehensive comparisons.

Image-to-image translation. Unpaired image-to-image translation models learn a mapping between two or more domains. We train a StarGAN v2 [CUYH20] model on our dataset, where decades are domains. Because the quantitative metrics between our model and StarGAN are so similar, we show a detailed qualitative comparison between our model and StarGAN in the supplementary material, over a set of individuals balanced by ethnicity and gender. We see that StarGAN has a poor understanding of skin tone and overall identity, which is critical for our real-world applications. Results on another model, DRIT++ [LTM*20], are in the supplementary material.

Attribute-based editing. We consider each decade as a facial attribute and compare against recent works performing attribute-based edits. While many attributes in prior work are binary [SYTZ20; HHLP20], our decade attribute has multiple classes. By comparison, age-based transformation is more similar to our problem setting, as age is often broken up into K bins [OSF*20]. We compare against SAM [APC21], a recent age-based transformation technique that also operates in the StyleGAN space.

Language-guided editing. Weakly supervised methods that leverage powerful image-text representations (e.g. CLIP [RKH*21]) have demonstrated impressive performance in portraits editing. We compare against: (i) StyleCLIP [PWS*21], which learns latent directions in StyleGAN’s $\mathcal{W}+$ space for a given text prompt, and (ii) StyleGAN-nada (StyleNADA) [GPM*21], which modifies the weights of the StyleGAN generator based on textual inputs. For both models, we used the text prompt “A person from the [XYZ]0s”, where [XYZ]0 is one of FTT’s 14 decades (1880-2010).

	Method	FID↓	KMMD ↓	DCA ₀ ↑	DCA ₁ ↑	DCA ₂ ↑	ID _{acc} ↑
FFHQ	StyleCLIP	254.39	1.87	0.08	0.18	0.36	0.99
	StyleNADA	312.06	2.03	0.10	0.30	0.38	0.96
	Ours	69.46	0.43	0.50	0.81	0.91	0.93
FTT	StarGAN v2	68.05	0.40	0.38	0.75	0.89	0.97
	SAM	96.52	0.72	0.51*	0.85*	0.89*	1.00
	StyleCLIP	108.25	0.85	0.07	0.21	0.36	1.00
	Ours	66.98	0.40	0.47	0.78	0.90	0.99

Table 1: Quantitative Evaluation. We compare performance against SOTA techniques on FFHQ (top three rows) and on our test set (bottom four rows). Our method outperforms others in terms of most metrics. *Note that SAM uses the decade classifier during training and therefore the DCA metric is skewed in this case, as we further detail in the text.

For StyleCLIP, we use models trained on both FFHQ and FTT. Because StyleGAN-nada is designed for out-of-domain changes, we experiment with how well it can modify the generator from the FFHQ space to various decades in our dataset. We use 100 FFHQ photos. We compare these two baselines to how well our model can transform FFHQ images to each of the 14 decades.

Conditioning on a single model vs. multiple models. We train one model for each decade as done in StyleAlign [WNSL21] because it is difficult to disentangle each decade’s style in a single model. The modifications we made to SAM [APC21], such as adding a decade classifier, can be considered a straightforward way to condition decade labels on a single model’s latent space. In addition, StyleCLIP [PWS*21] also uses a single finetuned model to transform images within a GAN’s latent space. Our results show that these single-model methods struggle to understand the style of each decade, as compared to our multiple-model approach.

Time Travel Rephotography. Similar to our method, Time Travel Rephotography [LZY*21] is designed to imagine what a historical person would have looked like in another time period (and can only perform a transformation to modern imagery). However, the authors focus primarily on image restoration as opposed to changing a person’s style or fashion. We see that Time Travel Rephotography improves camera quality and lighting instead of changing hairstyles and other stylistic features as our model does. Technically, [LZY*21] inverts their photos into a pretrained StyleGAN on FFHQ instead of training new models on data for each decade. We provide visual comparisons between our model and Time Travel Rephotography in the supplementary material.

5.2. Metrics

Visual quality. We use the standard FID [HRU*17; Sei20] metric as well as the Kernel Mean Maximum Discrepancy Distance (KMMD) [WGB*20] metric. As image quality varies across decades, we compute scores between real and edited portraits separately for each decade and then average over all decades. Because FID can be unstable on smaller datasets [CF19], similar to prior work [NH19; WGB*20], we measure KMMD [WGB*20] on Inception [SVI*16] features. Experimentally, we find that these two

scores are highly sensitive to an image’s background. Therefore, we compute the scores on images of size 256×256 cropped to 160×160 pixels.

