On The Local Geometry of Deep Generative Manifolds

Ahmed Imtiaz Humayun [1 2] **Ibtihel Amara** [3 2] **Candice Schumann** [4] **Golnoosh Farnadi** [2] **Negar Rostamzadeh** [2] **Mohammad Havaei** [2]

Abstract

In this paper, we study theoretically inspired local geometric descriptors of the data manifolds approximated by pre-trained generative models. The descriptors – local scaling (ψ), local rank (ν), and local complexity (δ) — characterize the uncertainty, dimensionality, and smoothness on the learned manifold, using only the network weights and architecture. We investigate and emphasize their critical role in understanding generative models. Our analysis reveals that the local geometry is intricately linked to the quality and diversity of generated outputs. Additionally, we see that the geometric properties are distinct for out-of-distribution (OOD) inputs as well as for prompts memorized by Stable Diffusion, showing the possible application of our proposed descriptors for downstream detection and assessment of pre-trained generative models.

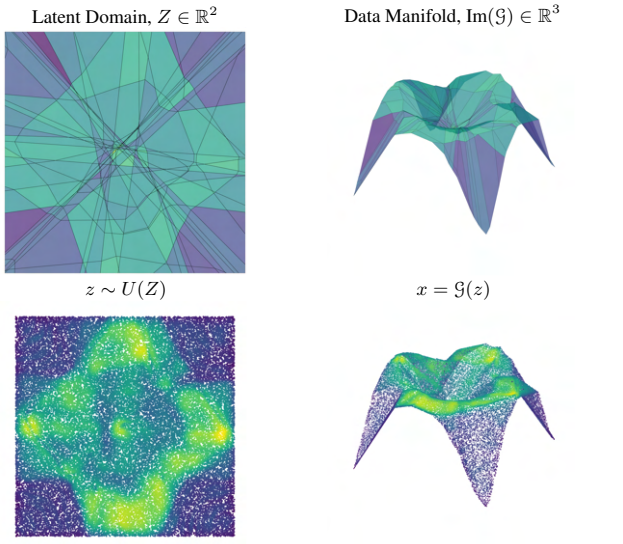

Figure 1. Analytical visualization of the piece-wise linear manifold (top-right) learned by a toy MLP generator $\mathcal{G} : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ with Width 20 and Depth 3. The generator subdivides the input space (top-left) into convex regions and affinely maps each region to the output (top-right). Regions are colored by the local scaling ψ_ω ($\uparrow$ for blue), induced by their corresponding affine maps. Uniform latent distribution (bottom-left) and generated samples (bottom-right) colored by a Kernel Density Estimate (KDE) of the output distribution ($\downarrow$ blue), shows a strong correlation between KDE and local scaling induced by the generator. *Local scaling is indicative of the un-normalized uncertainty of a latent vector.*

1. Introduction

In recent years, deep generative models have emerged as a powerful tool in machine learning, capable of synthesizing realistic data across diverse domains (15; 16; 21). However, evaluating these models remains a critical challenge that extends beyond assessing mere sample quality and conditional alignment to broader concerns about algorithmic bias and responsible generation. Current methods for evaluating generative models heavily rely on data-driven approaches or human evaluators, both of which have inherent limitations. Data-dependent assessments can be biased by the availability and representativeness of the data, while human evaluators introduce subjectivity, leading to potential biases in the assessment process. Data-dependent assessments, while valuable, are inherently constrained by the availability and representativeness of the data used. This limitation can inadvertently introduce bias into the evaluation process, as models may be unfairly penalized or favored based on the specific data they are evaluated against. Similarly, human evaluators, though often considered the gold standard, inevitably bring their own subjective biases to the assessment, leading to potential inconsistencies and unfair comparisons between models (3).

In light of these challenges, our paper poses a fundamental question; *Can we use local geometric characteristics of a generator data manifold for assessment?*. This question is inspired by prior research that has demonstrated the utility of curvature in various aspects of generative modeling, such as improving variance estimation (1), latent space

[1]Rice University [2]Google Research [3]McGill University [4]Google Deepmind. Correspondence to: Ahmed Imtiaz Humayun <imtiaz@rice.edu>, Mohammad Havaei <mhavaei@google.com>.

Accepted as an extended abstract for the Geometry-grounded Representation Learning and Generative Modeling Workshop at the 41st International Conference on Machine Learning, ICML 2024, Vienna, Austria.

