
Several questions of visual generation in 2024

Shuyang Gu
Microsoft Research Asia
cientgu@gmail.com

Abstract

This paper does not propose any new algorithms but instead outlines various problems in the field of visual generation based on the author's personal understanding. The core of these problems lies in *how to decompose visual signals*, with all other issues being closely related to this central problem and stemming from unsuitable approaches to signal decomposition. This paper aims to draw researchers' attention to the significance of **Visual Signal Decomposition**.

Due to the limitations of the author's knowledge, the paper does not cover certain critical issues, such as security issues with generative models and data annotation challenges. Moreover, it is possible that many of the opinions in this paper are prejudiced or even incorrect. I sincerely welcome readers to refute them or to discuss with me.

Question 1: What's the goal of generative models?

In my view, generative models are designed to "create what one envisions". The generation of digital signals requires three milestones:

1. The first milestone involves translating the user's thoughts into a format interpretable by a computer [43, 55, 20]. Specifically, this entails identifying the modality of the signal to be generated, which could range from text, 3D objects, to videos, or other forms. Subsequently, one must ascertain the exact state that is to be generated under the given conditions. From the perspective of manifold learning, this process corresponds to first determining the dimensionality D of the encompassing space, followed by identifying the target distribution P_t that the model seeks to emulate.
2. In the field of generative modeling, numerous researchers are engrossed in the challenge of how to model P_t . The principal obstacle is the inherent complexity of the target distribution, which often makes it exceedingly arduous to model. Consequently, researchers are continuously seeking out models with stronger modeling capabilities. Specifically, within the realm of visual generation, prevalent methodologies in recent years include Energy-Based Models [30], Variational Autoencoders (VAEs) [26], Generative Adversarial Networks (GANs) [12], diffusion models [19, 44], and many others.
3. Simultaneously, many researchers further exploring more advanced objectives, investigating how to generate results with greater efficiency [46, 34] and enhanced interpretability [48, 11]. These pursuits have substantial implications for practical deployment, security and the development of responsible AI.

Question 2: The problem of visual signal decomposition.

As discussed in Question 1, the endeavor to model the distribution P_t presents a significant challenge, especially when the goal is to generate convincing textual content or videos. These types of data represent low-dimensional manifolds embedded within extremely high-dimensional spaces, making the direct fitting of these distributions with a single network impractical. Therefore, we need to break

down the complex distribution modeling problem into multiple simpler problems and solve each subproblem separately. This gives rise to a question: *How to effectively decompose this complex distribution modeling problem?*

Language decomposition

The success of large language models is largely attributed to the effective decomposition of textual signals. Consider the task of modeling a text sequence $A = x_0, x_1, x_2, \dots$, where x_t denotes the token at the t -th position from front to back. This can be segmented into a sequence of conditional modeling subtasks based on the position: $p(x_0), p(x_1|x_0), p(x_2|x_1, x_0), \dots$. Large language models employ autoregressive models to approximate these conditional distribution mappings $p(x_t|x_{i < t})$. A critical aspect of this approach is that, for natural language, the decomposed subtasks are intrinsically interrelated. For example, a phrase is "I love swimming," it might appear at the beginning or in the middle of a sentence. In other words, it can appear in any subtask. With a sufficiently large dataset, $p(x_2 = \text{"swimming"}|x_{i < 2} = \text{"I love"})$ and $p(x_7 = \text{"swimming"}|x_{i < 7} = \text{"you love playing basketball, I love"})$ represent two tasks that are very interrelated and can serve as a form of data augmentation to one another. In other words, different subtasks are **"equivariant"**. Let's give "equivariant" a rigorous mathematical definition.

Definition 1 Assume the target distribution is $P_t(x)$, we split the signal into multiple sub-tasks: $p(x_0), p(x_1|x_0), p(x_2|x_1) \dots p(x_{t+1}|x_t) \dots$. For the t -th conditional probability fitting task $p(x_{t+1}|x_t)$, we adopt a network θ_t to fit it. For any two tasks t and k and two state samples s and \hat{s} , if

$$p_{\theta_t}(x_{t+1} = \hat{s}|x_t = s) = p_{\theta_k}(x_{t+1} = \hat{s}|x_t = s) \quad (1)$$

or

$$p_{\theta_t}(x_{t+1} = \hat{s}|x_t = s) = p_{\theta_k}(x_{k+1} = \hat{s}|x_k = s) \quad (2)$$

we refer to this signal decomposition as *equivariant*.