Decade style. We evaluate how well the generated samples capture the style of the target decade using a EfficientNetB0 classifier [TL19] that we trained separately. Using the classifier, we define the Decade Classification Accuracy (DCA). We follow prior works [GJ09] and report three metrics: DCA₀, DCA₁ and DCA₂, where DCA_p measure the accuracy within a tolerance of p decades.

Identity preservation. We use the Amazon Rekognition service to measure how well a person’s identity has been preserved in generated portraits. Their COMPAREFACES operation outputs a similarity score between two faces. As a metric, we report ID_{acc} – the fraction of successful identity comparisons. We consider a comparison to be successful if its similarity score is above a certain threshold (set empirically to 1.0).

5.3. Results

We present test set performance for all methods in Table 1 and qualitative results in Figure 6. Results on the *full* test set (1400 samples) are provided using an interactive viewer as part of our supplementary material. We present additional results on StyleGAN-nada’s effect on FFHQ input images in Figure 7. We see that our method performs well numerically in terms of all metrics. Although StarGAN also performs well in FID and KMMD, the style changes by StarGAN are mostly about color; there are few changes to makeup, hair, and beards, whereas our model can perform such changes. Our method also has fewer artifacts. As illustrated in Figure 7, modifications from StyleGAN-nada are more caricature-like than realistic. As a result, StyleGAN-nada performs poorly with respect to FID, KMMD, and DCA metrics. For StyleCLIP and SAM, the identity preservation is near 1.0 because their changes are generally limited to color. Nonetheless, our model still successfully matches input and transformed images in most cases (for 93% and 99% of samples, for FFHQ and FTT images, respectively), while generating significant changes in styles. While SAM performs the best with respect to DCA, we believe this is advantaged because SAM used the same classifier during training as a loss function. In fact, in Figure 6, SAM demonstrates little change across decades. We suspect that the classifier is leading SAM to overfit to noise, instead of truly changing an image’s style.

From our results we can discern interesting details that provide insight into style trends across time. For example, as illustrated in the top left example in Figure 6, we see that the individual adopts a bob haircut, one popularized by Irene Castle, and strongly associated with the flapper culture of the Roaring Twenties. Later on, we notice more contemporary hair styles. Finally, in the 2010s, we see that women tend to adopt longer hairstyles. Despite these generalizations, the portraits remain well-conditioned on the input, reflecting an individual’s identity and aspects of their personal style. For instance, the individual on the left with the long mustache maintains facial hair across the decade transformations, although in very different styles. In the bottom left example, we observe glasses that change style over time. Not only does our model generate realistic transformations, but also captures the nostalgia of various time periods.