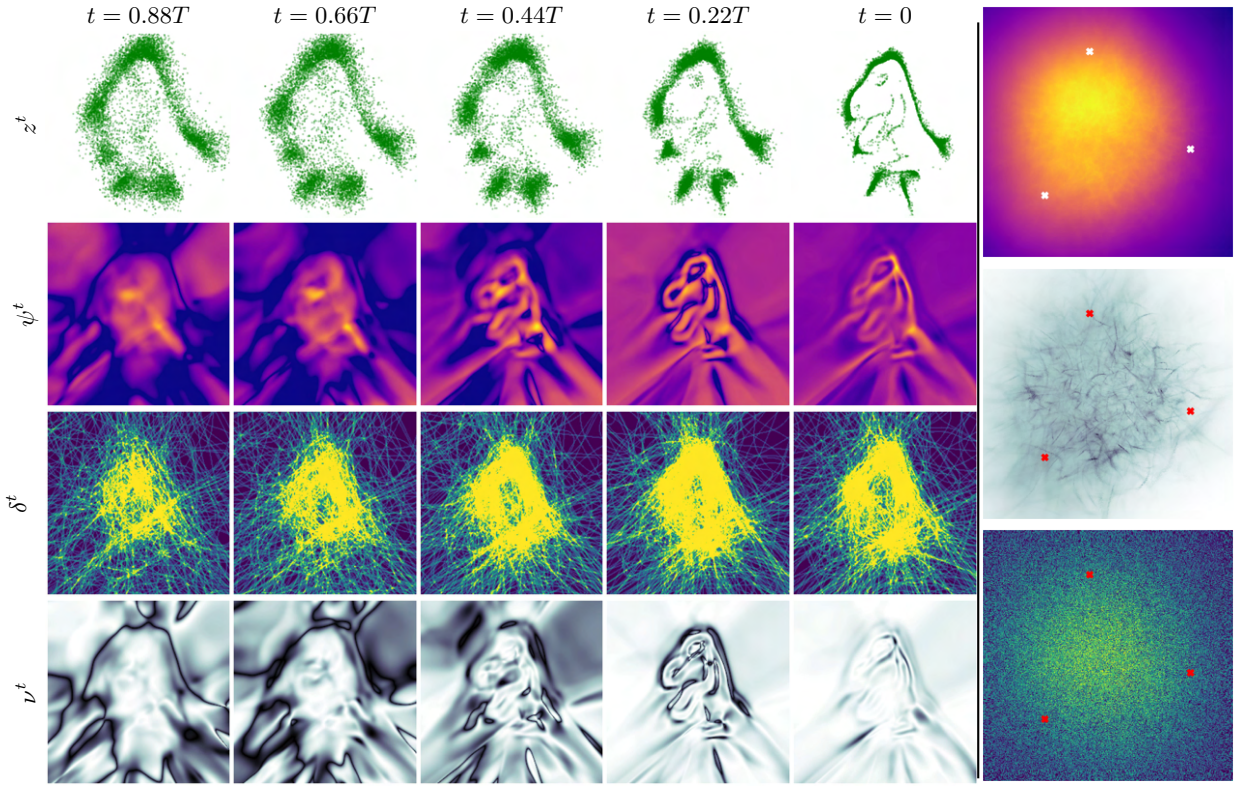

Figure 2. Geometric descriptors for 1) left-panel: a Denoising Diffusion Probabilistic Model trained on a 1-manifold embedded in $\mathbb{R}^2$ and 2) right-panel: the Stable Diffusion (21) VQGAN decoder computed for a 2D subspace of the latent space that passes through the denoised latents of "a fox", "a cat" and "a dog" (marked 'X'). We see that with decreasing t in the reverse diffusion process, local scaling ψ^t and local rank ν^t decreases, while local complexity δ^t increases on the domain of a denoised gaussian distribution (z^T) at denoising step t (z^t). Note that ψ^t is proportional to the change in density between z^t and z^{t-1}. The variance of ψ^0 and ν^0 across $\mathbb{R}^2$ is significantly lower than $t \neq 0$. For the stable diffusion latent space subspace visualized here, we see that ψ,δ is higher and ν is in the convex hull of the three denoised latents. For ν we see pockets of low rank regions in the convex hull, indicating that latents outside the latent space domain of the VQGAN has considerably higher rank. See appendix for decoded images from the subspace.

sampling (13), and controlling attributes in generated samples (8).

We introduce a new self-assessment framework for generative models using local geometry based on the theory of continuous piecewise affine spline generators (11; 2). We define three geometric measures: local scaling (ψ), local rank (ν), and local complexity (δ) to capture uncertainty, dimensionality, and smoothness, respectively. Our experiments show these correlate with quality, aesthetics, diversity, and bias in generated data. Additionally, the framework allows for out-of-distribution detection and model comparison. Specifically, we make the following contributions: (A) Define local descriptors (scaling, rank, complexity) based on data manifold geometry (Section 2). (B) Investigate how these descriptors correlate with quality, aesthetics, and diversity in generated data, revealing potential for bias detection (Section 3).

2. Geometric Descriptors of the Learned Data Manifold

Consider a generative network $\mathcal{G}$, which can be the decoder of a Variational Autoencoder (VAE) (17), the generator of a Generative Adversarial Network (GAN) (6) or an unrolled denoising diffusion implicit model (DDIM) (22). Suppose, $\mathcal{G} : \mathbb{R}^E \to \mathbb{R}^D$ is a deep neural network containing with L layers, input space dimensionality E and output space dimensionality D.

Let, the generator layer operations and activation functions be continuous piece-wise linear, e.g., fully-connected, convolution, or pooling layers and leaky-ReLU or ReLU activations. Therefore, the data manifold or image of the generator $\text{Im}(\mathcal{G})$ is formed via a continuous piece-wise affine mapping of the latent space $\mathbb{R}^E$ onto the output space $\mathbb{R}^D$, with the data manifold having a dimensionality of at most E everywhere. The data manifold can thus be expressed as the union

of sets:

$$\text{Im}(\mathcal{G}) = \bigcup_{\omega \in \Omega} \{\boldsymbol{A}_\omega z + \boldsymbol{b}_\omega : z \in \omega\}, \tag{1}$$

where, Ω is the partition of the latent space $\mathbb{R}^E$ into piecewise-linear regions, $\boldsymbol{A}_\omega$ and $\boldsymbol{b}_\omega$ are parameters of the affine mapping from latent space vectors $z \in \omega$ to the data manifold. Here $\boldsymbol{A}_\omega, \boldsymbol{b}_\omega$ are functions of the neurons/parameters of the network, such that entries in $\boldsymbol{A}_\omega$ are non-zero only if the corresponding activation function is active for all the vectors in ω. $\boldsymbol{A}_\omega, \boldsymbol{b}_\omega \forall \omega \in \Omega$ are therefore parameters of a continuous piecewise affine spline function that represents the generator (2).

Illustrative Example: Toy generator trained on a $f : \mathbb{R}^2 \to \mathbb{R}^3$ task. To illustrate the idea of a continuous piecewise linear manifold, we train a toy generator $\mathcal{G} : \mathbb{R}^2 \to \mathbb{R}^3$ with depth 3 and width 20, to map a 2-dimensional latent space onto a toy 2-manifold in a 3-dimensional output space. In Fig. 1-top, we present analytically computed visualization (10) of the piece-wise linear manifold learned by the generator as well as the latent space partition Ω represented by dark lines. Every black line represents a non-linearity of the function that folds/bends the latent space while going from $\mathbb{R}^2$ to $\mathbb{R}^3$. Therefore, the black lines are knots of the continuous piecewise affine spline generator. Each convex region ω formed by the intersection of the black lines, is mapped to $\text{Im}(\mathcal{G})$ via per region parameters as described in Equation 1. Given such a piece-wise affine manifold we propose the following local geometric descriptors:

Local scaling, ψ. For a continuous piece-wise linear manifold produced by generator $\mathcal{G}$ the local scaling is constant for every region ω and can be expressed as:

$$\psi_\omega = \sum_i^k \log(\sigma_i), \tag{2}$$

where $\{\sigma_i\}_{i=0}^{i=k}$ are k non-zero singular values of $\boldsymbol{A}_\omega$. Consider a uniform distribution at the input of $\mathcal{G}$. The output density on every region of $\text{Im}(\mathcal{G})$ would therefore be proportional to e^{ψ_ω}, with lower likelihood for higher e^{ψ}. ψ_ω can therefore be considered an non-normalized log-uncertainty measure for a given generator region ω. We demonstrate the effect of ψ in Fig. 1-top, where each region in the top row are colored by e^{ψ}.

Local complexity, δ. The local complexity of a continuous piece-wise linear Deep Neural Network, is a measure of the density of linear regions in a neighborhood (7). Therefore, the local complexity of $\text{Im}(\mathcal{G})$, for a latent vector z can be written as:

$$\delta_z = \sum_{\forall \omega \in V_z} \mathbb{1}_\omega$$
$$\text{where } V_z = \{x \in \mathbb{R}^E : ||Bx - Bz||_1 < r\}.$$

Here, B is an arbitrary wide orthonormal matrix, $||.||_1$ is the ℓ_1 norm operator and r is a radius parameter denoting the size of the locality to compute δ for. We use the method described in (14) to estimate the number of regions ω inside ℓ_1-ball V_z centered on z.