It can be observed that the language decomposition is independent of position. For any token x_t or phrase $\{x_t, x_{t+1}\}$, the probability of their occurrence at the t -th or k -th position in a sentence is nearly identical.¹ This observation aligns with Equation 2, thus suggesting that the language decomposition exhibits equivariance. Therefore, employing a single model to approximate these diverse but related tasks generally does not lead to conflicts, and in fact, it can be highly beneficial for the overall modeling of the data.

Image patch decomposition

This strategy of decomposition and modeling has made remarkable success in language domains. However, replicating this approach to decompose image through spatial position presents distinct challenges. Early attempts [6, 38] involved segmenting images into spatial patches to create sequences that could be processed by autoregressive models. Yet, unlike natural language, image patch inherently lack the "equivariance" characteristic. As illustrated in Figure 1, while there is a continuity within the blocks of a single row, such continuity is absent between the last block of one row and the first block of the subsequent row. In addition to continuity, other dataset-specific challenges persist, such as the propensity for human subjects to be centered within images. Therefore, employing a single, universal model to encapsulate all these varying distributions often results in conflicts. Moreover, the model must endeavor to learn an intricate joint distribution mapping problem that composed of multiple different distributions, which undermines the principle of decomposing complex distributions for simpler modeling. Although the integration of position-embedding can mitigate some of these conflicts, it is not a panacea for the issue.

Depth decomposition

Beyond spatial division, some researchers have explored splitting image data along the depth dimension. This approach is intuitively appealing, given the large amount of information present at each spatial location in an image, such as the need for a three-channel RGB representation with 8 bits per

¹Although the word "I" may have a higher probability of appearing in the first position than others, given the typically extensive length of linguistic texts, it can be approximated as identical.

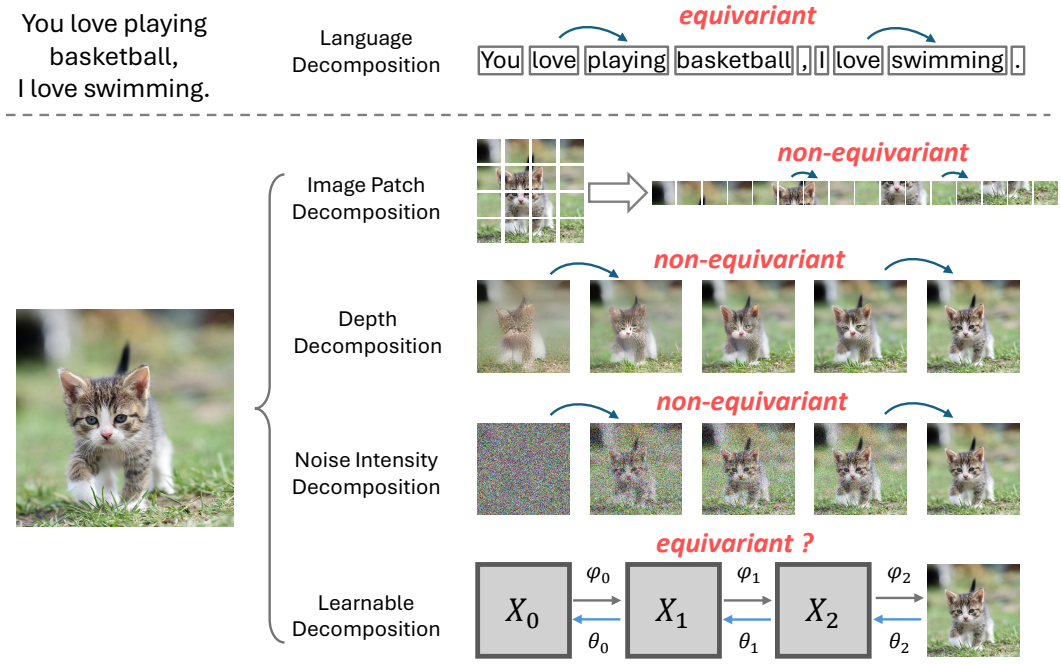


Figure 1: Illustration of Language and Visual Signal Decomposition. The language decomposition is equivariant, while most image decomposition are non-equivariant.

