



HAL
open science

Agnostic latent diversity enhancement in generative modeling

Mariia Zameshina, Mathurin Videau, Alessandro Leite, Marc Schoenauer, Laurent Najman, Olivier Teytaud

► **To cite this version:**

Mariia Zameshina, Mathurin Videau, Alessandro Leite, Marc Schoenauer, Laurent Najman, et al.. Agnostic latent diversity enhancement in generative modeling. 2024. hal-04661473v2

HAL Id: hal-04661473

<https://hal.science/hal-04661473v2>

Preprint submitted on 28 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Agnostic latent diversity enhancement in generative modeling

Mariia Zameshina
Meta AI
Univ Gustave Eiffel, CNRS, LIGM

Mathurin Videau
Meta AI
TAU, Inria Saclay, LISN

Alessandro Leite
TAU, Inria Saclay, LISN

Marc Schoenauer
TAU, Inria Saclay, LISN

Laurent Najman
Univ Gustave Eiffel, CNRS, LIGM

Olivier Teytaud
Meta AI

Abstract

Generative modeling methods can generate images from textual or visual inputs. However, diversity in the generated images persists as a major challenge of the existing approaches. In this work, we address this issue head-on by demonstrating that: (a) the diversity of a generated batch of images is intrinsically linked to the diversity within the latent space; (b) leveraging the geometry of the latent space, we can establish an effective metric for quantifying diversity; and (c) employing this insight allows one to achieve a significantly enhanced diversity in image generation beyond the capabilities of traditional random independent sampling. This advancement is consistent across a variety of generative models, including Generative Adversarial Networks and Latent Diffusion Models.

1. Introduction

Latent generative modeling involves learning a latent representation of the data that captures its underlying structure and using this representation to generate new data points. Specifically, the term “latent” refers to hidden variables or features that are inferred from the observed data without being directly observed. In this context, latent generative modeling aims to learn a model that map observed data into a latent space (LS), where each point represents a set of underlying features or factors that explain the variation in the data. This latent space has typically a lower dimension than the original one, which helps one capturing essential features while reducing redundancy. Examples of latent generative models include Variational Autoencoders (VAEs) [9] and Generative Adversarial Networks (GANs) [5]. A major application of latent generative modeling is text-to-image synthesis [14, 16], frequently referred to as “text2image”. This is a type of generative modeling where the goal is to generate realistic images from textual descriptions. In this context, a model takes a textual description as input and gen-

erates an image that corresponds to that description. While latent generative modeling works well for text2image [16], it suffers from diversity loss and mode collapse [1, 6, 8, 10]. To alleviate these issues, we propose to replace the regular, pure random generator of point configurations by another generator, that produces a well-distributed point configuration. In particular, well-distributed point configurations [12] refer to the arrangements of points in a space where the points are spread out evenly and uniformly, often with respect to certain criteria or constraints.

The concept of well-distributed point configurations depends on the context and the specific requirements of the application. Some common criteria for defining well-distributed point configurations include:

- (1) **Uniformity**: Points should be distributed uniformly across the space, meaning that there are no regions with significantly higher or lower point density compared to others.
- (2) **Packing density**: Points should be packed densely enough to cover the space adequately but not so densely that they become clustered or overlapping. Achieving an optimal packing density depends on the dimensions of the space and the desired properties of the point configuration.
- (3) **Symmetry**: In some cases, symmetry or regularity in the arrangement of points is desired, especially in geometric modeling or tessellation applications.
- (4) **Smoothness**: Points should be distributed smoothly across the space, without abrupt changes in density or clustering.

Well distributed point configurations have been applied to many settings [12]. The literature on well distributed point configurations contains only few tools for the case of small point sets in high-dimensional spaces, and our goal is to fill this gap. Figure 1 provides a general scheme for the application of well distributed point configurations to latent generative modeling to improve the diversity of generated images, instead of 50 random independent images usually used by existing approaches. More precisely, a well distributed point configurations is generated in the LS, so

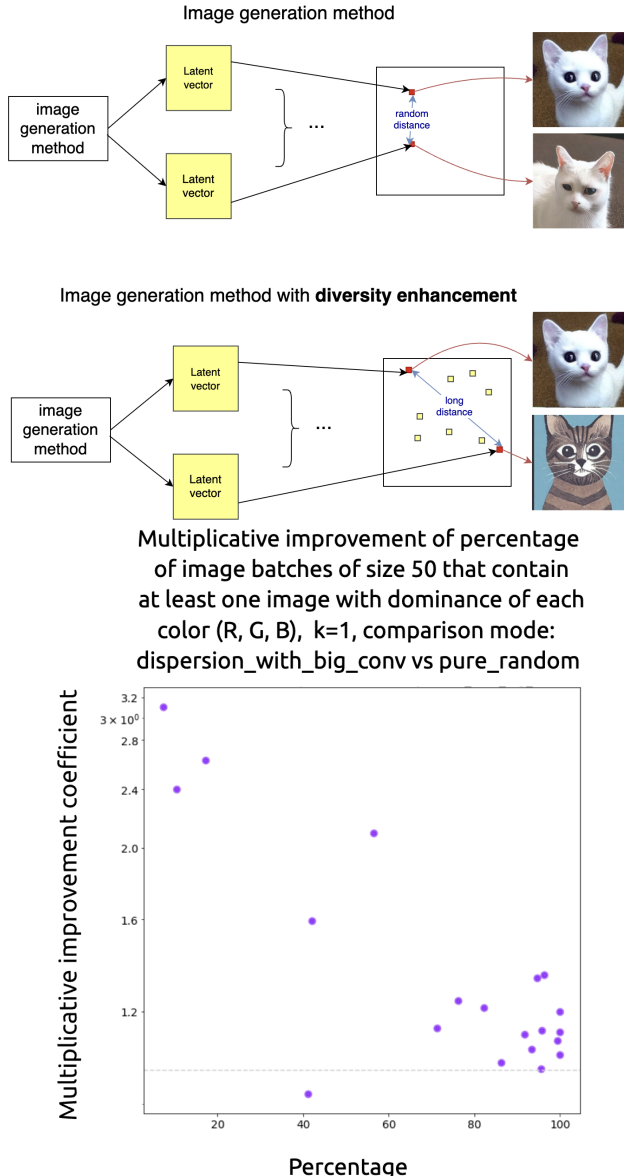


Figure 1. Top: General schema for diversity enhancement. A vanilla image generation method is compared to a method with enhanced diversity. Instead of being randomly generated, latent vectors are chosen on the basis of their pairwise distance, which results in better image diversity. Bottom: Comparison of dispersion-with-big-conv and standard LDM1 (pure-random) for various prompts: the presented numbers are the multiplicative improvement of the percentage of batches containing images with at least 2 different dominant colors for the following parameters: $K = 1$, batch size = 50. In general, we get better improvement for least represented groups.

that the batch of generated images is more diverse than in the vanilla case. In the present paper, we analyze the max-pooling method [20], compared to a vast sampling of tools for high-dimensional well distributed point configurations.

After careful examination, we show that dispersion combined with convolution (a) approximates the performance of max-pooling in the low-dimensional case and (b) performs better in large latent spaces used in modern latent diffusion models.

We note that the method described in [4] includes another strategy to ensure diversity. It considers a different context, and the proposed strategy is entangled in the diffusion process [19]. Due to this entanglement, [4] cannot be applied to other methods, such as LDM2 [13].

The rest of this paper is organized as follows. Section “Figures of merit for point configurations: artificial metrics in the LS” introduces the metrics for assessing the quality of point configurations. We explore the metrics designed to evaluate the distribution of images in Section “Figures of merit for the diversity of image”. Section “Generators of well distributed point configurations” describes the tools for generating well-distributed point configurations, while Section “Experimental results” details the experimental results derived from these tools. “Integration in public codes and reproducibility” discusses how our methodologies have been integrated into existing codebases. Finally, a mathematical analysis of our findings is provided in Section “Mathematical analysis”.¹

2. Measuring diversity

In this section, we present tools for measuring the point configurations diversity in the LS. These tools are instrumental in the development of our sampling methods. Likewise, we use them as a proxy to measure image diversity.