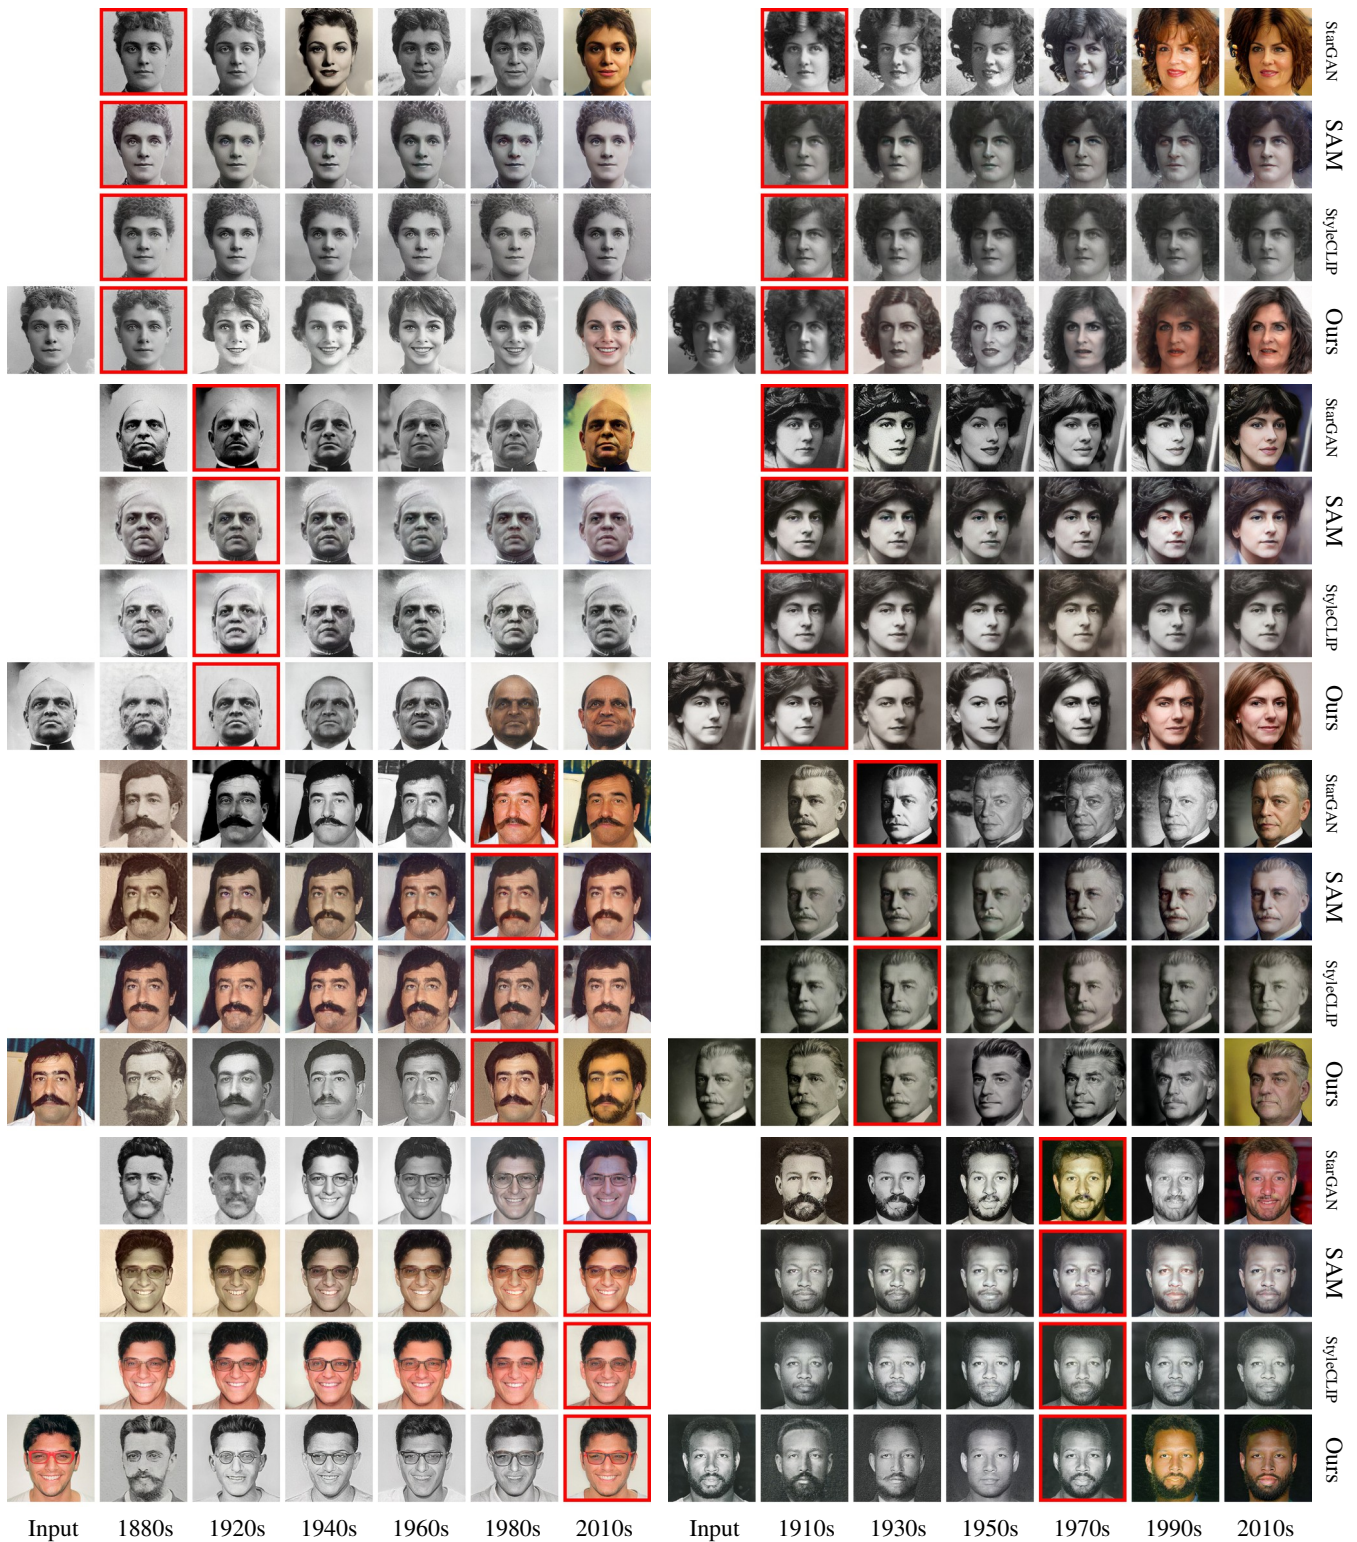


Figure 6: Qualitative Results. Above we compare results generated by baselines and our technique. The red box indicates the inversion of the original input. We observe that our approach allows for significant changes across time while best preserving the input identity. While SAM and StarGAN are able to stylize images, these changes are mostly limited to color. StyleCLIP struggles to generate meaningful changes. Please refer to the supplementary material for qualitative results on the full test set.




Figure 7: StyleGAN-nada Results. Although StyleGAN-nada [GPM*21] produces some style changes across decades, all images are transformed similarly. For example, all 1980s portraits adopt a frizzy hairstyle.



Figure 8: Ablations. We see significant improvement in terms of identity preservation after adding the identity loss and TMT during training. The masking procedure alleviates artifacts caused by other regions in the image (such as the hat in the example above) by focusing the model’s attention on facial details and allowing for larger modifications in other regions. See Table 2 for descriptions of the ablation labels.

5.4. Ablations

We present ablations on components in our approach in Table 2 and Figure 8. Specifically, we train five ablated models: (1) without the identity loss (for learning decade models), (2) without the blended image obtained using layer swapping, (3) without TMT, (4) without the identity loss (during TMT), and (5) without masking the images. The first ablation is akin to finding the nearest neighbor of an image in the StyleGAN \mathcal{W} space using latent projection. As shown in the first row of the table, our baseline method already captures a decade’s style well. However, the images are not aligned with respect to an input’s identity, which is reflected in its low ID_{acc} score. As a result, $\mathcal{L}_{id}^{(I)}$, $\mathcal{L}_{id}^{(II)}$, and TMT are necessary for identity