Local rank, ν. Similar to ψ_ω, local rank ν_ω is also constant per-region and can be expressed as:

$$\nu_\omega = \exp\left(-\sum_i p_i \log(p_i)\right)$$
$$\text{where } p_i = \frac{\sigma_i}{\sum_i \sigma_i} + \text{eps}.$$

Here, σ_i are singular values of $\boldsymbol{A}_\omega$ and eps$= 10^{-30}$ is a constant. The local rank ν_ω therefore denotes the dimensionality of the locally linear subspace on the data manifold to which ω is mapped to.

Figure 3. Images with the lowest (left) and highest (right) local rank ν from a set of 20000 randomly sampled ImageNet dataset samples. *Low rank images contain simpler textures for every class compared to the high rank samples.* This is because for images with higher local rank, the learned manifold is higher dimensional locally therefore allowing higher independent degrees of variations for the generated images.

What do these descriptors describe? Metrics like local scaling, local complexity, as well as the local rank have previously been used in the context of Deep Neural Networks (7; 14), Self-Supervised Learning (5) as well as Generative Modeling (18; 12). The local scaling for any given region of a continuous piecewise linear function denotes the change of density by the input-output mapping (12), with higher scaling resulting in higher uncertainity. The local complexity of a network has been shown to quantify the smoothness

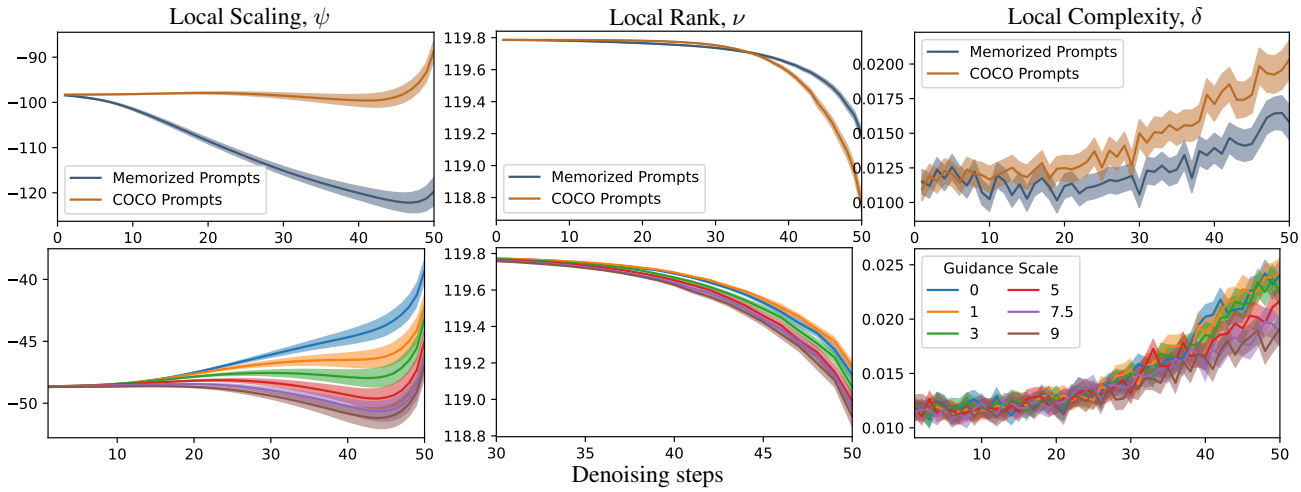

Figure 4. Geometric descriptors computed for the VQGAN decoder, during 50 stable diffusion denoising steps, for (top) 100 COCO and 100 memorized prompts (23) with guidance scale 7.5 and (bottom) 100 COCO prompts with varying guidance scales. Each prompt is generated for 4 seeds and we use the same set of seeds for different guidance scales as well as prompt types. Shaded region represents 95% confidence interval. *We see that the local geometry trajectories are discriminative of memorization, as well as increased quality when using stronger classifier free guidance.*

or expressivity of the function learned by a DNN (19). The local rank for any given region, is the rank of the linear operation for that region. Therefore it denotes the local dimensionality of the learned manifold.

Going back to Fig. 1, we sample uniformly from the latent space Z and fit a gaussian kernel density estimator on the samples generated in $\mathbb{R}^3$. In Fig. 1-right, we present uniform samples from the latent space, as well as samples from the data manifold colored by the kernel density estimate (KDE). Therefore local scaling ψ is inversely proportional to the density of the learned distribution by a generative model, or proportional to the uncertainity of the model for any latent vector.

3. Exploring Generative Model Manifolds using Descriptors

In this Section we will be exploring the geometry of the data manifolds learned by various generative models, e.g., denoising diffusion probabilistic models (9), latent diffusion models like Stable Diffusion (21).

DDPM trained on toy $f : \mathbb{R}^2 \to \mathbb{R}^2$ generation task. We train a denoising diffusion probabilistic (9) model on a toy dataset[1] and present the local complexity δ^t, local scaling ψ^t and local rank ν^t at denoising timesteps $t \in \{6, 17, 28, 39, 50\}$ in Fig. 2-(Left Panel). For any given timestep, the local scaling ψ^t is proportional to the local concentration of density going from timestep t to $t-1$. We see that ψ^t for smaller t is low on the data manifold, and higher in the regions adjacent to the data manifold. The local complexity is higher around the data manifold denoting higher expressivity (20) of the network, increasing with decreasing t. Local rank is lower around the data manifold until the last few denoising steps, when the variance of ψ and ν across the input space diminishes.

[1] http://www.thefunctionalart.com/2016/08/download-datasaurus-never-trust-summary.html

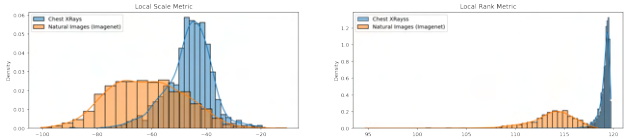

Figure 5. Chest X-ray vs Imagenet local scaling and local rank for Stable Diffusion. *Local geometry is descriptive of whether a dataset is out of distribution.*

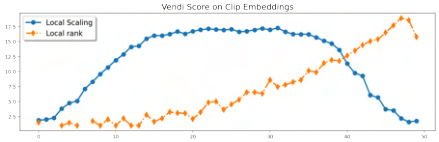

Figure 6. Vendi Score Computed on Clip Image Embeddings for 50k Generated Samples of ImageNet Classes. The y-axis represents the Vendi score, while the x-axis displays bins for sorted values of local scaling (blue) and local rank (orange), increasing from left to right. See Figure 8 and Figure 9 in the Appendix for samples of generated images corresponding to each bin for local scaling and local rank respectively.

Local geometry of Stable Diffusion We explore how the geometric descriptors vary for Stable Diffusion v1.4 (21) which is a text-to-image latent diffusion model. We compute the geometric descriptors for the decoder only. Note that in Section 2, we introduce the local descriptors for piece-wise linear manifolds only. However, for foundation models like Stable Diffusion, it is generally less likely to be piece-wise linear. In such cases $\boldsymbol{A}_\omega$ can be approximated by taking the jacobian of the input-output mapping of $\mathcal{G}$ evaluated at a point $z \in \omega$ of the latent space. For non-piecewise linear generators, this would be equivalent to a first-order approximation of the descriptor around a vector z. Instead of computing the full jacobian matrix, we do a random orthonormal projection of the output manifold onto a 120 dimensional subspace for Stable Diffusion, and compute the input-output jacobian considering the random projection as the output. The projection is kept fixed for all latent vectors z. For δ, we use the method described in (14) and consider 20 dimensional ℓ_1 balls of radius $r = 1e-4$ as the neighborhoods to compute δ.