2.1. Figures of merit for point configurations: artificial metrics in the LS

Let μ be the Lebesgue measure on \mathbb{S}^d , the unit sphere in \mathbb{R}^d . Given a set S of distinct points in \mathbb{S}^d , $S = \{s_1, \dots, s_n\}$, defines the empirical measure associated to $S \subset \mathbb{S}^d$ as $\hat{\mu}_S(C) = \frac{\#C \cap S}{\#S}$. The set of spherical caps for a threshold r is $C_r = \{\{x \in \mathbb{S}^d; x \cdot u \geq r\}; u \in \mathbb{S}^d\}$. All the spherical caps for a same threshold r have the same measure. C_r is nontrivial if $r \in [-1, 1]$, and usual values lie in $r \in [0, 1)$. This provides **spherical cap diversity** measures for a point set S , namely mappings of the form $u \mapsto \#\{x \in \mathbb{S}^d; x \cdot u \geq r\}$ which carries the Lebesgue measure to C_r . **Covering: average and worst case.** We use an average covering by default, i.e., $\text{covering}(S) = \mathbb{E}_{x \in \mathbb{S}^d} \inf_{s \in S} \|x - s\|$. We also consider a worst case covering, namely $\sup_{x \in \mathbb{S}^d} \inf_{s \in S} \|x - s\|$. **Packing/Dispersion.** The packing, $\text{packing}(S) = -\inf_{1 \leq i < j \leq \#S} \|s_i - s_j\|$, is

¹All experiments were conducted on Jean Zay servers. Meta affiliated authors provided expertise and code and acted in an advisory role. We refrain from experimenting with pre-existing third-party generative models, such as Stable Diffusion, and instead use a large diffusion model (2.2B parameters) trained on an internal dataset of image-text pairs.

also based on a distance. The minus is used so that all our metrics are to be minimized. We also consider an average packing, namely $avg\text{-packing}(S) = \mathbb{E}_{x,y \in S, x \neq y} x \cdot y$. **Riesz.** The Riesz potential is also based on a distance, with $Riesz(S) = -\sum_{1 \leq i < j \leq \#S} \|s_i - s_j\|^{-s}$. We consider $s = \frac{1}{2}, s = 1, s = 2$. **Extending metrics using convolutions.** Typically, the LS is a space of tensors of a given shape. For example, some latent diffusion models use $64 \times 64 \times 4$, where 64×64 corresponds to spatial coordinates and 4 to the number of unordered channels. Different works [18, 20] have shown the importance of the spatial proximity of the coordinates. Likewise, the authors in [20] propose to group pixels in the LS by groups of 8×8 . In this work, we propose to use convolutions. Thus, instead of measuring a distance between x and y , we consider the distance between $conv(x, k)$ and $conv(y, k)$, where k corresponds to a kernel for a Gaussian blurring of some radius (8 by default) on the spatial coordinates. In our setting, BigConv and MiniConv correspond to a radius of 24 and 2. **Combining metrics.** We also consider the average between all the above metrics without considering any weight. More metrics are presented in the appendix in the section ‘‘Figures of merit for point configurations: additional artificial metrics’’.

2.2. Figures of merit for the diversity of images

We now consider diversity measures which are convenient for batches of images. Given a batch B of images and a set S of classes, the diversity $Div_{S,B}$ (to be maximized) typically comprises the number of elements in S which contain at least one element of B . This diversity strategy can be used for different criteria, discussed below. For a given generator of batches, B is a random variable, and by abuse of notation, $Div_{S,B}$ is used for the expectation $\mathbb{E}Div_{S,B}$. We can consider a binary criterion with value 1 if the batch reaches some predefined threshold and 0 otherwise. For example, we can consider the probability that we have at least 2 or 3 classes (up to all classes) present in the batch, i.e., $Div_{S,B}^c = P(Div_{S,B} > c)$.

Color diversity. To assess the color diversity in an image batch B , we employ a method from [20] that involves extracting color information from each image in the batch using the RGB color model, which represents colors as a combination of red, green, and blue channels. We compute the mean value for each channel ($red(I)$, $blue(I)$, $green(I)$) in a given image I , and identify whether one of these colors is predominantly present in the image: red (resp. green, blue) is said to be *dominant for a parameter K* if $red(I) \geq K \cdot green(I)$ and $red(I) \geq K \cdot blue(I)$. $K > 1$ corresponds to rarer classes, which can be relevant. We can consider color diversity as follows:

(a) $Div_{\{C_r, C_g, C_b\}, B}$, where $C_r = \{I; red(I) > green(I) \wedge red(I) > blue(I)\}$, and C_g and C_b the same with *green* or *blue* in the place of *red*.

(b) The case above leads to $Div_{S,B} \leq 3$: we can have max 6 classes (instead of 3) by considering both which color is maximum and which color is minimum. In these examples, except for equality cases, a class represents a partition.

(c) We can increase further the number of classes by considering classes of the form $C_{red,green,K}$, denoting cases in which $red(I) \geq \max(green(I), blue(I))$ and $red(I) > K \cdot \min(green(I), blue(I))$, and $green(I) \leq blue(I)$: there are 18 such classes when considering all colors and $K \in \{1., 1.1, 1.2\}$. For short, we denote this by $Div_{18}(B)$.

(d) We also consider a discretization of the vectors ($red(I)$, $green(I)$, $blue(I)$): each discretized value leads to a class/partition. For example, we can consider the color diversity $Div_{depth=d}$ for various depth values d . Then, the classes are chosen as follows. The image is reduced to a single pixel, encoded in three channels with 2^d possible values. Thus, the maximum number of classes is $2^{3 \times d}$. $Div_{depth=d}(B)$ is the number of classes observed at least once in B .

3. Generators of well distributed point configurations

The state-of-the-art has predominantly focused on datasets with low-dimensional characteristics, often yielding bounds that are non trivial only for large sample sizes. In the context of this paper, methodologies tools such as Halton [7] and Sobol [17] are not applicable due to the high-dimensional nature of the LS under study. We highlight four main methods in this paper: **covering**, **dispersion-with-big-conv**, **max** and **max-pooling**. Table 10 in the section ‘‘Additional information on methods’’ describes the other strategies.

(a) The **covering** generator maximizes the covering criterion defined above.

(b) The **dispersion-with-big-conv** method minimizes the worst case dispersion, modified by convolution with a kernel size (24, 24).

(c) In the **max** and **max-pooling method** [20], the ‘‘**max**’’ setting comprises a maximum number of iterations in randomly searching for a new vector that would have a maximal minimal distance to all the already selected vectors in the batch. The parameters of this method are the target size of the batch B and a maximum number of iterations N_{max} . In the ‘‘**max-pooling**’’ version, the distance is computed after processing the vectors by the average pooling on blocks of size 8×8 on the spatial coordinates, which down-samples the vector size to $8 \times 8 \times 4$ in the case of a LS shape of $64 \times 64 \times 4$.

4. Experimental results

We emphasize in red our recommended methods: the **covering** method in the non-spatial case (i.e., only channels as non-trivial dimensions) and the **dispersion-with-big-conv** one in the spatial case, in green the **max-pooling** method,

Model name	Model type	Input type	Output type	Latent Space
LDM2	text2image	Text (prompt)	Image	$512 \times 512 \times 4$
LDM1	text2image	Text (prompt)	Image	$64 \times 64 \times 4$
SpecialGAN	class to image	None	Image	$1 \times 1 \times 128$

Table 1. We consider spherical domains in the LS.

and in gray a **pure random strategy**.

In many cases, we present a multiplicative improvement as a metric for the impact of our method. The multiplication factor for dispersion-with-big-conv is the proportion of batches created by dispersion-with-big-conv which provides satisfactory results, divided by the same proportion with the vanilla (default, classical, i.e. pure random sampling) method.