channel. On the other hand, the high spatial dimensionality of images often necessitates the use of Variational Autoencoders (VAE) [26] to reduce the dimension. Combining these two aspects, notable methods that partition data by depth include VQVAE2 [39] and RQVAE [31]. Firstly, such approaches may encounter the issue of "invalid encoding," which we will elaborate in Question 3. Furthermore, unlike natural language, the depth dimension also does not exhibit "equivariance". Taking RQVAE as an example, at the same spatial position, earlier tokens represent low-frequency information while later tokens represent high-frequency information. Therefore, applying autoregressive models with shared parameters [31, 52] to model these varied distribution mappings can result in conflicts. An alternative approach, such as MUSE [5], proposes segmenting the problem into a smaller number of unique subtasks, each modeled independently without parameter sharing. However, as the data distributions become increasingly complex, this may necessitate a larger set of subtasks, leading to a surge in the required model parameters and potentially exacerbating the "invalid encoding" problem.

Noise intensity decomposition

Diffusion models propose another intriguing approach to signal decomposition: characterizing images through a series of progressively denoised image sequences. For a given image x_0 in a dataset, noise is sequentially added via a Markov process to produce a sequence $x_0, x_1, x_2, \dots, x_N$, with x_N being nearly pure noise, with minimal remnants of the original image information. This process effectively breaks down the task of modeling the image distribution into N denoising subtasks: $p(x_t|x_{t+1})$, for $t = [0, 1, \dots, N - 1]$. While all subtasks are concerned with denoising, parameter sharing among them might seem practical in theory. However, in practice, typical noise-adding strategies can lead to discrepancies during the denoising stages, especially when noise levels differ significantly [16]. These non-equivariant tasks lead to a predicament akin to depth-based decomposition challenge: employing a model with shared parameters to fit complex data distribution mappings is a daunting task for the model’s capability. If parameters are not shared, this could inflate the model size rapidly. Some studies, such as eDiff-I [1], have attempted to balance the parameter efficiency with the complexity of fitting distributions from an implementation standpoint. Additionally, the reparameterization trick [35, 41, 33] has proven to be an extremely important technique to unify the output distributions of different denoising task, which alleviates the conflicts across different noise intensities. However, it does not eliminate the discrepancy of the input distributions. These concerns regarding noise intensity conflicts within diffusion models will be further explored in Question 4.

Learnable decomposition

Upon reviewing the diffusion models, we find that the extent of conflicts is determined by the chosen noise strategy, which is typically predefined manually. Therefore, some researchers have attempted to define superior noise strategies, endeavoring to ensure a certain level of similarity in the denoising process across varying noise levels. Notable works in this area include Flow Matching [33] and consistency models [46]. Concurrently, other researchers are exploring whether the adding noise strategy can be learned rather than preset. Noteworthy advancements in this domain comprise Variational Diffusion Models (VDM) [25] and Diffusion Schrödinger Bridge (DSB) [9, 51], although not all such works were initiated pursued this objective. In particular, VDM focuses on learning the coefficients for adding Gaussian noise, which somewhat narrows the potential for learning to mitigate conflicts. Meanwhile, studies based on the Schrödinger Bridge paradigm employ specialized networks to learn the adding noise process, iteratively approximating the entropy-regularized optimal transport. However, current learnable decomposition methods are not designed according to the characteristic of "equivariance". Future work may need to consider this as a prior, to constrain the network learning in the learnable decomposition. Moreover, a significant trade-off arises when substituting a predefined SDE with a network for noise learning: it's challenging to leverage the reparameterization to unify output distributions, which is a pivotal technique to mitigate conflicts across different noise intensities in practical applications. Although recent efforts [42, 51] have been made to address these issues, they still remain insufficient for practical application.

Extended Discussion

From the signal decomposition perspective, the debate over whether autoregressive (AR) models, diffusion models, or other model architectures are superior for visual generation is not particularly productive. The essential consideration is how the signal is decomposed and the selection of a generative paradigm whose inductive bias is conducive to the chosen decomposition approach. Ideally, there may be two paradigms for decomposition: one resembles the approach taken by language models, which simplifies complex data distributions into a series of simpler, conditional data distributions that exhibit "equivariance". The other strategy involves breaking down the data into multiple independent distribution problems, which can be viewed as a special instance of "equivariance".

