
All Roads Lead to Rome? Investigating Representational Similarities in Generative Image Models

Charumathi Badrinath¹ Usha Bhalla¹ Alex Oesterling¹ Suraj Srinivas¹ Himabindu Lakkaraju¹

Abstract

Do different generative image models secretly learn similar underlying representations? We investigate this by measuring the latent space similarity of four different models: VAEs, GANs, Normalizing Flows (NFs), and Diffusion Models (DMs). Our methodology involves training linear maps between frozen latent spaces to “stitch” arbitrary pairs of encoders and decoders and measuring output-based and probe-based metrics on the resulting “stitched” models. Our main findings are that linear maps between the latent spaces of performant models preserve most visual information even when latent space sizes differ; for CelebA models, *gender* is the most similarly represented probe-able attribute. Finally we show on a Normalizing Flow that latent space representations converge early in training.

1. Introduction

Recent literature on deep networks has suggested that all models are converging to the same representation of “reality”. Huh et al. (2024) show that the embedding spaces of performant vision and language models learn similar notions of “distance” between datapoints of the appropriate modalities. Bansal et al. (2021) “stitch” together layers of pairs of powerful vision networks via a simple map and show that the combined network has minimal loss in performance.

Asperti and Tonelli (2022) investigate whether a similar conclusion holds for *latent spaces* of generative image models by extending the model stitching framework. They train a *linear* map between the latent spaces of a Variational Autoencoder (VAE) and Generative Adversarial Network (GAN) and show that this stitched VAE-encoder-GAN-decoder model is able to produce reconstructions of images

¹Harvard University. Correspondence to: Charumathi Badrinath <charumathi.badrinath@gmail.com>.

Accepted by the *Structured Probabilistic Inference & Generative Modeling workshop* of ICML 2024, Vienna, Austria. Copyright 2024 by the author(s).

that look similar to the originals and have a low RMSE.

However, despite previous work, it remains unclear how well these conclusions hold for other families of generative image models, especially those with vastly different latent sizes. Further, it is not clear which *specific* visual attributes are represented similarly by models and which aren’t. In our work, we address these drawbacks by (1) extending the study to include NFs and DMs, (2) including probe-based metrics to capture semantics of representations. On CelebA we find that powerful models can be stitched via their latent spaces to reconstruct an image with little loss in quality. In addition to color, pose, and lighting, attributes related to *gender* are most similarly represented. Finally, we explore the question of how quickly the local minimum corresponding to this shared latent space is found, by seeing how early in training latent space probe accuracy plateaus. On an NF, we find that this happens less than 20% of the way into training.

Our overall contributions are:

1. In §3.2, we propose two sets of metrics (reconstruction-based, and probe-based metrics) grounded in a model “stitching”-based procedure to measure the similarity of generative image model latent spaces.
2. In §4, we perform extensive experiments across four generative image model classes trained on CelebA (VAEs, GANs, NFs and DMs), and find significant similarity among their latent spaces on our metrics, and especially in their representation of gender. On the Normalizing Flow model, we find that this common latent space structure is learned early in training.

2. Related Works

2.1. Exploring Latent Spaces

Discovering latent directions. Several approaches have been proposed to discover vectors corresponding to high-level image concepts (e.g. hair color) in the GAN latent space. Supervised methods include subtracting the average latent vectors corresponding to images with and without a binary attribute (Larsen et al., 2016) and finding hyperplanes in the latent space separating images with different attribute

values (Shen et al., 2020). One unsupervised method uses PCA on the GAN latent space to discover the most salient edit directions (Härkönen et al., 2020). These techniques have been applied to find edit directions in the latent spaces of VAEs (White, 2016) and NFs (Kingma & Dhariwal, 2018) as well. Our work *compares* the structure of different models’ latent spaces and is thus orthogonal to the exploration of latent spaces in isolation.

2.2. Comparing Model Representations

Metrics. Canonical correlation analysis (Morcos et al., 2018) and centered kernel alignment (Kornblith et al., 2019) are metrics that have been used to measure representational similarity between layers of different models. Another work compared distances between sets of datapoints in different models’ embedding spaces (Huh et al., 2024).