The supplemental material presents additional results in terms of artificial metrics, as presented in the section “Figures of merit for point configurations: artificial metrics in the LS”. We focus here on diversity as measured on the images. In real-life, the batch size varies from a few units to thousands (when people have converged to a prompt they like, and create a big batch overnight). Consequently, we perform experiments with various batch sizes. We consider generative models (Table 1), which convert a latent tensor and possibly an input into an image: $image = model(input, latent)$. Given a point set and a classifier m with values in a finite set, we consider various diversity metrics $Div_{m,S}$ of the set I_S of images $I_S = \{model(input, latent); latent \in S\}$, with $S = \{s_1, \dots, s_n\}$. $Div_{m,S}$ is typically equal to the cardinal of $\{m(i); i \in I_S\}$ (on average, for a stochastic S). We refer to the section “Figures of merit for the diversity of images” for all criteria.

We distinguish, in the following subsections two main categories of LSs: with and without spatial coordinates.

4.1. SpecialGAN: no spatial coordinate, 128 channels

We consider a LS with 128 channels and no spatial coordinate, with SpecialGAN [3]. We present the color diversity results in Figure 2 (more results are illustrated in Figure 8 in the appendix). Overall, the covering method appears to be a reasonable criterion for this context, but the success is moderate, and a greater gap is observed for latent diffusion models, as in the next section.

4.2. Large Diffusion Models: spatial coordinates 64x64, and 4 channels

A Large Diffusion Model (LDM1) is a text-to-image generation model that begins with a noisy LS representation and progressively refines it into a coherent image. The LS is $64 \times 64 \times 4$ in dimensions, which is derived from compressing larger, high-resolution images (e.g., $512 \times 512 \times 3$). This model employs a reverse diffusion process. It starts with a




Figure 2. SpecialGAN has 75 classes, studied separately in 75 columns. For each method (row) and each SpecialGAN class (col), we consider the average color diversity Div_{18} . The score for that method (e.g., 7.23 for pure-random) is then the average Div_{18} over the 75 classes. Methods are ordered by average score (the greater, the better; best at the bottom). Batch size 48: additional results and details in Figure 8 in the appendix. Reading guide: pure-random has 7.23 classes on average, evaluated on 175800 batches. Overall, as shown by the results in the appendix, our methods frequently outperform pure-random for SpecialGAN but the gap with pure-random is much lower than for latent diffusion models.

Mode	bird	butterfly	cat	horse	rose
max-pooling	100.00	100.00	94.83±1.23	98.74±0.71	95.10±1.17
max	100.00	100.00	91.91±1.64	97.44±1.10	94.88±1.31
covering	99.26±0.33	100.00	77.26±1.20	96.06±0.70	85.73±1.09
dispersion-with-big-conv	100.00	100.00	93.02±1.37	96.65±1.06	87.09±1.68
dispersion-with-conv	100.00	100.00	91.09±1.49	97.18±0.90	84.52±1.74
dispersion-with-mini-conv	99.70±0.30	100.00	86.71±1.70	97.40±0.94	86.23±1.56
dispersion	99.49±0.30	100.00	77.83±1.28	92.83±0.94	83.56±1.21
pure-random	99.67±0.33	100.00	74.32±1.89	90.49±1.52	84.62±1.73
Riesz _{...}	99.70±0.21	100.00	77.32±1.16	95.21±0.78	84.02±1.09

Table 2. LDM1 (LS 64x64). Comparison of different techniques for various prompts: the presented numbers are the percentage of batches containing images with at least 2 different dominant colors for the following parameters: $K = 1.1$, batch size = 50; the methods are sorted by the average percentage over the different prompts. 3s computational cost. More methods are compared in Table 12 in the Appendix. Both max-pooling and dispersion-with-big-conv perform significantly better than pure-random.

compressed noisy latent representation and then, gradually denoises it through a series of steps to form a detailed image. If a batch of images needs to be generated, the same process applies to each of them. We call this generation process pure-random. **Color diversity.** We aim to determine the dominant colors in each image in each batch of generated images. We define a dominant color based on the coefficient K , and calculate the diversity metrics. The results for $K = 1$ (easier context, less differences between methods) are presented in Figure 1 (bottom). Additional results are presented in Tab. 14 and in Figure 9 in the appendix. The results for $K=1.1$ are presented in Tab. 2 and in Figure 3 (top), and in Figure 10 in the appendix. Finally, the results for $K=1.2$ are presented in Table 3 and in Figure 3 (bottom), and in Figure 11 in the appendix.

4.3. Statistics in LDM3/ LDM2: LS 512x512x4, 3 seconds

Before integrating our work in main public codebases, we check that it also performs well in a different contexts, including low batch size (8) and a different image generator (Large-LDM). We limit the time to three seconds, which is negligible compared to the image generation process in




Figure 3. Multiplicative improvement of percentage of image batches of size 50 that contain images with dominance of different colors (at least 2 from 3): we compare dispersion-with-big-conv vs pure-random. X-axis = proportion of satisfactory batches (defined as: at least two of the three dominant colors for $K = 1$) in the pure random case. Y-axis: multiplication factor, i.e., proportion for dispersion-with-big-conv divided by the proportion for pure-random (i.e., > 1 means an improvement compared to pure-random). Top: $K = 1.1$. Bottom: $K = 1.2$. In general, we get better improvement for least represented groups (low percentage, i.e., left of each subplot). Each dot represents a prompt from the list in Section “List of prompts” M.

a MacBook Pro M1. We use the color diversity measure $Div_{depth=d}$ for various depth values d and compare in Table 4 (dispersion-with-big-conv) and Table 5 (max-pooling) the diversity of batches created by well distributed point configu-

Mode	bird	butterfly	cat	horse	rose
max-pooling	95.08±1.27	99.67±0.33	45.86±1.59	73.11±2.10	53.59±1.53
max	88.89±1.90	99.19±0.57	37.87±1.94	63.08±2.18	41.73±1.80
covering	76.93±1.25	76.30±1.22	9.86±0.97	21.66±1.21	36.80±1.11
dispersion-with-conv	77.53±1.82	75.43±1.91	15.51±1.74	31.03±1.79	40.65±1.66
dispersion-with-big-conv	77.34±1.94	77.41±1.87	16.61±1.79	22.30±1.97	41.72±1.66
dispersion-with-mini-conv	71.30±1.77	73.75±1.87	10.30±1.57	30.48±1.95	39.94±1.54
Riesz ₂	78.36±1.25	75.48±1.21	9.04±0.93	25.40±1.24	36.08±1.09
pure-random	71.80±1.85	75.25±1.89	6.42±1.33	24.59±1.86	37.82±1.71
dispersion	75.81±1.35	72.53±1.23	9.12±0.96	23.99±1.28	34.56±1.22

Table 3. LDM1 with LS 64x64. Comparison of different techniques for various prompts: the presented numbers are the percentage of batches containing images with 2 different dominant colors for the following parameters: $K = 1.2$, batch size = 50; the methods are sorted by the maximum average percentage across the different prompts. 3s computational cost. Both recommended methods perform well for this 64x64 LS (512x512 images), with better results for max-pooling. More methods are compared in Table 13 in the Appendix.

rations in the case of this big latent spaces. As a baseline, we also use the same images randomly grouped in batches of the same size. As demonstrated by the results, dispersion-with-big-conv is always beneficial. Sometimes, in section “Latent Diffusion Model: spatial coordinates 64x64, and 4 channels” with smaller latent spaces, max-pooling was better: however, Tab. 5 (also Tab. 9 in the appendix) shows that for the same parametrization (i.e., same pooling size 8×8) as in [20], or for a parametrization scaling linearly with the spatial coordinates (i.e., 64×64 -pooling as we switch from a 64×64 LS to a 512×512 LS), max-pooling fails in this high-dimensional context. We therefore recommend dispersion-with-big-conv, which works over all tested settings, including cases with no spatial coordinates (for which the convolution does not do anything); but then covering is usually better (see section “SpecialGAN: no spatial coordinate, 128 channels” in the appendix). We refer to the appendix for additional measures of diversity, such as LPIPS [21].