We argue that the difficulty of achieving equivariance in image decomposition is not entirely due to the fact that images are two-dimensional data while language is one-dimensional. Recent work [56] attempts to encode images to one-dimensional tokens, but these one-dimensional tokens are neither independent nor equivariant.

While learnable decomposition methods theoretically hold the potential of achieving this "equivariance", their practical application is fraught with challenges currently. Another feasible approach might be to integrate various signal decomposition techniques to simplify the data distribution. For instance, videos can be decomposed into temporally "equivariant" frames, which can then be further broken down based on "noise intensity" [2] or "image patches" [28]. Similarly, MUSE [5] initially decomposes the image signal along the depth dimension and then tackles the distribution mapping from the "noise intensity" dimension.

We argue the signal decomposition is the fundamental problem, and many of the subsequent questions can be seen as extensions of it, aims to alleviate the issue of non-equivariance in the current visual signal decomposition.

Question 3: The tokenization problem.

Current mainstream generative models for images and videos predominantly adopt a two-stage approach: initially encoding the data into a compact, low-dimensional representation, then modeling this compressed distribution. The purpose of the first compression phase is to streamline the data distribution while preserving as much of the original information as possible. This simplification is intended to alleviate the complexity faced during the model fitting phase that follows. In the context of textual data, the dimensionality reduction can be considered lossless. In contrast, visual data compression, whether through Autoencoders (AE) or Variational Autoencoders (VAE), is inherently lossy. However, the assertion that "the less lossy the compression, the better" is not necessarily valid.




Figure 2: Illustration of the invalid encoding issue. The top two figure represent the reconstruction FID of RQVAE on ImageNet with different depths, the bottom figure shows the visualization results on FFHQ. The reconstruction will not continue to improve after the dominate stage.

Depth	4	8	16
Invalid encoding ratio	0.312	0.4	0.571

Table 1: The longer the encoding length, the higher probability of encountering the invalid encoding issue.

A typical example is that signals compressed with AE might reconstruct better than those with VAE, but the complexity retained in the compressed data distribution still poses challenges for the fitting process in the second stage. Therefore, researchers introduce regularization constraints [53, 40, 14] in the latent space during compression to prevent the data distribution from becoming overly complex.

The conflict between reconstruction fidelity and fitting difficulty is a familiar challenge across various fields. In audio processing, continuous audio signals are typically encoded into tokens of length 16, with the fitting phase often concentrating on the foremost 8 tokens [54]. Similarly, in the realm of image processing, GLOW [27] and VDM++ [24] also discovered that training on 5-bit images yields better results than using full 8-bit depth. These observations underscore the potential benefits of employing an adaptive-length encoding strategy to better balance the precision of reconstruction against the complexity of the second stage fitting task.

A typical example of variable-length encoding is the RQVAE [31], which iteratively encodes the reconstruction error within the latent space, aims for achieving increasingly more accurate image reconstructions. However, as illustrated in Figure 2, we observed that deeper encoding does not always correlate with better reconstruction quality. We refer to this as the "invalid encoding" issue. To investigate it, we conducted a series of comprehensive experiments, modifying network architectures, learning rates, loss function weights, and codebook sizes, in addition to measuring the frequency of "invalid encoding" across a range of encoding lengths. Table 1 illustrates that the longer the encoding length, the higher probability of encountering the invalid encoding issue. However, we have not ascertained any overarching conclusions regarding the specific conditions that invariably lead to this problem. Below, we offer an intuitive but not rigorous explanation for this issue:

Let D represent the decoder and I represent the original input image. The encodings at different depths are denoted by x_0, x_1, \dots, x_N , where N is the depth of encoding, we assume it to 4 in this case. Thus, the reconstruction loss L for the RQVAE can be considered as the combination of the

following four reconstruction losses:

$$\begin{cases} L_0 = \|D(x_0) - I\|_2 \\ L_1 = \|D(x_0 + x_1) - I\|_2 \\ L_2 = \|D(x_0 + x_1 + x_2) - I\|_2 \\ L_3 = \|D(x_0 + x_1 + x_2 + x_3) - I\|_2 \end{cases} \quad (3)$$