Geometric descriptors can be used to perform out-of-distribution detection. In Figure 5 we present the local scaling and local rank computed for Imagenet (considered in-distribution) and Chest X-ray (considered out-of-distirbution) images encoded in the latent space of Stable Diffusion. Both local scaling and rank are discriminative of Chest X-ray images.

We also explore the relationship between local scaling, local ranking, and diversity using the well-known diversity score, the Vendi score (4). The results are illustrated in Figure 6. As shown in this figure, for both local scaling and local rank vendi score correlates positively with increase in the descriptor values. The observed drop in the Vendi score for high local scaling is anticipated; at very high uncertainty levels, the model targets regions with very low probability in the learned manifold, effectively collapsing to the strongest anti-mode.

In Figure 3 we show how that the textural diversity of generated images is connected with ν computed at stable diffusion latents for the decoder. Images with higher ν, has higher frequency textures. In Figure 4 we show that connections between local geometry and memorization/quality-diversity of generation.

4. Conclusion & Future Directions

In this paper, we proposed a novel assessment approach to evaluate generative models using geometry-based descriptors – local scaling (ψ), local rank (ν) and local complexity (δ) - effectively while utilizing only the model's architecture and weights. Our approach characterizes uncertainty, dimensionality, and smoothness of the learned manifold without requiring original training data or human evaluators. Our experiments demonstrated how these descriptors relate to generation quality, diversity, and memorization. While using the geometry of manifolds offers a novel approach to self-assess generative models, we acknowledge two main limitations that warrant further investigation. First, the geometry of the learned manifold is inherently influenced by the training dynamics of the model. A deeper understanding of this relationship is needed to fully leverage geometric analysis for model assessment and improvement. Second, the computational complexity of our method, particularly the calculation of the Jacobian matrix, poses a practical challenge, especially for large-scale models. Future work should explore more efficient algorithms or approximations to address this limitation.

References

[1] Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent space oddity: on the curvature of deep generative models. *arXiv preprint arXiv:1710.11379*, 2017.

[2] Randall Balestriero et al. A spline theory of deep learning. In *International Conference on Machine Learning*, pages 374–383. PMLR, 2018.

[3] Elizabeth Clark, Tal August, Sofia Serrano, Nikita Haduong, Suchin Gururangan, and Noah A Smith. All that's' human'is not gold: Evaluating human evaluation of generated text. *arXiv preprint arXiv:2107.00061*, 2021.

[4] Dan Friedman and Adji Bousso Dieng. The vendi score: A diversity evaluation metric for machine learning. *Transactions on Machine Learning Research*, 2023.

[5] Quentin Garrido, Randall Balestriero, Laurent Najman, and Yann Lecun. Rankme: Assessing the downstream performance of pretrained self-supervised representations by their rank. In *International Conference on Machine Learning*, pages 10929–10974. PMLR, 2023.

[6] I. J Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems*, pages 2672–2680. MIT Press, 2014.

[7] Boris Hanin and David Rolnick. Complexity of linear regions in deep networks. *arXiv preprint arXiv:1901.09021*, 2019.

[8] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. *Advances in neural information processing systems*, 33:9841–9850, 2020.

[9] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[10] Ahmed Imtiaz Humayun, Randall Balestriero, Guha Balakrishnan, and Richard G Baraniuk. Splinecam: Exact visualization and characterization of deep network geometry and decision boundaries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3789–3798, 2023.

[11] Ahmed Imtiaz Humayun, Randall Balestriero, and Richard Baraniuk. Magnet: Uniform sampling from deep generative network manifolds without retraining. In *International Conference on Learning Representations*, 2021.

[12] Ahmed Imtiaz Humayun, Randall Balestriero, and Richard Baraniuk. MaGNET: Uniform sampling from deep generative network manifolds without retraining. In *ICLR*, 2022.

[13] Ahmed Imtiaz Humayun, Randall Balestriero, and Richard Baraniuk. Polarity sampling: Quality and diversity control of pre-trained generative networks via singular values. In *CVPR*, pages 10641–10650, 2022.

[14] Ahmed Imtiaz Humayun, Randall Balestriero, and Richard Baraniuk. Deep networks always grok and here is why. *arXiv preprint arXiv:2402.15555*, 2024.

[15] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019.

[16] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proc. CVPR*, pages 8110–8119, 2020.

[17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[18] Line Kuhnel, Tom Fletcher, Sarang Joshi, and Stefan Sommer. Latent space non-linear statistics. *arXiv preprint arXiv:1805.07632*, 2018.

[19] B. Poole, S. Lahiri, M. Raghu, J. Sohl-Dickstein, and S. Ganguli. Exponential expressivity in deep neural networks through transient chaos. *Proc. Neural Inf. Process. Syst. (NIPS'16)*, 2016.

[20] B. Poole, S. Lahiri, M. Raghu, J. Sohl-Dickstein, and S. Ganguli. Exponential expressivity in deep neural networks through transient chaos. *arXiv preprint arXiv:1606.05340*, Dec. 2016.

[21] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. 2022 ieee. In *CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, 2021.

[22] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

[23] Yuxin Wen, Yuchen Liu, Chen Chen, and Lingjuan Lyu. Detecting, explaining, and mitigating memorization in diffusion models. In *The Twelfth International Conference on Learning Representations*, 2024.

A. Appendix / supplemental material

B. Broader Impact Statement

Our proposed framework for assessing and guiding generative models through manifold geometry offers several potential benefits to society. By providing a more objective and automated approach, we can significantly reduce the cost and time associated with human evaluation, making the auditing and mitigation of biases in large-scale models more accessible and efficient. This has implications for promoting fairness and equity in AI systems, particularly in domains where biases can have significant societal consequences.

Furthermore, our approach can empower researchers and practitioners to better understand the relationship between the geometry of learned representations and various aspects of model behavior, such as generation quality, diversity, and bias. This deeper understanding can inform the development of more robust and reliable generative models, leading to advancements in various fields, including art, design, healthcare, and education.

However, we recognize that our approach is not without limitations and potential risks. While it can be a valuable tool for identifying and mitigating biases, it should not and cannot fully replace human annotators, especially in high-risk domains where human judgment and contextual understanding are crucial. Our method focuses on reducing costs and improving the auditing process, but it should not be used as a standalone approach.

Moreover, the increased automation enabled by our approach raises concerns about the potential displacement of human annotators, leading to job losses and economic disruptions. While our method addresses some aspects of model evaluation, it is not comprehensive and cannot assess all facets of model behavior. Therefore, it should be used with caution and in conjunction with other evaluation methods, including human expertise.