4.4. Image quality

We start the discussion with gradient-based methods (described in the section “Additional information on methods”). We observe a big quality loss for images generated from latent variables created by those approaches. This can be explained as follows. As noted in [20], points in the LS with large norm, though they are theoretically possible in the Gaussian generator of latent variables, are rare, hence the diffusion does not work when the norm is large. We need points which have (for example) norm roughly \sqrt{d} in dimension d . This is aligned with the results observed in [20], as we operate within a normalized space. However, a second property of almost all generated points is that local averages (in the LS with spatial coordinates) are always close to 0. We should not have more local uniformity than the traditional random points. We observe that the gradient naturally “pushes” neighboring points in the same direction, leading to uniform local areas that diffusion approaches cannot handle.

Prompt	Depth	dispersion-with-big-conv	Baseline
a beautiful woman	1	2.16 (± 0.05)	2.08 (± 0.02)
	2	2.16 (± 0.05)	2.10 (± 0.02)
	3	2.75 (± 0.08)	2.63 (± 0.02)
	4	4.9 (± 0.1)	4.74 (± 0.03)
a handsome man	1	3.23 (± 0.06)	3.09 (0.02)
	2	3.23 (± 0.06)	3.08 (± 0.02)
	3	3.83 (± 0.08)	3.61 (± 0.03)
	4	5.42 (± 0.1)	5.22 (± 0.03)
a man and a woman dancing together	1	3.31 (± 0.07)	3.22 (± 0.02)
	2	3.31 (± 0.07)	3.22 (± 0.02)
	3	4.95 (± 0.09)	4.93 (± 0.03)
	4	6.2 (± 0.1)	6.17 (± 0.03)
a mma fight between two pharaohs	1	2.24 (± 0.04)	2.20 (± 0.01)
	2	2.24 (± 0.04)	2.22 (± 0.01)
	3	3.55 (± 0.07)	3.45 (± 0.02)
	4	5.46 (± 0.1)	5.33 (± 0.03)
a unicolor image	1	2.80 (± 0.07)	2.70 (0.02)
	2	4.79 (± 0.09)	4.46 (± 0.03)
	3	7.05 (± 0.09)	6.57 (± 0.03)
	4	7.84 (± 0.04)	7.71 (± 0.01)
unicorn+dragon dancing in a magical garden under a rainbow	1	3.35 (± 0.08)	3.23 (± 0.02)
	2	3.35 (± 0.08)	3.22 (± 0.02)
	3	3.35 (± 0.08)	3.23 (± 0.02)
	4	3.8 (± 0.1)	3.64 (± 0.03)

Table 4. Experiments on Large LDM (LS 512x512). Diversity Div_{depth} for various depth levels, for dispersion-with-big-conv and pure-random. All experiments with 800 images. Even on this completely different setting, results are positive in all cases. Our approach was never detrimental.

This leads to weird images, as illustrated in Figure 4.




Figure 4. Example of image obtained by LDM1 with points obtained by gradient-based methods, even after standardizing the norm: images do not make any sense.

We refer to the appendix for additional results confirming the quality preservation, with Brisque [11] estimates.

4.5. Discussion

The case without spatial coordinates is dominated by covering, which is intuitively simple and satisfactory. With spatial coordinates (as in most latent diffusion models), overall, max-pooling from [20] performs very well in the setting of [20], i.e., with a relatively small LS 64x64x4. In the case of 512x512x4, it becomes comparable, or even weaker, than pure-random (Table 5 for the default parametrization; an-

Prompt	Depth	Max-Pooling	Baseline
a beautiful manga character killing animals	1	2.46 (± 0.08)	2.45 (± 0.02)
	2	2.46 (± 0.08)	2.47 (± 0.03)
	3	2.46 (± 0.08)	2.46 (± 0.02)
	4	3.7 (± 0.1)	3.65 (± 0.03)
a gothic witch laughing and playing with a bazooka in the middle of hell	2	1.08 (± 0.04)	1.08 (± 0.01)
	3	3.44 (± 0.09)	3.46 (± 0.02)
	4	3.1 (± 0.1)	3.09 (± 0.04)
	1	2.86 (± 0.09)	2.96 (± 0.03)
a man and a woman dancing together	2	2.86 (± 0.09)	2.97 (± 0.03)
	3	4.3 (± 0.2)	4.36 (± 0.04)
	4	5.3 (± 0.2)	5.47 (± 0.05)
	1	1.48 (± 0.07)	1.48 (± 0.02)
an incredible image	2	1.48 (± 0.07)	1.46 (± 0.02)
	3	2.9 (± 0.1)	2.88 (± 0.03)
	4	4.0 (± 0.2)	4.13 (± 0.05)

Table 5. Max-Pooling with 8x8 blocks vs random for 512x512 LS. We present the diversity for a discretization of image colors as detailed in Section “Statistics in LDM3/ LDM2: LS 512x512x4, 3 seconds” i.e. depth 4. 400 images per prompt. The success of max-pooling (compared to pure-random) is questionable in this high-dimensional LS context. Hence, our preference for the more robust dispersion-with-big-conv (see Tab. 4). We removed the depth values for which all images were in the same class.

Large-LDM	Large-LDM + dispersion-with-big-conv	No preference
27.2% %	33.8%	39.0%

Table 6. Human ratings: statistics of preferences between Large-LDM, Large-LDM with dispersion-with-big-conv. Counting no preference as 50%, this is $53.3 \pm 2.1\%$ preference for the quality of our images. While this is not a statistically significant improvement, the quality is at least preserved; furthermore, our other results show a significant improvement in diversity, in particular in difficult cases. The exact question is “Do you prefer the image on the right or on the left for prompt XXX ? (Left, Right, No opinion)”

other scaling of parametrization is presented in Table 9 in the appendix). Cap-pooling (proposed in [20]) sometimes has a huge computational time, making it irrelevant in our high-dimensional context.

5. Conclusions

A classical method for increasing the diversity is to automatically add text suffixes. However, as shown by many recent counter-examples recently [2], this has several drawbacks, such as unrealistic outputs. We applied well distributed point configurations for sampling LSs, for enhancing the diversity of generated images. Compared to [20], we consider a broader range of cases (cases without spatial coordinates and

bigger LSs), and propose new methods. For the applications without spatial coordinates, the covering method performs best overall; and, in contexts with spatial coordinates, two strong methods are max-pooling and dispersion-with-big-conv. We note that taking into account the topology of variables (by convolution or pooling) is essential when working on LSs with spatial coordinates, and carefully choosing the size of convolution matters. Max-pooling is doing something similar to convolution, though by blocks, and performs best for small LSs with spatial coordinates. For bigger LSs, our dispersion-with-big-conv performs best; and we did not find a case in which dispersion-with-big-conv is detrimental, so we recommend it until better metrics are found. The case of gradient-based methods shows that the selected metrics are not perfect: we need the LS to be roughly sampled as in the original methods, and black-box optimization methods do that, whereas gradient-based methods do not. A nice achievement would be to design a metric that would simultaneously ensure diversity and quality: for the moment, only black-box optimization of our metrics was good at preserving quality (as shown by the human ratings) while improving diversity. We note that strong computational budgets do not have a clear positive impact, and results with gradient-based methods (numerically excellent but leading to poor quality) even suggest that over-optimizing can be detrimental: we recommend a time budget always tiny compared to the image generation part, and even when using large batch sizes (> 500) we never use more than a few seconds. Our method modifies a very specific part of image generation, namely the sampling of the latent part: it can be applied to GANs or latent diffusion models, and is compatible with other methods for increasing diversity. It has no drawback in terms of image quality.