Model Stitching. One branch of work compares representations via “stitching” where a simple map is learned between layers of two different models (Lenc & Vedaldi, 2015). One work stitched together neural networks of different strengths at an intermediate layer, concluding that good models tended to converge on similar representations and that stitching strong models to weak models improved their performance (Bansal et al., 2021). Work on stitching the latent spaces of generative models has also been explored. One paper showed that it is possible to construct a *linear* map between the latent space of image encoders and text decoders to obtain text describing an image, demonstrating that models learn similar representations of multi-modal data (Merullo et al., 2023).

The work most closely related to ours showed that it is possible to linearly stitch the latent spaces of a VAE and GAN and reconstruct an image with little loss in quality (Asperti & Tonelli, 2022). Our work extends this idea to more model architectures, performing a thorough pixel-space and latent-space analysis to characterize what types of models learn similar representations and what specific concepts they represent similarly.

3. Methodology

3.1. Models and Datasets

Model	GAN	VAE	VQVAE	NF	DM
Latent Dim	512	512	768	12288	12288

Table 1. Models and their latent space sizes. See Appendix A.2 for background and Appendix A.3 for more details on the model architectures and implementations used.

We experiment on 5 generative image models trained on CelebA (Liu et al., 2015) – a dataset of celebrity faces annotated with 40 binary attributes. Models (summarized

in Table 1) output images of size 64×64 with slightly different crops.

We use a GAN-generated dataset *CelebA-Synthetic* to train linear maps to and from the GAN latent space, since GANs lack an encoder. The “latent space” of the DM is taken to be the noised image after running the forward process using the DDIM scheduler with 50 timesteps (Song et al., 2020). As a baseline, we use a “random” encoder which maps images deterministically to arbitrary 512-dimensional standard Gaussian vectors.

3.2. Proposed Metrics

Model latent spaces are first stitched via linear maps (see Appendix A.4).

Reconstruction-Based Metrics. We qualitatively assess the image reconstructions of stitched models. We also compute the latent space MSE of the mapping, pixel-space RMSE of the reconstructed image, LPIPS score (Zhang et al., 2018) which measures the perceptual similarity between a pair of images, and the FID (Heusel et al., 2018) which measures how similar the distribution of reconstructed images is to the distribution of target images. Lower is better for all metrics.

Probe-Based Metrics. We train linear probes (see Appendix A.5) on the latent spaces of encoder-decoder models to detect the presence of binary attributes and report test set accuracy. We then measure the percentage of times probe predictions on latents encoded by the model whose latent space it was trained on and those mapped from other latent spaces match (higher is better). We also measure the percent change in probe accuracy on these sets of latents (near 0 is better).

4. Results

4.1. Assessing Stitched Model Reconstructions

We stitch the latent spaces of a VAE, VQVAE, NF, DM and GAN pairwise and compute the reconstruction-based metrics described in Section 3.2.

From Figure 1 we see that stitching the NF and VQVAE encoders to any model \mathbb{X} ’s decoder yields reconstructions resembling those generated by \mathbb{X} (boxed in red) and obtain relatively low latent MSE, RMSE, LPIPS and FID scores (Figure 3). Additional examples are in Appendix A.7.

Conversely, the DM and VAE encoders barely outperform the random encoder in terms of stitched reconstruction quality. Poor VAE performance can be attributed to significant information loss during encoding which is reflected in lossy reconstructions even with its own decoder. For the DM, we note that the “latent vectors” we use to construct our

mappings are heavily noised versions of the image rather than true encodings. The consequence of this is poor linear probe performance (Figure 4). Latent space maps pick up on remaining signal, which is sometimes enough to map to an area of the decoder model’s latent space producing images in the visual vicinity of a target image (Figure 2). The spatial nature of DM “latents” (the “representation” of a feature is a function of the pixel values in the region of the image that the feature manifests) makes the ability to stitch another model’s encoder to its decoder *at all* unexpected. However, it explains why models stitched to the DM decoder often produce reconstructions preserving the geometry of the target image but not the colors.