preservation. In addition, we find that using masks during TMT reduces artifacts in generated images as it allows for an accurate inversion in the facial region and larger modifications in other regions in the image. Empirically, we notice that this reduces noise and improves decade classification. While FID, KMMD and DCA scores remain similar across the ablations, our full model shows strong improvement in ID_{acc} , which is the main objective of $\mathcal{L}_{id}^{(I)}$ and our proposed TMT stage. We also experimentally find that $\mathcal{W}+$ spaces [RAP*21] are less well aligned than the \mathcal{W} space. More details are in the supplementary material.

$\mathcal{L}_{id}^{(I)}$	LS	TMT	$\mathcal{L}_{id}^{(II)}$	Mask	FID	KMMD	DCA ₀	DCA ₁	DCA ₂	ID_{acc}
×	×	×	×	×	69.51	0.45	0.49	0.79	0.92	0.61
✓	×	×	×	×	69.36	0.47	0.51	0.82	0.92	0.63
✓	×	×	×	×	68.18	0.45	0.50	0.82	0.92	0.72
✓	✓	✓	×	×	67.32	0.39	0.46	0.78	0.89	0.95
✓	✓	✓	✓	×	67.08	0.38	0.46	0.77	0.89	0.99
✓	✓	✓	✓	✓	66.98	0.40	0.47	0.78	0.90	0.99

Table 2: Ablation study evaluating the effect of the identity loss while learning decade models ($\mathcal{L}_{id}^{(I)}$) and during TMT ($\mathcal{L}_{id}^{(II)}$), using a blended image with layer swapping (LS), TMT, and masking the images during TMT (Mask).

6. Ethical Discussion

Face datasets—and the tasks that they enable, such as face recognition—have been subject to increasing scrutiny and recent work has shed light on the potential harms of such data and tasks [DRD*20; GNH19]. With awareness of these issues, our dataset was constructed with attention to ethical questions. The images in our dataset are freely-licensed and provided through a public catalog. We will include the source and license for each image in our dataset. As part of our terms of use, we will only provide our dataset for academic use under a restrictive license. Furthermore, our dataset does not contain identity information (and only includes one face per identity), and therefore cannot readily be used for facial recognition. Nonetheless, our dataset does inherit biases that are present in Wikimedia Commons. For instance, the data is gender imbalanced, containing a ratio of roughly 2 : 1 male to female samples (according to the binary gender labels available on Wikimedia Commons, which are annotated by Wikipedia contributors). While such biases can be mitigated by balancing the data for training and evaluation purposes, we plan to continue gathering more diverse data to address this underlying bias in the data. For additional details on various features of our dataset, please refer to the accompanying datasheet [GMV*21].