Building on this, we make two assumptions for simplifying the analysis. First, we hypothesize that the decoder functions as a linear transformation, allowing for a more tractable examination of the results. Second, in alignment with standard configurations, we assign equal loss weights to the four losses. With these assumptions, the calculation of the aforementioned reconstruction loss can be simplified as follows:

$$L_{all} = \|D(x_0 + \frac{3}{4}x_1 + \frac{2}{4}x_2 + \frac{1}{4}x_3) - I\|_2 \quad (4)$$

Therefore, the representation in the latent space that minimizes the image level reconstruction loss would be:

$$\arg \min X = x_0 + \frac{3}{4}x_1 + \frac{2}{4}x_2 + \frac{1}{4}x_3 \quad (5)$$

This cannot guarantee that $x_0 + x_1 + x_2 + x_3$ is closer to the $\arg \min X$ than $x_0 + x_1 + x_2$. Assuming the encodings at various depths share a common codebook and are independently and identically distributed, the latter sum would definitively be even closer to the ground truth than the former. Therefore, this leads to the invalid encoding issue.

Question 4: Is diffusion model a maximize likelihood model?

Autoregressive models are classical maximum likelihood models that facilitate various sophisticated tasks by calculating the likelihood function, including manipulating the generated results and assessing their quality. An interesting question arises: *Can diffusion models also be considered maximum likelihood models?* The initial study on Denoising Diffusion Probabilistic Models [19] originated from the maximum likelihood to derive the training loss function. [45] proposes a loss weight setting for ELBO training. Moreover, VDM++ [24] shows that optimizing with monotonic weightings effectively equates to maximizing the likelihood function with a distribution augmentation. However, in practical training, different loss weights are often employed. The current mainstream practice, exemplified by SD3 [10], does not fully embrace this principle.

Coincidentally, a similar perplexity appears in both the generation and evaluation process. During generation, it has been observed that directly sampling from the likelihood model $p_\theta(x_{t-1}|x_t, c)$ yields inferior results compared to those modified by classifier-free guidance [18]: $p_\theta(x_{t-1}|x_t) + \lambda(p_\theta(x_{t-1}|x_t, c) - p_\theta(x_{t-1}|x_t))$, where λ is the classifier-free guidance scale. This can be derived [49] as sampling from $p_\theta(x_{t-1}|x_t, c) * p_\theta(c|x_{t-1})^{\lambda-1}$. We can easily find that this adjustment merges the likelihood function with the posterior distribution, intimating that maximization of likelihood does not always equate to the best result. This issue is further evidenced during the evaluation phase, wherein models with lower Negative Log Likelihood (NLL) scores do not consistently correspond to the most aesthetically pleasing visual results or the lowest Fréchet Inception Distance (FID) metrics [19]. This leads to a subtle yet critical question: *Why maximize the likelihood not necessarily lead to optimal results?*

Here is a possible understanding. As illuminated in [21], score matching is closely related to the maximization of non-normalized likelihood. Typically, score matching can avoid the proclivity in maximum likelihood learning to assign equal probabilities to all data points. In certain special cases, such as Multivariate Gaussian distributions, they are equivariant. VDM++ elucidates that training with monotonic loss weights $w(t)$ indeed equates to maximizing the ELBO for all intermediate states. The particular weighting indicates the varying significance of different noise levels on the ultimate model performance. However, as discussed in Question 2, the image data lacks the "equivariance". In practical training, The **difficulty** of learning the likelihood function varies with

noise intensity; intuitively, the greatest difficulty appears at medium noise levels [10], where the likelihood function tends to be less accurately learned. During the generation process, the use of classifier-free guidance can be interpreted as a rectification for the suboptimally learned likelihood function. This is particularly evident in [29], the classifier-free guidance is extremely crucial at medium noise levels. During model evaluation, given that tasks at different noise levels have varying degrees of **importance** to the final result, applying uniform weighting to these NLL losses could not effectively measure the quality of the final generated output.

Question 5: For diffusion model, how to balance the conflicts among different SNRs?