Finally, Figure 2 shows reconstructions of a test image from CelebA-Synthetic obtained by stitching various models’ encoders to the GAN decoder. Due to the strength of the GAN decoder, all reconstructions resemble a face. Furthermore, mapping to the vicinity of the true GAN latent is sufficient to produce a reconstruction preserving most characteristics of the target image (e.g. age, hair and skin color, expression) supporting results from literature (Asperti & Tonelli, 2022).

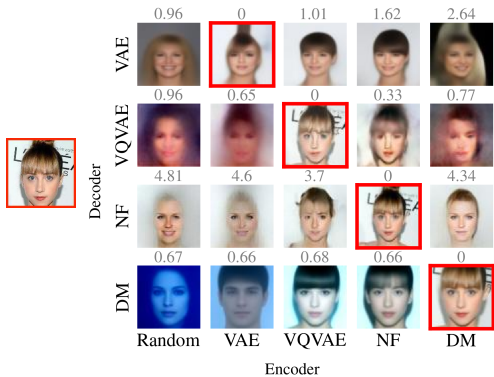


Figure 1. Reconstruction of a CelebA image (left) using stitched models. Latent space MSEs of mapped latents are displayed above each image. The stitched models yielding the best reconstructions use the VQVAE and NF encoders.



Figure 2. Reconstruction of a CelebA-Synthetic image using various encoders stitched to the GAN decoder. The leftmost image is the ground truth. Stitched models using the VQVAE and NF encoders yield the closest reconstructions.

4.2. Identifying Similarly Represented Attributes

We stitch the latent spaces of a VAE, VQVAE, NF, DM and GAN pairwise and compute the probe-based metrics

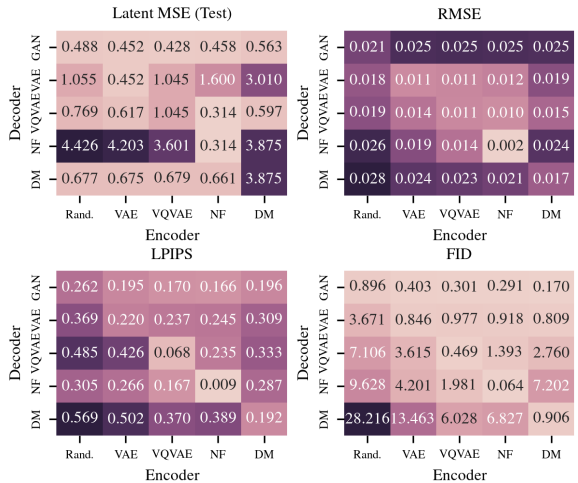


Figure 3. Heatmaps of latent-space MSE, pixel-space RMSE, LPIPS and FID for images reconstructed by stitched models. We see relatively low values of these metrics for models using the NF and VQVAE encoders, corroborating the results of Figure 1.

described in Section 3.2 on 10 interesting CelebA attributes. Results for all 40 attributes are in Appendix A.6.

From Figure 4 we see that linear probes for the gender-related attributes Male and Heavy Makeup achieve 80% accuracy on the NF, VAE and VQVAE latent spaces. In Figure 5 we see that NF probes produce similar predictions on latents produced by its own decoder and latents encoded and mapped from the VQVAE latent space (and vice versa), in agreement with the results from Section 4.1.

Figure 5 also shows a relatively high percentage of matches between probes on NF, VAE, and VQVAE latents, and their counterparts mapped from the GAN latent space. Since the GAN latent space is also known to be linearly separable (Larsen et al., 2016), we conclude that latent space “stitchability” of two models depends on the linearity of both models’ latent representations.

We hypothesize that attributes that are represented similarly in all models’ latent spaces are either uniformly represented in the training dataset (e.g. Smiling), or explain much of the variance in an image (e.g. Male and related attributes). Interestingly, post-map probe match percentages are high for gender-related attributes, particularly Heavy_Makeup and Male, even for maps originating from the non-linear DM latent space indicating that signal corresponding to these attributes is least likely to be lost in the noising process.