C. Extra Figures

Figure 7. ImageNet samples ordered along the columns (from left to right), with increasing local scaling ψ of the Stable Diffusion decoder learned manifold. We observe that ImageNet samples with lower values of ψ contain simpler backgrounds with modal representation of the object category. Conversely for higher ψ we have increasing diversity both in background and foreground features. This is indicative that for non-piecewise linear manifolds like that learned by Stable Diffusion, the local scaling ψ is still inversely proportional to uncertainty.

Figure 8. ImageNet generated samples ordered along the columns (from left [-100.47] to right [-2.288]), with increasing local scaling ψ of the Stable Diffusion decoder. Similar to ImageNet real samples, generated samples with lower values of ψ contain simpler backgrounds with modal representation of the object category. Conversely for higher ψ we have increasing diversity both in background and foreground features. We also observe that with very high ψ values (very uncertain samples) the model focuses on lowest probability regions in the manifold corresponding to strongest anti-mode

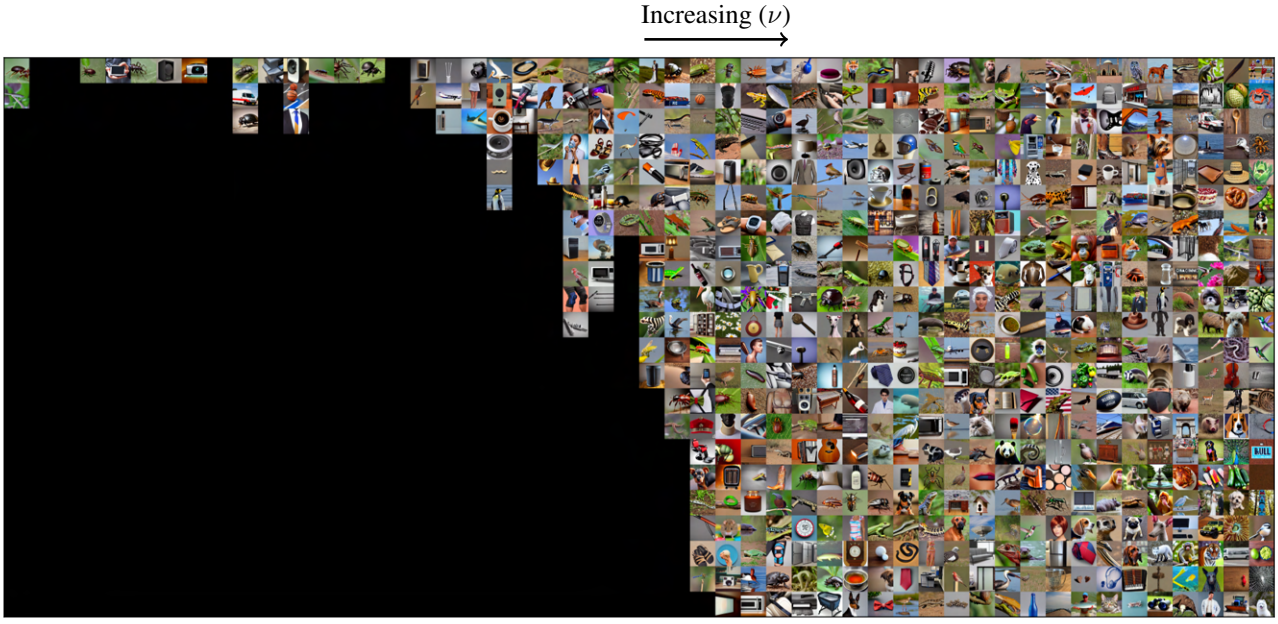

Figure 9. ImageNet generated samples ordered along the columns (from left [109.287] to right [119.752]), with increasing local rank ν of the Stable Diffusion decoder.