5.1. Future work

A recent paper [18] proposed to use user preferences for guiding local perturbations in the LS: combining such approaches with our tools is a natural further work. The present work is a step in very high dimensional well distributed point configurations: maybe the same methods could be applied to other unexplored areas of applied mathematics.

5.2. Limitations

Our work can not create images that can not be created by the original model, or concepts unavailable in the original training data. Our method has a big impact for rare classes in latent diffusion models, but a moderate (though beneficial) impact on GANs. For small latent spaces, Max-Pooling is frequently better than our proposal.

References

- [1] Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. In *ACM Conference on Fairness, Accountability, and Transparency*, pages 1493–1504, 2023. 1

- [2] Nick Bonyhady. Black George Washington, 2024. Accessed on March 1st, 2024. 6
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019. 4
- [4] Gabriele Corso, Yilun Xu, Valentin De Bortoli, Regina Barzilay, and Tommi S. Jaakkola. Particle guidance: non-I.I.D. diverse sampling with diffusion models. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [5] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1):53–65, 2018. 1
- [6] Kathleen C Fraser, Svetlana Kiritchenko, and Isar Nejadgholi. Diversity is not a one-way street: Pilot study on ethical interventions for racial bias in text-to-image systems. *ICCV*, 2023. 1
- [7] J.H. Halton. On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numerische Mathematik*, 2:84–90, 1960. 3
- [8] Cheuk Ting Ho. Stable diffusion: Why are diverse results so hard to come by? In *anaconda.com*, 2023. 1
- [9] Diederik P Kingma, Max Welling, et al. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4):307–392, 2019. 1
- [10] David Marwood, Shumeet Baluja, and Yair Alon. Diversity and diffusion: Observations on synthetic image distributions with stable diffusion. *arXiv:2311.00056*, 2023. 1
- [11] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing*, 21(12):4695–4708, 2012. 6
- [12] Harald Niederreiter. *Random Number Generation and quasi-Monte Carlo Methods*. Society for Industrial and Applied Mathematics, 1992. 1
- [13] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis, 2023. 2
- [14] 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. 1
- [15] Jeremy Rapin and Olivier Teytaud. Nevergrad - A gradient-free optimization platform. github.com/facebookresearch/nevergrad, 2018. 20
- [16] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 1

- [17] I. M. Sobol. On the systematic search in a hypercube. *SIAM Journal on Numerical Analysis*, 16(5):790–793, 1979. [3](#)
- [18] Mathurin Videau, Nikolai Knizev, Alessandro Leite, Marc Schoenauer, and Olivier Teytaud. Interactive latent diffusion model. In *Genetic and Evolutionary Computation Conference*, pages 586–596, 2023. [3](#), [7](#)
- [19] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Survey*, 56(4), 2023. [2](#)
- [20] Mariia Zameshina, Olivier Teytaud, and Laurent Najman. Diverse diffusion: Enhancing image diversity in text-to-image generation. *arXiv:2310.12583*, 2023. [2](#), [3](#), [5](#), [6](#)
- [21] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [5](#)

A. Figures of merit for point configurations: additional artificial metrics

Cap discrepancy. The spherical cap discrepancy or cap discrepancy, for a threshold r , is defined in the present paper as the variance $Disc_r = Var_{c \in C_r} \hat{\mu}_S(c)$. By default, we use $r = r_d = \frac{1}{\sqrt{d}}$, so that the cap discrepancy is for us $Disc_{r_d}$.

This cap discrepancy has the advantage that it is easy to compute. Other cap discrepancies might consider $E(\hat{\mu}_S(C) - \mu(C))^2$, which is equivalent.

Half-sphere discrepancy. The half sphere discrepancy is the special case $Disc_0$ of cap discrepancy.

B. Integration in public codes and reproducibility

We provide our tools as a module in Nevergrad (starting from version 0.15.0). Given a batch of latent variables, with a shape $(n, s_1, s_2, \dots, s_k, c)$ the code will assume that we have n latent variables with shape $(s_1, s_2, \dots, s_k, c)$ with s_i spatial coordinates and c the number of channels. A convolution will be applied on s_1, \dots, s_k (at least if $k \geq 1$) where the last coordinate will be considered as channels. The code is as follows:

```
import nevergrad as ng
my_latent = ng.common.quasi_randomize(my_latent)
```

C. Additional experimental results: artificial metrics in the LS






Figure 5. Comparison between generators of well distributed point configurations (rows) for artificial metrics (columns), in the case 64x64x4 corresponding to LDM1. Numbers = proportion of methods which did better than this method for this criterion, i.e. 50% of methods did better than dispersion-with-big-conv for the metric packing-with-conv. Batch size 48. Scores = average per row, used for ranking, so that low values (at the top) are better.

Figures 5 to 7 present artificial metrics. Each column corresponds to a different figure of merit for well distributed point sets. Each row corresponds to a method. Each number $q_{m,c}$ in the heatmap corresponds to the proportion of methods doing better than method m for that metric c : the lower the better. Methods (rows) are ranked by averages of $q_{m,c}$ and this average is mentioned next to their name. Figure 5 presents results for an image LS corresponding to LDM1, namely 64x64x4. Our observations are straightforward. Good designs in terms of spherical cap discrepancy with convolution, or worst-case packing with convolution, are difficult to obtain without optimizing that specific criterion. Spherical cap discrepancy is very low (i.e. good) when we optimize the covering, and more unexpectedly, designs created by LHS also perform well for this criterion. 2-antithetic (simple centered symmetry) is excellent for half-sphere discrepancy (including with convolution), average packing (also including with convolution). Covering or dispersion are good for worst case packing, and more unexpectedly latin hypercube sampling is also good for this criterion. Figures 6 and 7 present results for channels only, i.e., shapes of the form $1 \times 1 \times c$. Also, we note that no artificial metric is really good at predicting if a point configuration is good for the image generation that follows. Packing with average looks not too bad, but gradient method are a counterexample: it turns out that the way we optimize a criterion has an impact on the usability of generated points by latent diffusion models. This shows that criteria are not perfect (see Section “Image quality”).

As expected, antithetic sampling performs well for half-spheres discrepancy and for average packing, and optimizing the sum of all metrics is effective for being reasonably good on all metrics. Most methods are worse than random for at least one



metric: this shows the complexity of high-dimensional point configurations.



points, shape (1,1,3)

48 points, shape (16,16,2)


384




6 points, shape (32, 32, 4)

96 points, shape (32,32,4)


Figure 6. Additional comparisons between generators of well distributed point configurations (rows) for artificial metrics (columns): similar to Figure 5 but with different batch sizes and shapes. Low values (at the top) are better.




12 points, shape (32,32,2)



12 points, shape (32,32,3)



50 points, shape (64,64,3)



12 points, shape (64,64,4)

Figure 7. Additional comparison between generators of well distributed point configurations (rows) for artificial metrics (columns). Similar to Figure 5 and Figure 6 but with different batch sizes and shapes. Low values (at the top) are better.

D. Additional results on SpecialGAN

Figure 8 extends Figure 2 (diversity Div_{18} for SpecialGAN) by considering several batch sizes.




Figure 8. Extension of results with SpecialGAN in Figure 2, for batch size 24, 48, 96, 192, and 384. The 75 SpecialGAN classes are studied separately in 75 columns. For each method (row) and each SpecialGAN class (col), we consider the average color diversity Div_{18} . The score for that method is then the average Div_{18} over the 75 classes. Methods are ordered by average score (the greater, the better; best at the bottom). Results are good for the covering method on average (or almost equal performance), but the difference is small, compared to cases with spatial latent classes for which there is a big difference. 22200 refers to the time budget in centiseconds.

E. Results with 3s, 6 cores, Latent Diffusion Model with LS 64x64x4

Results are also positive with a limited computational power (as in Tables 7 and 8).