Finally, we note in Figure 4 that probes trained on less linearly separable latent spaces are often *more* accurate when applied on latents mapped from a more linearly separable latent space. This explains the phenomenon of the corresponding reconstructions having more “extreme” versions

DM	57	51	64	57	66	65	60	70	56	52
NF	85	78	87	86	89	91	71	87	67	75
VAE	71	78	62	84	82	73	74	68	68	60
VQVAE	87	85	88	86	92	91	76	92	75	75
DM → NF	-17	-25	-37	-16	-12	-30	-11	-19	-7	-26
DM → VAE	-1	-14	-8	-29	-21	-21	-18	-23	-8	-9
DM → VQVAE	-14	-21	-34	-17	-15	-19	-10	-20	-13	-18
NF → DM	21	23	23	26	13	15	16	11	14	3
NF → VAE	9	0	24	-5	-1	1	6	7	7	18
NF → VQVAE	4	-2	1	-8	-3	-5	7	-1	0	-5
VAE → DM	8	43	-9	26	18	-3	6	-8	5	-3
VAE → NF	-12	-2	-34	-3	-11	-27	-1	-26	13	-20
VAE → VQVAE	-12	-9	-37	-6	-13	-20	3	-27	1	-14
VQVAE → DM	28	39	29	31	18	29	16	15	1	13
VQVAE → NF	3	3	3	-2	4	2	7	2	7	-10
VQVAE → VAE	19	1	25	-7	0	6	6	8	5	13
Map										
	Bald	Blond_Hair	Eyeglasses	Heavy_Makeup	Male	No_Beard	Pale_Skin	Smiling	Wavy_Hair	Young
Attribute										

Figure 4. The top 4 rows show latent space linear probe accuracy on various binary attributes. The NF and VQVAE have the most linearly separable latent spaces. The remaining rows show change in accuracy of probe trained on model \mathbb{X} 's latent space when making predictions on latents encoded by \mathbb{X} and latents mapped from $\mathbb{X}' \rightarrow \mathbb{X}$. Probe accuracy increases on latents mapped from a more linearly separable to less linearly separable latent space.

of the attributes present in the target image. We see an example of this in Figure 1 where the reconstructions produced by the NF and VQVAE stitched to the DM have very distinct, block-like bangs.

DM → NF	65	72	59	78	81	64	66	73	69	64
DM → VAE	66	67	48	63	60	57	60	60	64	57
DM → VQVAE	75	72	66	77	82	74	74	77	76	72
GAN → DM	67	60	53	57	47	67	54	59	54	70
GAN → NF	69	71	64	75	76	74	63	78	76	75
GAN → VAE	76	76	72	76	75	78	81	79	81	75
GAN → VQVAE	72	80	82	87	81	73	83	90	76	69
NF → DM	60	62	56	67	65	64	60	76	50	50
NF → VAE	71	84	73	81	79	83	81	83	81	69
NF → VQVAE	92	90	89	87	95	93	88	99	84	84
VAE → DM	55	57	44	53	57	62	54	57	49	52
VAE → NF	67	80	60	81	74	71	69	67	73	69
VAE → VQVAE	77	82	61	82	76	75	83	69	79	67
VQVAE → DM	48	60	61	62	57	65	54	77	53	53
VQVAE → NF	87	87	91	90	94	92	75	92	81	80
VQVAE → VAE	76	89	70	84	78	83	83	80	82	78
Map										
	Bald	Blond_Hair	Eyeglasses	Heavy_Makeup	Male	No_Beard	Pale_Skin	Smiling	Wavy_Hair	Young
Attribute										

Figure 5. Percentage of times a probe trained on model \mathbb{X} 's latent space produces the same prediction on latents encoded by \mathbb{X} and latents mapped from $\mathbb{X}' \rightarrow \mathbb{X}$. Attributes correlated with gender are represented similarly by nearly every model.