As discussed previously, diffusion models differ from autoregressive models in text generation in that they do not maintain "equivariance" across various subtasks. Some studies categorize diffusion models by the noise intensity and explicitly utilize the Mixture of Experts (MOE) strategy for model fitting. Works such as eDiff-I [1] and SDXL [37] exemplify this approach, with each model do not share parameters. The key of these approaches is the strategic partition of tasks since subtasks not only conflict with each other but also possess relevance. By leveraging these interrelationships, it is possible to improve the efficiency of model convergence and inhibit the exponential increase [1] of model parameters.

Other methods attempts to reconcile the conflicts among different noise intensities without increasing the number of parameters. According to VDM++ [24], the training objective is the combination of loss weighting and importance sampling:

$$\mathcal{L}(x) = \frac{1}{2} \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \mathbf{I}), \lambda \sim p(\lambda)} \left[\frac{w(\lambda)}{p(\lambda)} \|\hat{\epsilon}_\theta(z_\lambda; \lambda) - \epsilon\|_2^2 \right] \tag{6}$$

where λ denotes the log of Signal-to-Noise Ratios (SNRs), x denotes the training image, z_λ is the noisy image at noise intensity λ . $w(\lambda)$ and $p(\lambda)$ denotes the loss weight and sampling frequency at noise level λ .

Therefore, to balance various noise intensities, one might adjust the loss weights or implement importance sampling for different SNRs. MinSNR [16] is a notable study that carefully designs the loss weights, aiming to circumvent conflicts by pursuing Pareto optimal optimization directions. Studies such as SD3 [10] and HDiT [8] have empirically found that increasing the weights for SNRs in the medium range can lead to improved outcomes. As depicted in Equation 6, adjusting the loss weights $w(\lambda)$ has similar effects of modifying the sampling frequency $p(\lambda)$. However, in practice, increasing the loss weight for important tasks is equivariant to raising the learning rate, while enhancing the frequency can be seen as dedicating more computational resources (Flops), which often results in better performance.

From another perspective, performing importance sampling on different noise levels can be considered as designing a noise schedule [15], or a type of signal decomposition, as we discussed in Question 2. When applying independent and identically distributed (i.i.d.) Gaussian noise across different spatial positions, previous studies have indicated the necessity to adapt the noise schedule based on token length [7], and ensuring there is no signal leakage in the final step [32, 50]. [15] empirically introduced a noise schedule for training that achieves more efficient convergence and has proven its effectiveness across diverse conditions. However, it may need to adjust hyperparameters based on the target distribution, and lacks analysis of the conflicts during inference. The author speculate that abandoning the use of i.i.d. Gaussian noise for signal decomposition might be a fundamental approach to resolving the conflicts.

Question 6: Is there a scaling law for diffusion models?

Upon rethinking the remarkable achievements of Large Language Models (LLMs), one pivotal factor is the principle of the scaling law. This naturally raises a question: *Do diffusion models in visual generation also conform to a scaling law?*




Figure 3: The diffusion scaling law rely on importance weighted loss. For the left figure, $y = 0.369 * x^{-0.030}$. For the right figure, $y = 0.303 + 1.4 * 10^5 * x^{-0.140}$.

The challenge in addressing this question lies in the lack of an evaluation metric that aligns with human perception to assessing the performance of the model. In language modeling through autoregressive models, the principle of "equivariance" facilitates the evaluation of model performance by assigning equal weight to the negative log likelihood loss across different subtasks during both training and evaluation phases. In contrast, the subtasks within diffusion models lack this equivariance; they each contribute to the final generated output in varying degrees. Therefore, a simple, equal-weighted aggregation of losses fails to capture the full extent of a generative model's capabilities.

To uncovering an appropriate metric, the first attempt is to construct an importance coefficient for various subtasks. The novel noise schedule defined in [15] can be considered as assigning "difficulty coefficients" to distinct tasks. We treat it as an "importance coefficient", and employing it to weight the losses at different noise intensities. For training text2image models, we employed COYO [4] dataset, which consists of 700M text-image pairs. We trained four models with parameter counts of 32.28M, 128.56M, 454.98M, and 671.32M respectively. For convenience, we denote them as S, M, L, and XL. All the models are trained with 1024 batchsize. We use the "importance weighted loss" as metric to measure the model's performance. Following the approach outlined in [47], we utilize the formula from [22] to estimate the performance of the XL model based on the S, M, L models. The results are shown in Figure 3. The left figure predicts the model's performance according to the number of parameters, while the right figure predicts the model's performance in relation to the number of training iterations. It can be seen that both predictions are quite accurate, although the scale of verification is relatively limited due to resource constraints. Nevertheless, it's essential to note that, there is no direct evidence to affirm that the metric aligns with human judgment.