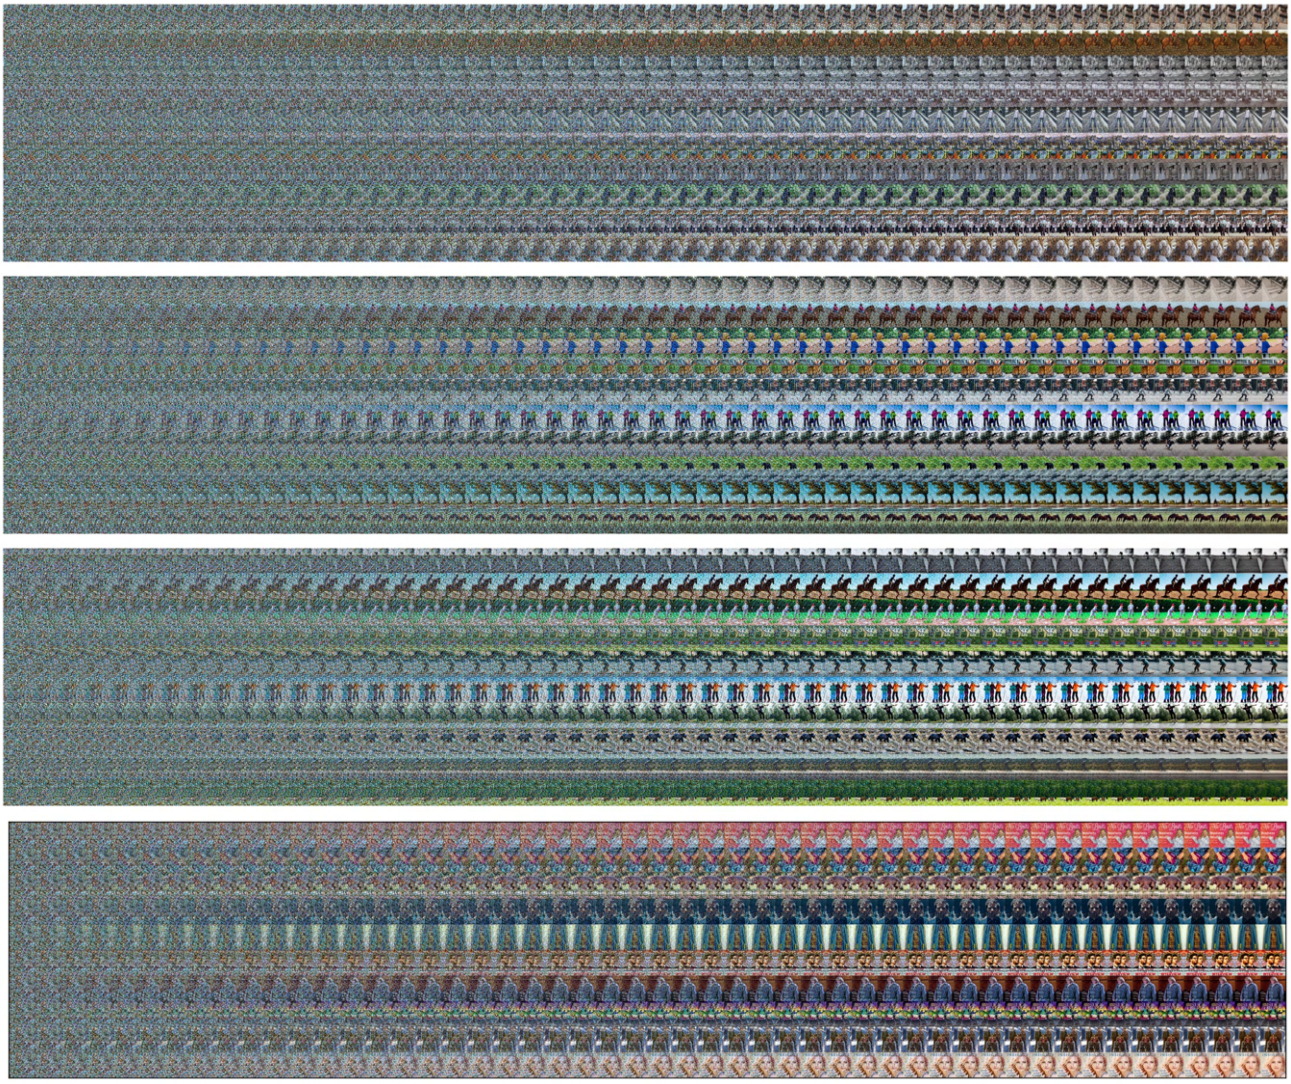

Figure 10. Images generated during 50 diffusion denoising steps for top to bottom, COCO prompts generated with guidance scale 1,5,9 and memorized prompts generated with guidance scale 7.5. Higher guidance scale images, as well as memorized images, tend to resolve faster during the denoising process.

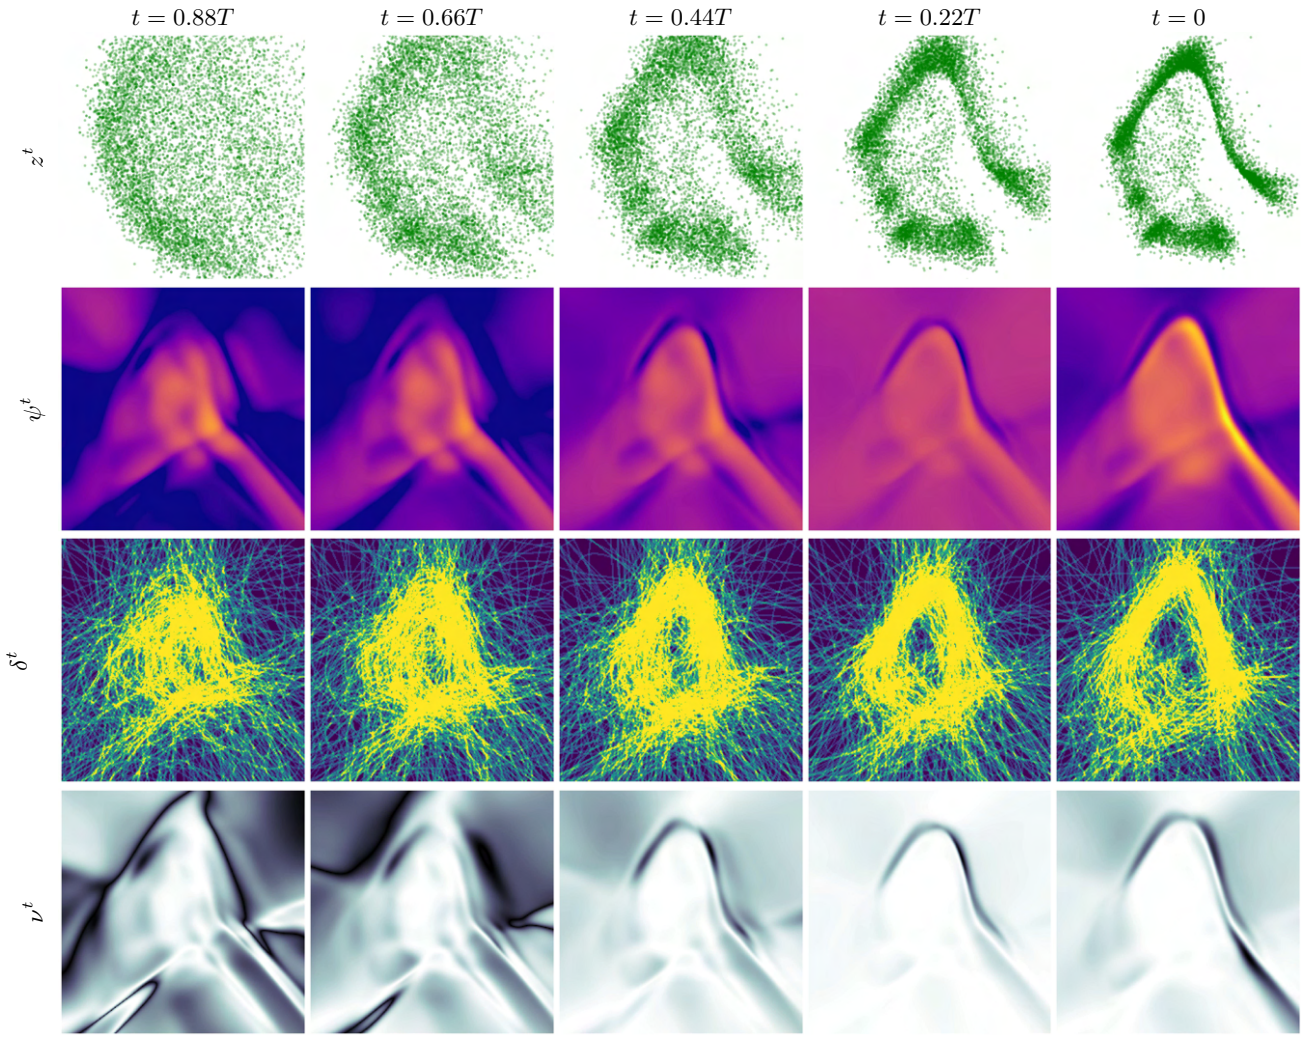

Figure 11. DDPM presented in Figure 2-left-panel, earlier in training after 11395 optimization steps.

Figure 12. Decoded images for the three latent vectors used to determine the 2D subspace in the stable diffusion latent space, presented in Figure 2-right-panel.

Figure 13. Decoded images (right) using 32 latents (left) along a straight line on 2D subspace. Each image bounding box (right) is color coded according to the corresponding latent vector (left).