4.3. Observing Latent Space Structure Over Training

We train the Glow NF (Kingma & Dhariwal, 2018) on CelebA for 26 epochs checkpointing every 5 epochs. We train linear probes for 10 interesting attributes on the frozen latent spaces and measure how their accuracy changes as model training progresses.




Figure 6. Test set accuracy of linear probes on NF latent space frozen at increasing training epochs. Probe accuracy on most attributes seems to plateau after epoch 6.



Figure 7. Reconstructions generated by encoding an image with a partially trained NF trained and decoding with a fully trained NF. Reconstruction quality improves with encoder training.

Figure 6 shows that probe accuracy begins to plateau for all attributes by epoch 6 of 26. Attributes correlated with gender as well as Smiling (the similarly represented attributes identified in Section 4.2) change minimally in their representations after epoch 1 of 26. Figure 7 shows that there is slight improvement in images reconstructed from latents encoded by increasingly trained NF models, indicating that the latent space still changes over time, albeit minimally after epoch 11. Thus, we hypothesize that a “universal” representation is natural for models to learn and that later training epochs improve the decoder given a somewhat fixed latent representation.

5. Discussion and Future Work

This work advances the hypothesis that as (unconditional) generative image models become more expressive, their latent space representations (given that they are meaningful) become more similar to one another. We find that this shared representation is learned early on in model training, indicating that it is quite “natural” in describing the data. This suggests that the latent spaces of increasingly expressive models may be converging to an “optimal” data representation in the limit, which has deep implications for how we construct world models.

Similarity of latent space representations also means that edit directions, unsafe regions, or areas of bias found in the latent space of one model, can be mapped onto any other model via a simple linear map. It also allows for more powerful image editing by stitching models with strong encoders to models with strong decoders via their latent spaces.

Future work can investigate whether similar results hold for models trained on datasets containing images from multiple

classes with more variance, and can probe similarity of representation of non-binary attributes. We expect to see similar results given that the models used are able to produce realistic samples from the data distribution. We also wish to extend our analysis to pipelines composed of *multiple* foundation models, for example, the latent diffusion model (Rombach et al., 2022) which trains a DM on the VQVAE latent space. Since it is the *first* foundation model in this pipeline that learns a latent representation of the data, we expect to see similar results to this paper. Finally, work has shown that it is possible to find a linearly separable latent space for DMs (Kwon et al., 2023); we can validate that this latent space is similarly structured to those studied in our paper.

Acknowledgements

We would like to thank the GRaM workshop’s ELLIS Mobility Grant as well as the SPIGM workshop for supporting CB’s travel and attendance at ICML.

References

- Asperti, A. and Tonelli, V. Comparing the latent space of generative models, 2022.
- Bansal, Y., Nakkiran, P., and Barak, B. Revisiting model stitching to compare neural representations, 2021.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks, 2014.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium, 2018.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models, 2020.
- Huh, M., Cheung, B., Wang, T., and Isola, P. The platonic representation hypothesis, 2024.
- Härkönen, E., Hertzmann, A., Lehtinen, J., and Paris, S. Ganspace: Discovering interpretable gan controls, 2020.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks, 2019.
- Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., and Aila, T. Training generative adversarial networks with limited data, 2020.
- Kim, S. glow-pytorch. <https://github.com/rosinality/glow-pytorch>, 2020.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes, 2022.
- Kornblith, S., Norouzi, M., Lee, H., and Hinton, G. Similarity of neural network representations revisited, 2019.
- Kwon, M., Jeong, J., and Uh, Y. Diffusion models already have a semantic latent space, 2023.
- Larsen, A. B. L., Sønderby, S. K., Larochelle, H., and Winther, O. Autoencoding beyond pixels using a learned similarity metric, 2016.
- Lenc, K. and Vedaldi, A. Understanding image representations by measuring their equivariance and equivalence, 2015.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Merullo, J., Castricato, L., Eickhoff, C., and Pavlick, E. Linearly mapping from image to text space, 2023.
- Morcos, A. S., Raghu, M., and Bengio, S. Insights on representational similarity in neural networks with canonical correlation, 2018.
- Pandey, K., Mukherjee, A., Rai, P., and Kumar, A. VAEs meet diffusion models: Efficient and high-fidelity generation. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. URL https://openreview.net/forum?id=-J8dM4ed_92.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows, 2016.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models, 2022.
- Shen, Y., Gu, J., Tang, X., and Zhou, B. Interpreting the latent space of gans for semantic face editing, 2020.
- Song, J., Meng, C., and Ermon, S. Denoising diffusion implicit models. *arXiv:2010.02502*, October 2020. URL <https://arxiv.org/abs/2010.02502>.
- van den Oord, A., Vinyals, O., and Kavukcuoglu, K. Neural discrete representation learning, 2018.