There are also ethical considerations relating to the risks of using portrait editing for misinformation. Our task is perhaps less sensitive in this regard, since our explicit goal is to create fanciful imagery that is clearly anachronistic. That said, any results from such technology should be clearly labeled as imagery that has been modified.

7. Conclusion

We present a dataset and method for transforming portrait images across time. Our *Faces Through Time* dataset spans diverse geographical areas, age groups, and styles, allowing one to capture the

essence of each decade via generative models. By learning a family of generators and efficient tuning offsets, our two-stage approach allows for significant style changes in portraits, while still preserving the appearance of the input identity. Our evaluation shows that our approach outperforms state-of-the-art face editing methods. It also reveals interesting style trends existing in various decades. However, our method still has limitations. As with any data-driven technique, our results are affected by biases that exist in the data. For instance, females with short hair are less common at the beginning of the 20th century, which may yield gender inconsistencies when transforming a short-haired modern female face to these early decades, including unexpected changes in visual features often associated with gender. In the future, we plan to explore methods that can improve consistency, perhaps by devising a way to jointly optimize models for different decades that better enforces consistency among them. Finally, we envision that future uses of our data could go beyond the synthesis tasks we consider in our work, and explore the combination of both analysis and synthesis.

8. Acknowledgements

This work was supported in part by the National Science Foundation (IIS-2008313).