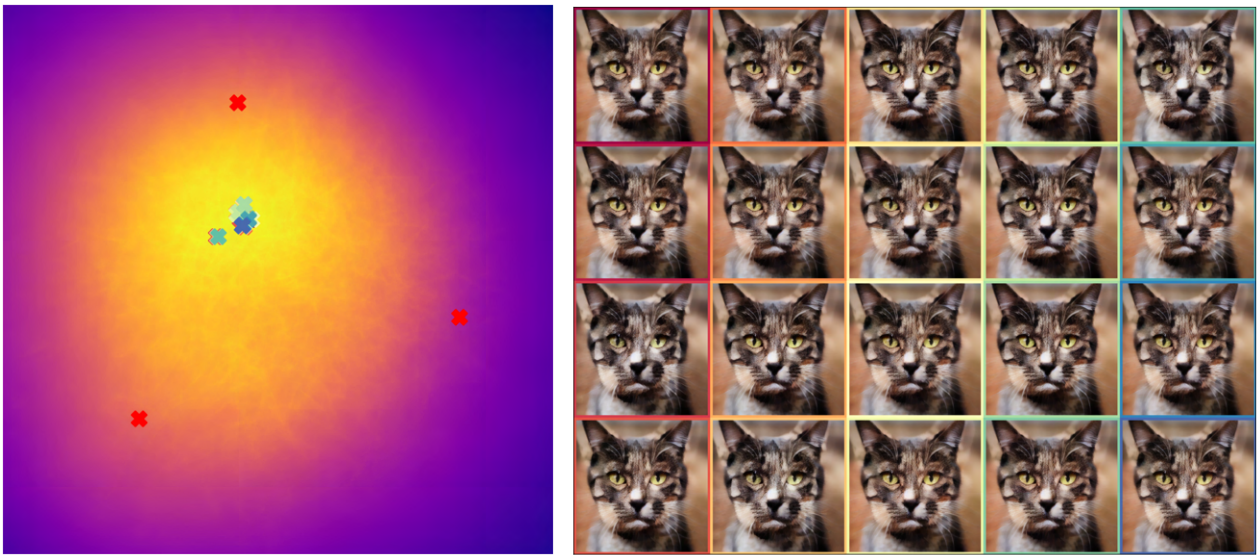

Figure 14. Decoded images (right) using 20 latents (left) from the 2D subspace, with highest ψ. Each image bounding box (right) is color coded according to the corresponding latent vector (left).

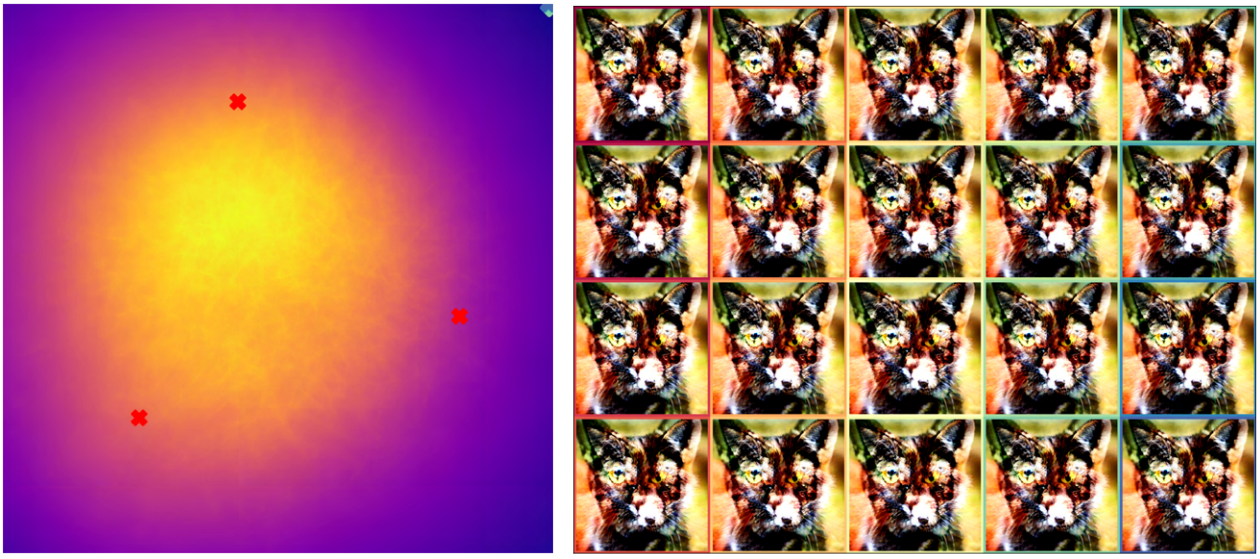

Figure 15. Decoded images (right) using 20 latents (left) from the 2D subspace, with lowest ψ. Each image bounding box (right) is color coded according to the corresponding latent vector (left). Selected latents lie outside the domain of the VQGAN latent space.

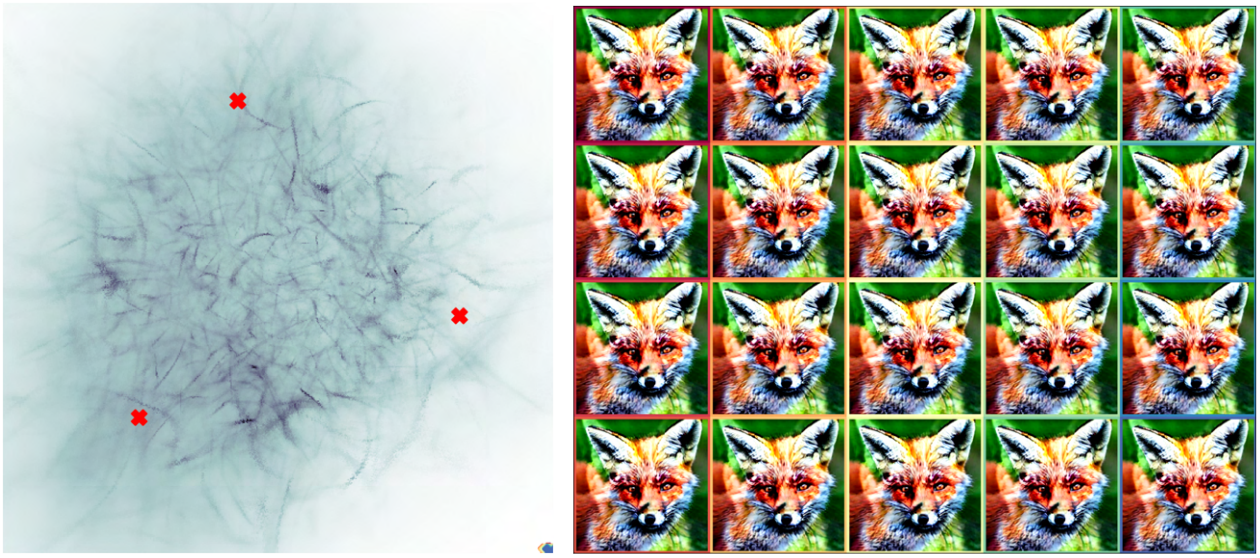

Figure 16. Decoded images (right) using 20 latents (left) from the 2D subspace, with highest ν. Each image bounding box (right) is color coded according to the corresponding latent vector (left). Selected latents lie outside the domain of the VQGAN latent space.