von Platen, P., Patil, S., Lozhkov, A., Cuenca, P., Lambert, N., Rasul, K., Davaadorj, M., Nair, D., Paul, S., Liu, S., Berman, W., Xu, Y., and Wolf, T. Diffusers: State-of-the-art diffusion models. URL <https://github.com/huggingface/diffusers>.

Wang, Z., Zheng, H., He, P., Chen, W., and Zhou, M. Diffusion-gan: Training gans with diffusion. *arXiv preprint arXiv:2206.02262*, 2022. URL <https://arxiv.org/abs/2206.02262>.

White, T. Sampling generative networks, 2016.

Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric, 2018.

A. Appendix

A.1. Source Code

All code for this paper is available in the following GitHub repository: <https://github.com/charumathib/thesis-latent-spaces>.

A.2. Background on Models

All generative image models learn to “sample” images following their the distribution of their training set. However, the four main families of models – Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), Normalizing Flows (NFs) and Diffusion Models (DMs) – differ drastically in how they do so.

A VAE (Kingma & Welling, 2022) consists of an encoder network mapping images to latent vectors of lower dimension, and a decoder neural network that maps latent vectors to images with the objective of minimizing image reconstruction error. A KL-Divergence term is added to the objective to encourage a marginally Gaussian latent space distribution.

A GAN (Goodfellow et al., 2014) is comprised of a generator network (decoder) that learns to generate images from marginally Gaussian latent samples and a discriminator network whose objective is to distinguish images from the training set from generated images. Both networks improve by playing a minimax game, ultimately yielding a generator capable of producing realistic images.

An NF (Rezende & Mohamed, 2016) samples a latent variable from a simple distribution and repeatedly applies a series of invertible functions to this variable until it is transformed into an image. The marginal likelihood of each datapoint can be derived via the change of variables theorem for distributions and the objective is the negative data log likelihood.

Finally, a DM (Ho et al., 2020) progressively adds Gaussian noise to an image producing a sequence noisy samples that become isotropic to a Gaussian in the limit (“forward process”). It then learns how to produce samples from the “reverse process” to gradually transform pure noise into an image. DMs dont have an obvious latent space in the same way as VAEs, GANs and NFs since marginally Gaussian “latents” are created via progressive noising rather then by a transformation that would force a representation to be learned.

A.3. Models Used

The pre-trained GAN we use is from the Diffusion-GAN codebase (Wang et al., 2022) and follows the StyleGAN2-ADA architecture (Karras et al., 2020). The pre-trained VAE is from the DiffuseVAE codebase (Pandey et al., 2021). The pre-trained VQVAE (van den Oord et al., 2018) from the Hugging Face `diffusers` library (von Platen et al.). The NF implementation is from an independent source (Kim, 2020) follows the Glow architecture (Kingma & Dhariwal, 2018). The DM is from the DDIM codebase (Song et al., 2020).

A.4. Training Latent Space Linear Maps

Many of our experiments rely on the construction of linear maps between model latent spaces. For each pair of encoder-decoder models (VAE, VQVAE, NF, DM), we encode the first 9000 images from the CelebA test split into their respective latent spaces and use a closed-form linear regression solver from `sklearn` to predict one set of latents from the other (Pedregosa et al., 2011). To prevent overfitting in maps originating from the NF and DM latent spaces, we use ridge regression with manually tuned values of α as summarized in Table 2.