References

- [Ant19] ANTIC, JASON. *A deep learning based project for colorizing and restoring old images (and video!)* 2019. URL: <https://github.com/jantic/DeOldify3>.
- [APC21] ALALUF, YUVAL, PATASHNIK, OR, and COHEN-OR, DANIEL. “Only a Matter of Style: Age Transformation Using a Style-Based Regression Model”. *ACM Trans. Graph.* 40.4 (2021) 1–3, 6, 7.
- [AQW19] ABDAL, RAMEEN, QIN, YIPENG, and WONKA, PETER. “Image2StyleGAN: How to Embed Images Into the StyleGAN Latent Space?”. *Int. Conf. Comput. Vis.* 2019, 4431–4440 3.
- [AZMW21] ABDAL, RAMEEN, ZHU, PEIHAO, MITRA, NILOY J., and WONKA, PETER. “StyleFlow: Attribute-Conditioned Exploration of StyleGAN-Generated Images Using Conditional Continuous Normalizing Flows”. *ACM Trans. Graph.* 40.3 (May 2021). ISSN: 0730-0301. DOI: [10.1145/3447648](https://doi.org/10.1145/3447648). URL: <https://doi.org/10.1145/3447648> 1, 2.
- [BV99] BLANZ, VOLKER and VETTER, THOMAS. “A morphable model for the synthesis of 3D faces”. *Proceedings of the 26th annual conference on Computer graphics and interactive techniques.* 1999, 187–194 3.
- [CBPS20] COLLINS, EDO, BALA, RAJA, PRICE, BOB, and SÜSSTRUNK, SABINE. “Editing in Style: Uncovering the Local Semantics of GANs”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2020 3.
- [CF19] CHONG, MIN JIN and FORSYTH, DAVID. “Effectively Unbiased FID and Inception Score and where to find them”. *arXiv preprint arXiv:1911.07023* (2019) 7.
- [CKA*21] CUI, YUQING, KHANDELWAL, APOORV, ARTZI, YOAV, et al. “Who’s Waldo? Linking People Across Text and Images”. *Int. Conf. Comput. Vis.* Oct. 2021, 1374–1384 3.
- [CUYH20] CHOI, YUNJEY, UH, YOUNGJUNG, YOO, JAEJUN, and HA, JUNG-WOO. “StarGAN v2: Diverse Image Synthesis for Multiple Domains”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2020, 8185–8194 6.
- [CZP*18] CHEN, LIANG-CHIEH, ZHU, YUKUN, PAPANDREOU, GEORGE, et al. “Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation”. *Eur. Conf. Comput. Vis.* 2018 6.
- [DDS*09] DENG, JIA, DONG, WEI, SOCHER, RICHARD, et al. “ImageNet: A large-scale hierarchical image database.” *IEEE Conf. Comput. Vis. Pattern Recog.* 2009, 248–255. ISBN: 978-1-4244-3992-8 2.
- [DGXZ19] DENG, JIANKANG, GUO, JIA, XUE, NIANNAN, and ZAFEIRIOU, STEFANOS. “Arcface: Additive angular margin loss for deep face recognition”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2019, 4690–4699 4.
- [DRD*20] DROZDOWSKI, PAWEŁ, RATHGEB, CHRISTIAN, DANTCHEVA, ANTITZA, et al. “Demographic Bias in Biometrics: A Survey on an Emerging Challenge”. *IEEE Transactions on Technology and Society* 1.2 (June 2020), 89–103. ISSN: 2637-6415. DOI: [10.1109/TTS.2020.2992344](https://doi.org/10.1109/TTS.2020.2992344). URL: <http://dx.doi.org/10.1109/TTS.2020.2992344> 9.
- [EST*20] EGGER, BERNHARD, SMITH, WILLIAM AP, TEWARI, AYUSH, et al. “3d morphable face models—past, present, and future”. *ACM Trans. Graph.* 39.5 (2020), 1–38 3.
- [GJ09] GAUDETTE, LISA and JAPKOWICZ, NATHALIE. “Evaluation methods for ordinal classification”. *Canadian Conference on Artificial Intelligence.* Springer. 2009, 207–210 7.
- [GMV*21] GEBRU, TIMNIT, MORGENSTERN, JAMIE, VECCHIONE, BRIANA, et al. “Datasheets for datasets”. *Communications of the ACM* 64.12 (2021), 86–92 9.
- [GNH19] GROTH, PATRICK, NGAN, MEI, and HANAOKA, KAYEE. *Face recognition vendor test (fvrt): Part 3, demographic effects.* National Institute of Standards and Technology, 2019 9.
- [GPM*21] GAL, RINON, PATASHNIK, OR, MARON, HAGGAI, et al. *StyleGAN-NADA: CLIP-Guided Domain Adaptation of Image Generators.* 2021. arXiv: [2108.00946](https://arxiv.org/abs/2108.00946) [cs.CV] 3, 6, 9.
- [GRS*15] GINOSAR, SHIRY, RAKELLY, KATE, SACHS, SARAH, et al. “A century of portraits: A visual historical record of american high school yearbooks”. *Proceedings of the IEEE International Conference on Computer Vision Workshops.* 2015, 1–7 2, 4.
- [HG21] HSIAO, WEI-LIN and GRAUMAN, KRISTEN. “From Culture to Clothing: Discovering the World Events Behind A Century of Fashion Images”. *Int. Conf. Comput. Vis.* 2021 2.
- [HHL20] HÄRKÖNEN, ERIK, HERTZMANN, AARON, LEHTINEN, JAAKKO, and PARIS, SYLVAIN. “Ganspace: Discovering interpretable gan controls”. *arXiv preprint arXiv:2004.02546* (2020) 3, 4, 6.
- [HRU*17] HEUSEL, MARTIN, RAMSAUER, HUBERT, UNTERTHINER, THOMAS, et al. “GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium”. *Adv. Neural Inform. Process. Syst.* 2017 7.
- [JL17] JIANG, HUAIZU and LEARNED-MILLER, ERIK. “Face detection with the faster R-CNN”. *Int. Conf. on Automatic Face & Gesture Recognition.* 2017 3.
- [KAH*20] KARRAS, TERO, AITTALA, MIIKA, HELLSTEN, JANNE, et al. “Training Generative Adversarial Networks with Limited Data”. *Adv. Neural Inform. Process. Syst.* 2020 1, 2, 4.
- [KAL*21] KARRAS, TERO, AITTALA, MIIKA, LAINE, SAMULI, et al. “Alias-Free Generative Adversarial Networks”. *Adv. Neural Inform. Process. Syst.* 2021 3.
- [Kem16] KEMELMACHER-SHLIZERMAN, IRA. “Transfiguring Portraits”. *ACM Trans. Graph.* 2016 3.
- [KLA*20] KARRAS, TERO, LAINE, SAMULI, AITTALA, MIIKA, et al. “Analyzing and improving the image quality of stylegan”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2020, 8110–8119 1–3.
- [KLA19] KARRAS, TERO, LAINE, SAMULI, and AILA, TIMO. “A style-based generator architecture for generative adversarial networks”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2019, 4401–4410 2–4.
- [KNT17] KOWALSKI, MAREK, NARUNIEC, JACEK, and TRZCINSKI, TOMASZ. “Deep alignment network: A convolutional neural network for robust face alignment”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2017 4.
- [LEH13] LEE, YONG JAE, EFROS, ALEXEI A., and HEBERT, MARTIAL. “Style-Aware Mid-level Representation for Discovering Visual Connections in Space and Time”. *Int. Conf. Comput. Vis.* (2013), 1857–1864 2.