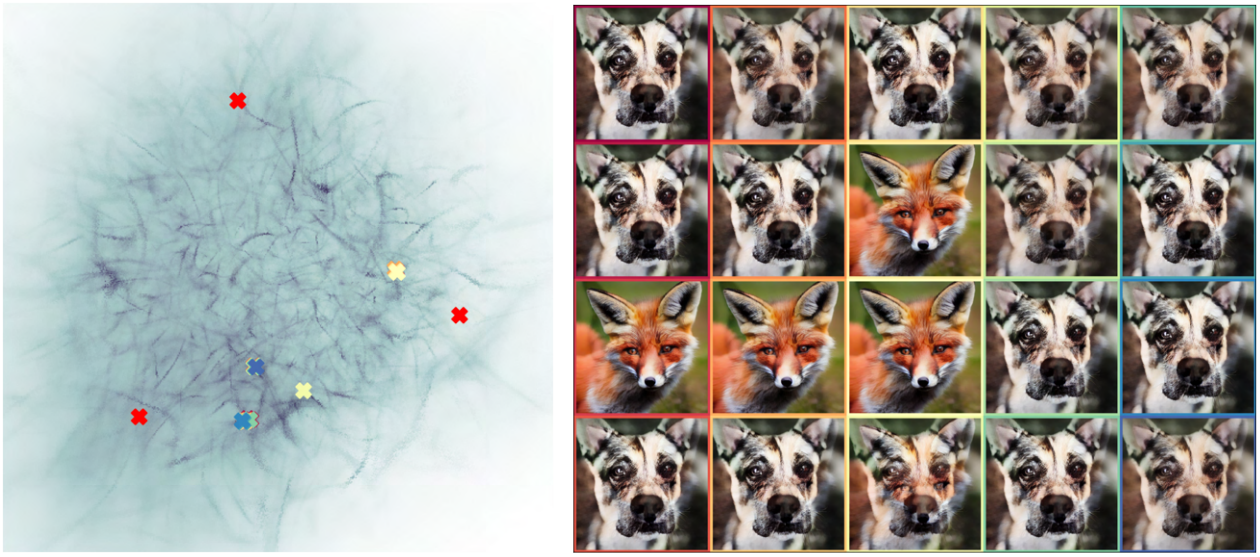

Figure 17. Decoded images (right) using 20 latents (left) from the 2D subspace, with lowest ν. Each image bounding box (right) is color coded according to the corresponding latent vector (left).

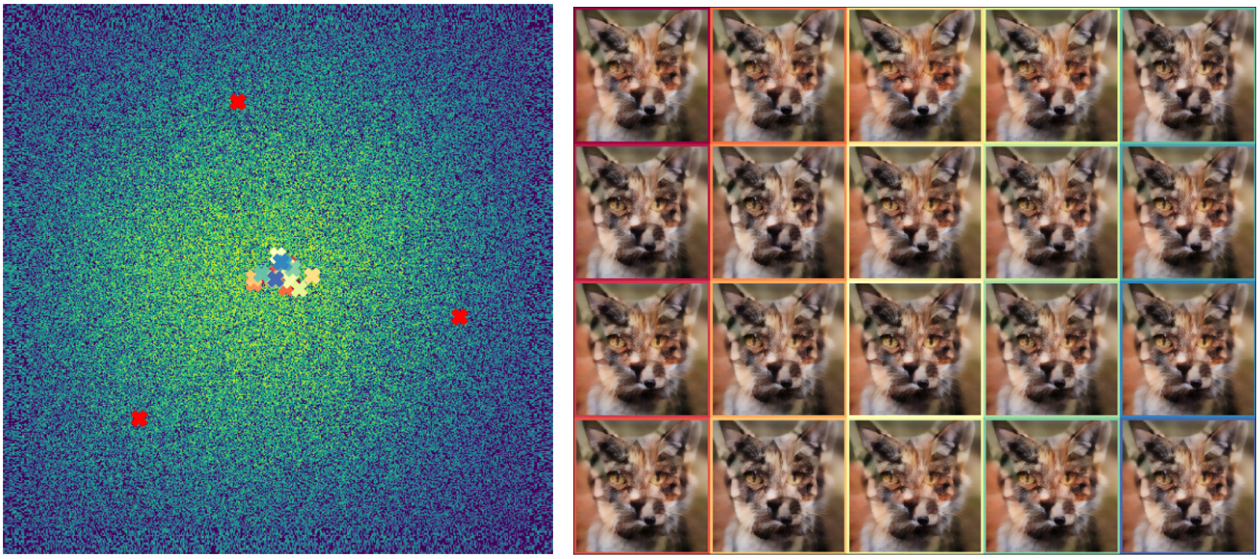

Figure 18. Decoded images (right) using 20 latents (left) from the 2D subspace, with highest δ. Each image bounding box (right) is color coded according to the corresponding latent vector (left).

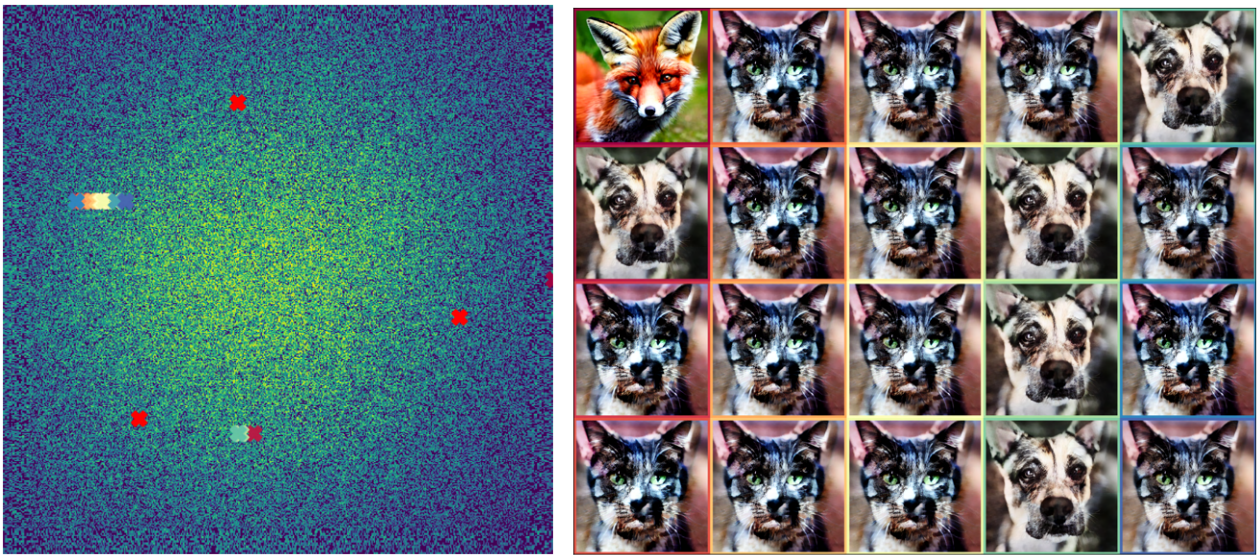

Figure 19. Decoded images (right) using 20 latents (left) from the 2D subspace, with lowest δ. Each image bounding box (right) is color coded according to the corresponding latent vector (left).