Map	α
DM \rightarrow GAN	2000
DM \rightarrow VAE	100
DM \rightarrow VQVAE	5000
DM \rightarrow NF	5000
NF \rightarrow GAN	50000
NF \rightarrow VAE	5000
NF \rightarrow VQVAE	50000
NF \rightarrow DM	50000

Table 2. Regularization parameters used in Ridge regression for each linear map. Maps that are not listed in this table were not regularized.

All Roads Lead to Rome?

Our training procedure differs slightly for maps to and from the GAN latent space. We use the “style space” as our latent space (a layer connected to the base latent space by an 8 layer MLP) (Karras et al., 2019). We use the first 9000 images in the CelebA-Synthetic dataset to train maps in the way described earlier; the latent used by the GAN to generate each image is treated as the image’s “encoding” in the GAN latent space. All maps are tested on a holdout set of 100 images from either CelebA or CelebA-Synthetic as appropriate.




Figure 8. Latent space map MSE on train (left) and test (right) sets.

A.5. Training Latent Space Linear Probes

We train linear probes to predict the value of each of the 40 binary attributes (e.g. Male) CelebA images are annotated with. To construct a balanced training set for an attribute’s probe, we first compute the number of images with and without that attribute in the first 9000 CelebA images. 80% of the minimum of these two numbers is how many positive and negative examples are included. Probes are trained using lasso regression in sklearn where $\alpha = 0.005$ for probes on the VAE and VQVAE latent spaces, 0.02 for probes on the DM latent space and 0.1 for probes on the NF latent space. (Pedregosa et al., 2011). We use a holdout set of 200 latents (100 with attribute, 100 without attribute) to evaluate probe performance.

A.6. Full Probe Results




Figure 9. Full version of Figure 4 (top) with results for all 40 attributes.




Figure 10. Full version of Figure 4 (bottom) with results for all 40 attributes.

All Roads Lead to Rome?

Map	5_o_Clock_Shadow	ArchedEyebrows	Attractive	Bags_UnderEyes	Bald	Bangs	BigLips	BigNose	BlackHair	BlondHair	Blurry	BrownHair	BushyEyebrows	Chubby	DoubleChin	Eyeglasses	Goatee	GrayHair	HeavyMakeup	HighCheekbones	Male	Mouth_SlightlyOpen	Mustache	NarrowEyes	No_Beard	OvalFace	PaleSkin	PointyNose	RecedingHairline	RosyCheeks	Sideburns	Smiling	StraightHair	WavyHair	WearingEarrings	WearingHat	WearingLipstick	WearingNecklace	WearingNecktie	Young
DM → NF	60	69	74	72	65	59	61	68	64	72	65	63	70	61	69	59	52	62	78	80	81	82	62	65	64	76	66	71	57	67	52	73	64	69	65	64	80	77	64	64
DM → VAE	59	67	63	63	66	63	53	56	70	67	46	60	57	56	48	48	46	63	63	56	60	59	61	48	57	55	60	48	60	51	42	60	54	64	60	69	65	59	50	57
DM → VQVAE	72	65	82	69	75	69	64	72	71	72	67	72	77	64	79	66	56	70	77	72	82	76	66	72	74	75	74	63	63	75	69	77	69	76	64	70	76	62	70	72
GAN → DM	47	57	60	52	67	57	52	51	65	60	60	59	56	61	54	53	61	59	57	64	47	63	64	55	67	54	54	57	64	57	53	59	59	54	59	64	51	60	51	70
GAN → NF	79	79	83	60	69	63	66	66	78	71	59	71	77	64	59	64	79	61	75	84	76	82	70	66	74	72	63	66	71	85	78	78	76	74	77	74	68	69	75	
GAN → VAE	71	81	78	65	76	79	59	73	76	76	56	82	71	62	67	72	61	74	76	82	75	81	62	67	78	70	81	77	76	69	67	79	63	81	69	81	80	64	71	75
GAN → VQVAE	79	83	88	79	72	78	75	75	78	80	63	76	79	75	57	82	72	81	87	94	81	76	67	68	73	91	83	78	67	86	67	90	61	76	83	71	83	76	75	69
NF → DM	67	63	57	63	60	57	54	56	65	62	49	54	57	54	55	56	63	61	67	71	65	70	56	53	64	56	60	60	53	61	61	76	47	50	65	59	63	61	66	50
NF → VAE	71	86	86	72	71	78	69	66	87	84	49	71	70	62	59	73	68	81	81	81	79	79	68	50	83	74	81	71	63	65	60	83	61	81	66	76	77	57	63	69
NF → VQVAE	90	83	86	85	92	92	75	87	85	90	79	87	85	80	86	89	84	85	87	89	95	90	83	80	93	86	88	78	83	90	93	99	88	84	87	90	90	81	87	84
VAE → DM	56	48	48	46	55	48	41	46	55	57	45	47	48	56	56	44	57	44	53	52	57	56	45	49	62	51	54	62	44	50	48	57	51	49	52	54	47	51	51	52
VAE → NF	61	74	80	68	67	67	67	66	70	80	66	66	67	63	68	60	56	64	81	68	74	65	54	61	71	69	69	63	63	73	50	67	57	73	73	69	81	68	64	69
VAE → VQVAE	64	65	81	70	77	69	59	77	79	82	63	72	61	61	63	61	56	74	82	65	76	61	56	59	75	67	83	69	66	66	54	69	69	79	70	71	75	62	66	67
VQVAE → DM	56	52	56	62	48	50	62	56	63	60	55	53	48	56	57	61	57	62	62	69	57	72	48	52	65	49	54	50	62	56	55	77	44	53	57	62	63	46	57	53
VQVAE → NF	80	84	82	86	87	86	73	77	77	87	70	76	83	79	78	91	84	82	90	92	94	88	78	77	92	82	75	81	80	90	79	92	74	81	82	85	87	83	81	80
VQVAE → VAE	77	87	85	74	76	84	64	75	86	89	63	80	75	56	65	70	73	86	84	82	78	78	64	52	83	78	83	74	67	69	64	80	63	82	74	79	83	68	67	78