Table 7. LDM1 with LS 64x64x4. Percentage of image batches of size 50 that contain at least one image with dominance of each color (Red, Green, Blue), $k = 1.1$. Number of cores: 6. Average latent vector generation time for different methods: 0s for pure-random, 2s for dispersion-with-big-conv, 3s for max-pooling, 3s for max. Even with very limited computational power, dispersion-with-big-conv improves the color diversity w.r.t pure-random.

Mode	butterfly	cat
ng_	75.69±2.41	12.70±2.11
dispersion-with-big-conv	80.28±3.79	8.05±2.68
dispersion	77.56±2.26	10.43±1.89
covering	73.45±1.96	12.50±1.73
rs_	75.34±0.92	9.36±0.76
greedy-dispersion-with-big-conv	67.37±3.24	16.98±3.03
dispersion-with-conv	70.75±3.13	13.19±3.08
greedy-dispersion-with-mini-conv	73.00±3.24	10.23±2.91
antithetic_pm	74.19±3.37	6.38±2.36
greedy-dispersion	73.96±3.31	6.25±2.54
dispersion-with-mini-conv	69.15±3.30	11.00±2.78
pure-random	74.12±3.53	6.00±2.23
Riesz_	6.60±0.42	2.19±0.25

Mode	butterfly	cat
dispersion-with-big-conv	44.78±11.47	0.00±1.00
ng_	39.18±7.76	1.59±3.06
antithetic-pm	36.96±10.85	1.11±3.74
greedy-dispersion-with-big-conv	37.89±10.57	0.00±1.00
rs_	36.99±2.99	0.84±0.03
covering	35.63±6.53	1.43±2.39
greedy-dispersion-with-mini-conv	36.84±10.69	0.00±1.00
dispersion-with-mini-conv	32.58±11.45	1.00±1.00
dispersion	32.83±7.67	0.49±0.06
greedy-dispersion	31.18±11.31	1.25±4.20
pure-random	29.63±12.21	0.00±1.00
dispersion-with-conv	27.45±10.96	0.00±1.00
Riesz_	3.04±1.04	0.03±0.01

Table 8. LDM1 with LS 64x64x4. Percentage of image batches of size 50 that contain at least one image with dominance of each color (Red, Green, Blue), $k = 1.2$. Number of cores: 6. Average latent vector generation time for different methods: 0s for pure-random, 2s for dispersion-with-big-conv, 3s for max-pooling, 3s for max. Dispersion-with-big-conv outperforms pure-random for the color diversity.

F. Additional experiments with dispersion-with-conv

Our main focus is on the more successful dispersion that one obtain with Big Conv. Thus, we check if variants (such as dispersion-with-conv) are also robust enough for being beneficial or at least non-detrimental. We observe in Figures 9, 10 and 11 that dispersion-with-conv becomes better and better for more difficult cases, compared to pure random. Additionally, when the classes are rare, the frequency has a multiplicative effect.

Diversity improvement for different time settings. Here we detail an experiment that aims at exploring the impact of time on color diversity in the generated images. The results are presented in Figures 12 and 14 for the “dispersion-with-conv” and “dispersion-with-big-conv” methods.

Multiplicative improvement of percentage of image batches of size 50 that contain at least one image with dominance of each color (Red, Green, Blue), comparison mode: dispersion_with_conv vs pure_random




Figure 9. LDM1 with LS 64x64x4. Comparison of dispersion-with-conv and standard LDM1 (pure-random) for various prompts: the presented numbers are the multiplicative percentage of batches containing images with all 3 dominant colors for the following parameters: $K = 1$, batch size = 50. In general we get better improvements (greater y-axis) for the most difficult cases (low x-axis value).

Multiplicative improvement of percentage of image batches of size 50 that contain images with dominance of different colors (at least 2 from 3) , $k = 1.1$, comparison mode: dispersion_with_conv vs pure_random




Figure 10. LDM1 with LS 64x64x4. Comparison of dispersion-with-conv and standard LDM1 (pure-random) for various prompts: the presented numbers are the multiplicative percentage of batches containing images with at least 2 of the 3 different dominant colors for the following parameters: $K = 1.1$, batch size = 50. In general we get better improvements (greater y-axis values) for least represented groups (lower x-axis values).

Multiplicative improvement of percentage of image batches of size 50 that contain images with dominance of different colors (at least 2 from 3) , $k = 1.2$, comparison mode: dispersion_with_conv vs pure_random




Figure 11. LDM1 with LS 64x64x4. Comparison of dispersion-with-conv and standard LDM1 (pure-random) for various prompts: the curves present the multiplicative percentage of batches containing images with at least 2 different dominant colors for the following parameters: $K = 1.2$, batch size = 50. In general we get better improvements (greater y-axis values) for least represented groups (lower x-axis values).





Figure 12. LDM1 with LS 64x64x4. This figure shows the variation in color diversity depending on the computational cost for the 'dispersion-with-conv' method. It compares the improvement in color diversity against the 'pure-random' baseline, demonstrating a slight increase in the percentage of image batches containing all dominant colors as the computational budgets progresses: we do not need a big computational power. In Section "Statistics in LDM3/ LDM2: LS 512x512x4", 3 seconds, we also show that 3s is enough for an impact on a larger LS 512x512x4 as the one of Large-LDM.

G. Additional experiments with max-pooling

The max-pooling method, successful in the 64x64x4 LS, is tested in the present 512x512x4 case. We test both 8x8 pooling (as in the 64x64 LS), and 64x64 pooling (i.e., scaling the size of blocks proportionally to the LS size).

We observe in Table 5 (with 8x8 blocks, so without scaling the block size when moving from 64x64 to 512x512) that results are not good for max-pooling when the LS is large.

We can also note in Table 9 that scaling the kernel size from 8x8 to 64x64 (when switching from a LS with spatial coordinates 64x64 to a LS with spatial coordinates 512x512) does not make max-pooling much better than random either (we tested a 8x8 kernel, i.e., no scaling proportionally to the LS spatial coordinates.)

Depth	max-pooling	pure-random
Prompt = a gothic witch laughing and playing with a bazooka in the middle of hell		
(num classes: 5, depth 3)		
3	3.300 (0.086897)	3.311667 (0.026412)
(num classes: 9, depth 4)		
4	3.020 (0.122857)	3.001667 (0.033236)
a unicolor image		
(num classes: 5, depth 1)		
1	2.700 (0.095831)	2.690 (0.029855)
(num classes: 11, depth 2)		
2	3.960 (0.133890)	3.933333 (0.038159)
(num classes: 26, depth 3)		
3	5.600 (0.145686)	5.648333 (0.041561)
(num classes: 83, depth 4)		
4	7.540 (0.081866)	7.311667 (0.030518)
Prompt=an incredible image		
(num classes: 3, depth 1)		
1	1.420 (0.070508)	1.480 (0.021215)
(num classes: 3, depth 2)		
2	1.420 (0.070508)	1.476667 (0.021470)
(num classes: 6, depth 3)		
3	2.720 (0.127903)	2.846667 (0.034424)
(num classes: 15, depth 4)		
4	4.120 (0.172946)	4.108333 (0.045497)

Table 9. max-pooling with 64x64 blocks vs random for 512x512 LS. We present the diversity for a discretization of image colors by depth as detailed in Section “Figures of merit for the diversity of images” and num classes refers to the number of classes that are represented at least once in our tests (num classes at most $2^{3 \times d}$ for a depth d , by definition of discretization at d bits with 3 colors).

H. Additional information on methods

Table 10 presents the methods used in the present paper, in particular those not developed in the main text.