- [LGZ*20] LIU, ANDREW, GINOSAR, SHIRY, ZHOU, TINGHUI, et al. “Learning to Factorize and Relight a City”. *Eur. Conf. Comput. Vis.* 2020 2.
- [LLWL20] LEE, CHENG-HAN, LIU, ZIWEI, WU, LINGYUN, and LUO, PING. “MaskGAN: Towards Diverse and Interactive Facial Image Manipulation”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2020 6.
- [LMB*14] LIN, TSUNG-YI, MAIRE, MICHAEL, BELONGIE, SERGE, et al. *Microsoft COCO: Common Objects in Context*. 2014. URL: <http://arxiv.org/abs/1405.0312>.
- [LMC*15] LEE, STEFAN, MAISONNEUVE, NICOLAS, CRANDALL, DAVID, et al. “Linking past to present: Discovering style in two centuries of architecture”. *IEEE International Conference on Computational Photography (ICCP)*. 2015 2.
- [LTM*20] LEE, HSIN-YING, TSENG, HUNG-YU, MAO, QI, et al. “DRIT++: Diverse Image-to-Image Translation via Disentangled Representations”. *Int. J. Comput. Vis.* (2020), 1–16 6.
- [LZY*21] LUO, XUAN, ZHANG, XUANER, YOO, PAUL, et al. “Time-Travel Rephotography”. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia 2021)* 40.6 (Dec. 2021). DOI: <https://doi.org/10.1145/3478513.3480485> 3, 7.
- [MBS17] MATZEN, KEVIN, BALA, KAVITA, and SNAVELY, NOAH. “StreetStyle: Exploring world-wide clothing styles from millions of photos”. *arXiv preprint arXiv:1706.01869* (2017) 2.
- [MMH*19] MALL, UTKARSH, MATZEN, KEVIN, HARIHARAN, BHARATH, et al. “GeoStyle: Discovering fashion trends and events”. *Int. Conf. Comput. Vis.* 2019 2.
- [MS14] MATZEN, KEVIN and SNAVELY, NOAH. “Scene Chronology”. *Eur. Conf. Comput. Vis.* 2014 2.
- [MSE17] MÜLLER-BUDACK, ERIC, SPRINGSTEIN, MATTHIAS, and EWERTH, RALPH. ““When Was This Picture Taken?” – Image Date Estimation in the Wild”. *Advances in Information Retrieval*. Apr. 2017, 619–625. ISBN: 978-3-319-56607-8 2.
- [NBLC20] NITZAN, YOTAM, BERMANO, A., LI, YANGYAN, and COHEN-OR, D. “Face identity disentanglement via latent space mapping”. *ACM Trans. Graph.* 39 (2020), 1–14 5.
- [NH19] NOGUCHI, ATSUHIRO and HARADA, TATSUYA. “Image generation from small datasets via batch statistics adaptation”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2019, 2750–2758 7.
- [OSF*20] OR-EL, ROY, SENGUPTA, SOUMYADIP, FRIED, OHAD, et al. “Lifespan age transformation synthesis”. *Eur. Conf. Comput. Vis.* Springer. 2020, 739–755 1, 3, 4, 6.
- [PA20] PINKNEY, JUSTIN N. M. and ADLER, DORON. “Resolution Dependent GAN Interpolation for Controllable Image Synthesis Between Domains”. *ArXiv abs/2010.05334* (2020) 3, 4.
- [PWS*21] PATASHNIK, OR, WU, ZONGZE, SHECHTMAN, ELI, et al. “StyleCLIP: Text-Driven Manipulation of StyleGAN Imagery”. *Int. Conf. Comput. Vis.* Oct. 2021, 2085–2094 2, 3, 6, 7.
- [RAP*21] RICHARDSON, ELAD, ALALUF, YUVAL, PATASHNIK, OR, et al. “Encoding in Style: a StyleGAN Encoder for Image-to-Image Translation”. *IEEE Conf. Comput. Vis. Pattern Recog.* June 2021 3, 4, 9.
- [RCR18] RUIZ, NATANIEL, CHONG, EUNJI, and REHG, JAMES M. “Fine-Grained Head Pose Estimation Without Keypoints”. *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.* June 2018 4.
- [RHGS15] REN, SHAOQING, HE, KAIMING, GIRSHICK, ROSS, and SUN, JIAN. “Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks”. *Adv. Neural Inform. Process. Syst.* 2015 3.
- [RKH*21] RADFORD, ALEC, KIM, JONG WOOK, HALLACY, CHRIS, et al. “Learning Transferable Visual Models From Natural Language Supervision”. *ICML*. 2021 3, 6.