The second strategy employs established metrics for evaluating generative models, with Fréchet Inception Distance (FID) [17] being the most prevalent. FID is designed to quantify the disparity between two distributions of data. However, when dealing with large-scale generative models and exceedingly complex data distributions, it becomes challenging to accurately capture the target distribution with a finite sample set, inevitably leading to biases in the FID score. Furthermore, FID assumes that the feature vectors extracted from the neural network follows Gaussian distribution, which introduces significant systematic errors. It is also important to note that some studies [36, 13, 3] have highlighted further issues associated with the FID metric.

To ensure alignment with human preferences, the most crucial approach involves using extensive human annotation as the benchmark for evaluation. Taking text2image generation as an example, one potentially approach is to collect a large number of high-quality text-image pairs. For the generative models to be tested, they can generate results based on the given text prompts. Users evaluate which image, the generated one or the ground truth image, better satisfies their preferences. In theory, as the quality of the models improves, this preference rate of generated results should converge towards 0.5. This preference rate could then act as an indicator for the scaling law, providing insights into how computational resources, model size, and data size affect the model's final performance. It is important to note that the quality of outputs from diffusion-based vision models is highly sensitive to the chosen inference strategy [23, 18], which differs significantly from large language models. With a metric that can capture the human preference, this factor merits further exploration.

Acknowledgements

I would like to thank Tiankai Hang for the consistent discussions during this work. I would also want to thank all the friends who have discussed these issues with me, including Dong Chen, Zhicong Tang, Zheng Zhang, Jianmin Bao, and Chengyi Wang. Moreover, we sincerely hope that readers will point out any errors or inadequacies in our understanding, or feel free to directly contact me, to collectively strive to resolve these challenges in the field of visual generation.

Finally, if you plan to cite this paper and use an abbreviation, hope you can use VSD or Visual Signal Decomposition. The author would greatly appreciate it.

References

- [1] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Qinsheng Zhang, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, et al. ediff-i: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022.
- [2] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.
- [3] Ali Borji. Pros and cons of gan evaluation measures. *Computer vision and image understanding*, 179: 41–65, 2019.
- [4] Minwoo Byeon, Beomhee Park, Haechon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. Coyo-700m: Image-text pair dataset. <https://github.com/kakaobrain/coyo-dataset>, 2022.
- [5] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation via masked generative transformers. *arXiv preprint arXiv:2301.00704*, 2023.
- [6] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *International conference on machine learning*, pages 1691–1703. PMLR, 2020.
- [7] Ting Chen. On the importance of noise scheduling for diffusion models. *arXiv preprint arXiv:2301.10972*, 2023.
- [8] Katherine Crowson, Stefan Andreas Baumann, Alex Birch, Tanishq Mathew Abraham, Daniel Z Kaplan, and Enrico Shippole. Scalable high-resolution pixel-space image synthesis with hourglass diffusion transformers. In *Forty-first International Conference on Machine Learning*, 2024.
- [9] Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.
- [10] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024.
- [11] Zigang Geng, Binxin Yang, Tiankai Hang, Chen Li, Shuyang Gu, Ting Zhang, Jianmin Bao, Zheng Zhang, Houqiang Li, Han Hu, et al. Instructdiffusion: A generalist modeling interface for vision tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12709–12720, 2024.
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- [13] Shuyang Gu, Jianmin Bao, Dong Chen, and Fang Wen. Giga: Generated image quality assessment. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 369–385. Springer, 2020.
- [14] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10696–10706, 2022.
- [15] Tiankai Hang and Shuyang Gu. Improved noise schedule for diffusion training. *arXiv preprint arXiv:2407.03297*, 2024.
- [16] Tiankai Hang, Shuyang Gu, Chen Li, Jianmin Bao, Dong Chen, Han Hu, Xin Geng, and Baining Guo. Efficient diffusion training via min-snr weighting strategy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7441–7451, 2023.
- [17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [18] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.

- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [20] Leigh R Hochberg, Daniel Bacher, Beata Jarosiewicz, Nicolas Y Masse, John D Simeral, Joern Vogel, Sami Haddadin, Jie Liu, Sydney S Cash, Patrick Van Der Smagt, et al. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375, 2012.
- [21] Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- [22] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [23] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022.
- [24] Diederik Kingma and Ruiqi Gao. Understanding diffusion objectives as the elbo with simple data augmentation. *Advances in Neural Information Processing Systems*, 36, 2024.
- [25] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in neural information processing systems*, 34:21696–21707, 2021.
- [26] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [27] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018.
- [28] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Rachel Hornung, Hartwig Adam, Hassan Akbari, Yair Alon, Vighnesh Birodkar, et al. Videopoet: A large language model for zero-shot video generation. *arXiv preprint arXiv:2312.14125*, 2023.
- [29] Tuomas Kynkäänniemi, Miika Aittala, Tero Karras, Samuli Laine, Timo Aila, and Jaakko Lehtinen. Applying guidance in a limited interval improves sample and distribution quality in diffusion models. *arXiv preprint arXiv:2404.07724*, 2024.
- [30] Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and Fugie Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
- [31] Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11523–11532, 2022.
- [32] Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and sample steps are flawed. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 5404–5411, 2024.
- [33] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- [34] Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14297–14306, 2023.
- [35] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pages 8162–8171. PMLR, 2021.
- [36] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11410–11420, 2022.
- [37] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sd-xl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023.
- [38] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International conference on machine learning*, pages 8821–8831. Pmlr, 2021.

- [39] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019.
- [40] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [41] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.
- [42] Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36, 2024.
- [43] Ben Shneiderman. The future of interactive systems and the emergence of direct manipulation. *Behaviour & Information Technology*, 1(3):237–256, 1982.
- [44] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [45] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in neural information processing systems*, 34:1415–1428, 2021.
- [46] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023.
- [47] Hui Su, Zhi Tian, Xiaoyu Shen, and Xunliang Cai. Unraveling the mystery of scaling laws: Part i. *arXiv preprint arXiv:2403.06563*, 2024.
- [48] Raphael Tang, Linqing Liu, Akshat Pandey, Zhiying Jiang, Gefei Yang, Karun Kumar, Pontus Stenetorp, Jimmy Lin, and Ferhan Ture. What the daam: Interpreting stable diffusion using cross attention. *arXiv preprint arXiv:2210.04885*, 2022.
- [49] Zhicong Tang, Shuyang Gu, Jianmin Bao, Dong Chen, and Fang Wen. Improved vector quantized diffusion models. *arXiv preprint arXiv:2205.16007*, 2022.
- [50] Zhicong Tang, Shuyang Gu, Chunyu Wang, Ting Zhang, Jianmin Bao, Dong Chen, and Baining Guo. Volumediffusion: Flexible text-to-3d generation with efficient volumetric encoder. *arXiv preprint arXiv:2312.11459*, 2023.
- [51] Zhicong Tang, Tiankai Hang, Shuyang Gu, Dong Chen, and Baining Guo. Simplified diffusion schrödinger bridge. *arXiv preprint arXiv:2403.14623*, 2024.
- [52] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *arXiv preprint arXiv:2404.02905*, 2024.
- [53] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- [54] Chengyi Wang, Sanyuan Chen, Yu Wu, Ziqiang Zhang, Long Zhou, Shujie Liu, Zhuo Chen, Yanqing Liu, Huaming Wang, Jinyu Li, et al. Neural codec language models are zero-shot text to speech synthesizers. *arXiv preprint arXiv:2301.02111*, 2023.
- [55] Jonathan R Wolpaw, Niels Birbaumer, Dennis J McFarland, Gert Pfurtscheller, and Theresa M Vaughan. Brain–computer interfaces for communication and control. *Clinical neurophysiology*, 113(6):767–791, 2002.
- [56] Qihang Yu, Mark Weber, Xueqing Deng, Xiaohui Shen, Daniel Cremers, and Liang-Chieh Chen. An image is worth 32 tokens for reconstruction and generation. *arXiv preprint arXiv:2406.07550*, 2024.