Generator	Method
Pure random	Randomly generate (normal sampling) n points.
Antithetic methods	
2-Antithetic	Randomly draw $n/2$ points $x_1, \dots, x_{n/2}$, add their opposite $-x_1, \dots, -x_{n/2}$
24-Antithetic (a.k.a. big block symmetries)	Define s the last tensor dimension (4 for 64x64x4), define $k = s!$, randomly draw n/k points, and for each of them consider the k permutations of the last tensor indices: this corresponds to channels.
65536-Antithetic (a.k.a block symmetries)	Split the first and second tensor dimensions in 4, get a partition of the tensor scalars into 16 blocks. This leads to 2^{16} symmetries by replacing any of these blocks by its opposite. Now, when we generate a point, we also consider these 2^{16} symmetries.
16-Antithetic	Split the first and second tensor dimensions in 2, get a partition of the tensor scalars into 4 blocks. This leads to 2^4 symmetries by replacing any of these blocks by its opposite. When we generate a point, we also consider these 2^4 symmetries.
Metric-based methods	
Without convolution: use the Euclidean norm. With convolution: use the Euclidean norm after convolution over the 2 first coordinates	
Greedy dispersion	Generate the first point at random, then each point maximizes its minimum distance to previous points.
Dispersion (packing)	Same initialization as greedy dispersion, and then optimize the dispersion globally: we maximize $\min_{1 \leq i \neq j \leq n} \ s_i - s_j\ $.
Covering	Same initialization as greedy dispersion, and then randomly move points as in K-means algorithms for optimal covering (min average squared distance to the domain): we minimize $\mathbb{E}_s \min_{1 \leq i \leq n} \ s - s_i\ ^2$.
Optimizing metrics by black-box optimization (Section “Figures of merit for point configurations: artificial metrics in the LS”)	
Without convolution: use the Euclidean norm. With convolution: use the Euclidean norm after convolution over the 2 first coordinates	
Random search	Exists for all metrics, and RS-ALL optimizes on average over all metrics
Nevergrad	Same, but with optimization by Nevergrad
Other methods	
LHS	For each variable, randomly draw n points for each coordinate, one in each of the n quantiles (randomly ordered) of the standard Gaussian distribution. Then project radially to \mathbb{S}^d .
Jittered	Partition the sphere using the signature of a point: the signature is the ranking (among $24=4!$ possible values) of the 4 sums of the channels over each of the $4 \times 4=16$ squares partitioning the 2 first spatial coordinates. Then, draw points in (randomly ordered) parts of the partition, one at a time; repeat if not enough.
Reduced-Jittered	Same with $2 \times 2=4$ squares partitioning the 2 first spatial coordinates.

Table 10. The point generators that we consider. All metrics are considered after normalization of points to norm 1 and all outputs are rescaled for a norm as expected by the latent2image converter. For illustration purpose, the numbers proposed as prefix for the Antithetic methods correspond to the degree of antithetic methods in the case of 64x64x4 tensors. Channel: a channel is a tensor of variables corresponding to each index of the last dimension (there are 4 channels in 64x64x4). Spatial coordinates: the coordinates except the last one which corresponds to channels. For RandomSearch and Nevergrad, detailed names of methods are provided in Table 11.

Most of our results are based on ad hoc methods for specific criteria or black-box optimization based on [15]. However, we also include methods using gradient-based optimization, for the Riesz potential methods. They have “Riesz.” as a prefix. Figure 4 shows a typical, weird result obtain by these gradient-methods. We explain why we get such results in Section “Image quality”.

Name	Method
RS-Pack	Random search for the Packing metric.
RS-Cap	Random search for the spherical Cap metric ($r = 1/\sqrt{d}$).
RS-Cc	Random search for the covering metric with Convolution
RS-Pac	Random search for the Packing, with Average and Convolution
RS-Mhc	Random search for the Metric with Half spherical caps ($r = 0$) with Convolution
RS-metric	Random search for the half spherical caps ($r = 0$)
RS-ALL	Random search for the sum of all metrics above
RS	Random search for the spherical cap metric $r = 0$
RS-Pc	Random search for the Packing, with Convolution
RS-Pa	Random search for the Packing, with Average
RS-RA	Random search for the Riesz potential with $s = 1$
RS-RA2	Random search for the Riesz potential with $s = 2$
RS-RA05	Random search for the Riesz potential with $s = 0.5$
RS-RAC	Random search for the Riesz potential with $s = 1$ with convolution
RS-RAC2	Random search for the Riesz potential with $s = 2$ with convolution
RS-RAC05	Random search for the Riesz potential with $s = 0.5$ with convolution

Table 11. Naming of our random search search for creating point configurations. Methods with NG as prefix instead of RS refer to Nevergrad counterparts. In many figures, `ng_` is an average of all methods starting with `ng_` and `rs_` is an average of all methods starting with `rs_`: the differences usually did not justify using more space. Methods with Riesz as a prefix are methods with gradient.

I. Bird's eye view

Figure 13 presents a bird's eye view of our approach.




Figure 13. Top: LDM1, the pure random case.

- (a) The process starts with a text prompt (e.g., "cat") and initial noise.
- (b) The noise is then processed through a series of latent vectors, each sized $64 \times 64 \times 4$.
- (c) Through the diffusion process, these latent vectors are transformed step by step into a coherent image of a cat with dimensions $512 \times 512 \times 3$ pixels.

Bottom: LDM1, Diversity Enhancement Mode with $x \times y \times z$ pooling.

- (a) Similar to the pure random process, it begins with the same text prompt and noise.
- (b) Instead of a single chain of transformations, multiple latent vectors undergo parallel diffusion processes.
- (c) An additional pooling step ($x \times y \times z$ pooling) is employed, where various features from the parallel latent vectors are combined or selected to enhance diversity.
- (d) A diverse set of latent vectors is chosen according to the required batch size.
- (e) This results in a diverse array of images, exemplified here by two different cat images, both with dimensions $512 \times 512 \times 3$ pixels, but with distinct visual characteristics.

J. Figure illustrating the impact of the computational cost

14 shows that a moderate computational budget does not prevent the method from being effective. We note that Section "Statistics in LDM3/ LDM2: LS $512 \times 512 \times 4$, 3 second" shows that three seconds are enough to obtain positive results even when using the large LS of Large LDM ($512 \times 512 \times 4$).




Figure 14. This figure illustrates the change in color diversity (LDM1 with LS 64x64x4) depending on the computational cost for the 'dispersion-with-big-conv' method. Similarly to the 'dispersion-with-conv' method, there is a noticeable improvement in the percentage of image batches with complete color representation compared to the 'pure-random' approach, increasing slightly with the computational budget.

K. Mathematical analysis

Given classes $C = \{C_1, \dots, C_m\}$, such that $\forall j, C_j \subset \mathbb{S}^d$, the average match $M(S, C)$ of a stochastic point configuration S of cardinal n is the expected cardinal $\mathbb{E}\#\{i; 1 \leq i \leq m \wedge C_i \cap S \neq \emptyset\}$. By definition, $0 \leq M(S, C) \leq m$. If C is a partition and $n > 0$, then $0 < M(S, C)$. So, we use $Div_{x \mapsto (\chi_{C_i})_{1 \leq i \leq m}} S = M(S, C)$. We call pure random sampling the random independent uniform sampling of \mathbb{S}^d .

For each sampling that we have defined, it is easy to create an example in which it performs vastly better than random, i.e., $M(S, C) \gg M(\text{pure-random}_n, C)$ (with n the cardinal of the stochastic point configuration S). We check whether, for some stochastic point configuration S of cardinal n , it is possible to have counter-examples with $M(S, C) < M(\text{pure-random}_n, C)$.

K.1. Properties

2-antithetic is perfect for half-sphere discrepancy, in the following sense: for all half-sphere, S puts half points in it. We can also understand why, for $n = 3$ and $d = 2$, the greedy dispersion method performs weakly: the 3 points are in the same half.

K.2. Counterexample for all antithetic systems

Define an antithetic sampling of n points as follows. The sampling depends on a homeomorphism π from \mathbb{S}^d to \mathbb{S}^d , such that there is $\pi_0 \subset \mathbb{S}^d$ of measure $1/k$ such that $\mathbb{S}^d = \bigcup_{0 \leq i < k} \{\pi^i(x); x \in \pi_0\}$. Assuming that k divides n , we sample n/k points $b_1, \dots, b_{n/k}$ uniformly in π_0 and the sampling is $\{\pi^i(b_j); i < k, j < n/k\}$. We note this sampling *antithetic* $_{n,\pi,\pi_0,k}$ (n points, k strata, π_0 first stratum, π homeomorphism covering all strata).