- [RMBC21] ROICH, DANIEL, MOKADY, RON, BERMANO, AMIT H, and COHEN-OR, DANIEL. “Pivotal Tuning for Latent-based Editing of Real Images”. *arXiv preprint arXiv:2106.05744* (2021) 3–6.
- [RR17] RUIZ, N. and REHG, J. M. “Dockerface: an easy to install and use Faster R-CNN face detector in a Docker container”. *ArXiv e-prints* (Aug. 2017). arXiv: 1708.04370 [cs.CV] 3.
- [SAL*20] STACCHIO, LORENZO, ANGELI, ALESSIA, LISANTI, GIUSEPPE, et al. “IMAGO: A family photo album dataset for a socio-historical analysis of the twentieth century”. *CoRR abs/2012.01955* (2020). arXiv: 2012.01955. URL: <https://arxiv.org/abs/2012.01955>.
- [Sei20] SEITZER, MAXIMILIAN. *pytorch-fid: FID Score for PyTorch*. <https://github.com/mseitzer/pytorch-fid>. Version 0.2.1. Aug. 2020 7.
- [SGTZ20] SHEN, YUJUN, GU, JINJIN, TANG, XIAOOU, and ZHOU, BOLEI. “Interpreting the latent space of gans for semantic face editing”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2020, 9243–9252 3.
- [SVI*16] SZEGEDY, CHRISTIAN, VANHOUCHE, VINCENT, IOFFE, SERGEY, et al. “Rethinking the Inception Architecture for Computer Vision”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2016, 2818–2826 7.
- [SWZJ16] SALEM, TAWFIQ, WORKMAN, SCOTT, ZHAI, MENGHUA, and JACOBS, NATHAN. “Analyzing Human Appearance as a Cue for Dating Images”. *IEEE Winter Conference on Applications of Computer Vision (WACV)*. 2016, 1–8 2.
- [SYTZ20] SHEN, YUJUN, YANG, CEYUAN, TANG, XIAOOU, and ZHOU, BOLEI. “Interfacegan: Interpreting the disentangled face representation learned by gans”. *IEEE Trans. Pattern Anal. Mach. Intell.* (2020) 1–4, 6.
- [SZ21] SHEN, YUJUN and ZHOU, BOLEI. “Closed-form factorization of latent semantics in gans”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2021, 1532–1540 2, 3.
- [TAN*21] TOV, OMER, ALALUF, YUVAL, NITZAN, YOTAM, et al. “Designing an Encoder for StyleGAN Image Manipulation”. *arXiv preprint arXiv:2102.02766* (2021) 3.
- [TL19] TAN, MINGXING and LE, QUOC V. “EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks”. *ArXiv abs/1905.11946* (2019) 7.
- [WGB*20] WANG, YAXING, GONZALEZ-GARCIA, ABEL, BERGA, DAVID, et al. “Minegan: effective knowledge transfer from gans to target domains with few images”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2020, 9332–9341 7.
- [WLS21] WU, ZONGZE, LISCHINSKI, D., and SHECHTMAN, ELI. “StyleSpace Analysis: Disentangled Controls for StyleGAN Image Generation”. *IEEE Conf. Comput. Vis. Pattern Recog.* (2021), 12858–12867 1–3, 6.
- [WNSL21] WU, ZONGZE, NITZAN, YOTAM, SHECHTMAN, ELI, and LISCHINSKI, D. “StyleAlign: Analysis and Applications of Aligned StyleGAN Models”. *ArXiv abs/2110.11323* (2021) 3, 4, 7.
- [WT20] WULFF, JONAS and TORRALBA, ANTONIO. *Improving Inversion and Generation Diversity in StyleGAN using a Gaussianized Latent Space*. 2020. arXiv: 2009.06529 [cs.CV] 3.
- [XZY*21] XIA, WEIHAO, ZHANG, YULUN, YANG, YUJIU, et al. *GAN Inversion: A Survey*. 2021. arXiv: 2101.05278 [cs.CV] 3.
- [YLLT16] YANG, SHUO, LUO, PING, LOY, CHEN CHANGE, and TANG, XIAOOU. “WIDER FACE: A Face Detection Benchmark”. *IEEE Conf. Comput. Vis. Pattern Recog.* 2016 4.
- [ZZI*17] ZHANG, RICHARD, ZHU, JUN-YAN, ISOLA, PHILLIP, et al. “Real-Time User-Guided Image Colorization with Learned Deep Priors”. *ACM Trans. Graph.* 9.4 (2017) 3.