Figure 11. Full version of Figure 5 with results for all 40 attributes.

A.7. Additional Examples of Stitched Model Reconstructions

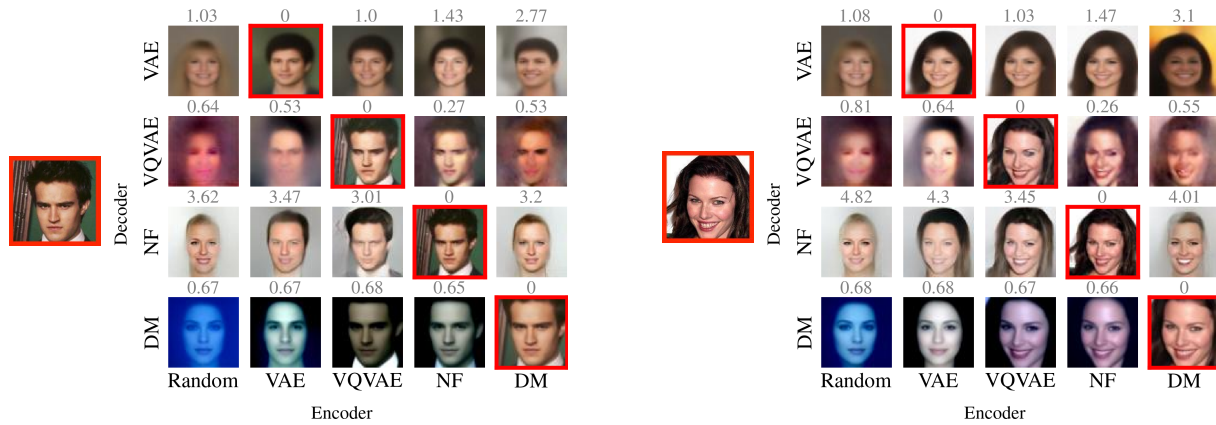


Figure 12. See caption of Figure 1.




Figure 13. See caption of Figure 2.