Then, as an example in which antithetic sampling performs very well, we consider classes $C = \{C_1, \dots, C_m\}$ with $m = k$ exactly matching the $\pi^i(\pi_0)$ (i.e., $C_i = \pi^i(\pi_0)$ with $C_0 := C_m$), then $n \geq k$ implies that all classes are present in a batch of size n . This is the best case.

Now, let us consider the counter-example. It will be the opposite of the previous case: classes C_1, \dots, C_m are equally distributed over the J_j , and with a special positioning related to π as follows.

Proposition 1 (Antithetic sampling has counter-examples). *Given $k > 1$, $n/k > 1$, n multiple of k , then for all antithetic sampling $a = \text{antithetic}_{n,\pi,\pi_0,k}$ of \mathbb{S}^d of cardinal n , there exists a partition $C = \{C_1, \dots, C_k\}$ of \mathbb{S}^d such that $M(a, C) < M(\text{pure-random}_n, C)$.*

Proof. Given $a = \text{antithetic}_{n,\pi,\pi_0,k}$ of cardinal $n > k > 1$ and n multiple of k , we build a partition $C = \{C_1, \dots, C_m\}$ of \mathbb{S}^d such that $M(a, C) < M(a, \text{pure-random}_n)$. If we consider classes with the same measure (i.e., $\forall i, j, \mu(C_i) = \mu(C_j)$), and each class C_j is the union $\bigcup_i \pi^i(C_j \cap \pi_0)$ for $i < k$. Therefore:

(a) The number of classes found in the batch is equal to the number of classes found in the intersection of the batch and of π_0 , i.e. $M(S, C) = M(S \cap \pi_0, C)$.

(b) This number is therefore equal to the number of classes found in a pure random batch of size n/k .

Therefore, $M(\text{antithetic}_{n,\pi,\pi_0,k}, C) = M(\text{pure-random}_{n/k}, C) < M(\text{pure-random}_n, C)$ if $k > 1$ and $n > k$. \square

K.3. Counterexamples for packing, average covering and worst-case covering

Consider $d = 2$ and $n \geq 2$. The maximum packing or covering (average or worst case) is equivalent to antithetic sampling and has the same counter-example as in Section ‘‘Counterexample for all antithetic systems’’.

K.4. No counterexample for jittered sampling

We consider Jittered sampling as follows:

(a) Consider a partition J_1, \dots, J_k of \mathbb{S}^d , with all J_j having the same measure.

(b) Assuming that k divides n , we randomly, uniformly and independently draw n/k points in each J_j : x_i is drawn in J_j if $k|(i - j)$.

We call this the jittered sampling of \mathbb{S}^d and denote it *jittered* $_{-J,k,n}$.

Proposition 2 (No counter-example for jittered sampling). *Consider a partition $C = \{C_1, \dots, C_m\}$ of \mathbb{S}^d . Consider a jittered sampling for a partition J of cardinal k with n multiple of k . Then, $M(\text{jittered}_{-J,k,n}, C) \geq M(\text{PR}_n, C)$.*

Proof. $M(\text{Jittered}, C)$ is, by definition, the sum over $i \in \{1, \dots, m\}$ of the $P(S \cap C_i \neq \emptyset)$, with S the jittered sampling:

$$M(\text{Jittered}, C) = \sum_{i \leq m} P(S \cap C_i = \emptyset).$$

With Jittered sampling, the probability q_i of missing class C_i is $q_i = 1 - P(S \cap C_i \neq \emptyset) = 1 - \pi_{j=1}^k (1 - P(x_j \in C_i))^{n/k}$. Consider a fixed i , with $p_{j,i} = P(x_j \in C_i)$, then we have Eq. 1 and Eq. 2:

$$\frac{\partial q_i}{\partial p_{j,i}} = -\frac{nq_i}{k(1 - p_{j,i})} \quad (1)$$

$$\sum_{j=1}^k P(x_j \in C_i) = k\mu(C) \quad (2)$$

We apply the Karush-Kuhn-Tucker conditions to Eqs. (1) and (2), leading to the existence of λ such that $\forall 1 \leq i \leq m, \forall 1 \leq j \leq k, (1 - p_{i,j}) = \lambda q_i$. This shows that the minimum of $\sum_i q_i$ is reached when, for each i , all the $p_{j,i}$ are equal, which means that $C_i \cap J_j$ has the same measure for all j : this is equivalent to the pure random case.

Therefore, $q_i \leq \mu(C_i)/\mu(\mathbb{S}^d)$, and $M(\text{Jittered}, C) = \sum(1 - q_i) \geq M(\text{PR}_n, C)$. \square

K.5. Counter-example for Latin Hypercube Sampling

Consider dimension 2 (the unit circle of dimension 1 in \mathbb{R}^2) and 4 points. Consider 2 classes: the class of a point (x, y) is the sign of $x \times y$. Consider, without loss of generality, that the values for the first coordinates are in increasing order, with the two first negative and the two last positive. Then, the probability of all points in the same class is $1/8$ for pure random. Therefore the expected number of classes is $\frac{1}{8} + 2 \times \frac{7}{8} = \frac{15}{8}$.

For Latin Hypercube Sampling, the following holds:

- (a) The probability of having the second point with the same class as the first one is $1/3$.
- (b) If the two first points have the same class, the probability of having the third point in the same class as the two first points is 100%.
- (c) If the three first points have the same class, the probability of having the fourth point in the same class as the three first points is 100%.

Therefore the average number of classes is $\frac{1}{3} + 2 \times \frac{2}{3} = \frac{5}{3}$.

$\frac{15}{8} > \frac{5}{3}$, hence the expected result.

L. Additional tables: results on latent diffusion models

Table 14 shows that even in an easy context ($K = 1$ implies that classes are less rare) our method remains beneficial.

Mode	bird	butterfly	cat	horse	rose
max-pooling	100.00	100.00	94.83±1.23	98.74±0.71	95.10±1.17
max	100.00	100.00	91.91±1.64	97.44±1.10	94.88±1.31
covering	99.26±0.33	100.00	77.26±1.20	96.06±0.70	85.73±1.09
dispersion-with-big-conv	100.00	100.00	93.02±1.37	96.65±1.06	87.09±1.68
dispersion-with-conv	100.00	100.00	91.09±1.49	97.18±0.90	84.52±1.74
dispersion-with-mini-conv	99.70±0.30	100.00	86.71±1.70	97.40±0.94	86.23±1.56
greedy-dispersion-with-mini-conv	100.00	100.00	83.14±1.67	95.88±1.12	85.00±1.65
greedy-dispersion-with-big-conv	99.73±0.27	100.00	79.18±1.88	94.77±1.25	87.65±1.58
ng_	99.63±0.22	100.00	76.91±1.06	94.76±0.74	84.18±1.02
antithetic.pm	100.00	100.00	79.14±1.85	93.49±1.36	89.75±1.52
lhs	100.00	100.00	79.70±1.95	93.31±1.38	83.88±1.69
greedy-dispersion	100.00	100.00	78.29±1.85	93.04±1.34	83.96±1.73
dispersion	99.49±0.30	100.00	77.83±1.28	92.83±0.94	83.56±1.21
jittered	100.00	100.00	76.73±1.43	93.97±0.94	86.07±1.26
rs_	99.62±0.12	100.00	75.80±0.62	92.97±0.47	83.63±0.58
pure-random	99.67±0.33	100.00	74.32±1.89	90.49±1.52	84.62±1.73
Riesz_	99.70±0.21	100.00	77.32±1.16	95.21±0.78	84.02±1.09
big-block-symmetry	97.59±0.88	100.00	65.23±1.79	81.90±1.78	70.33±1.75
block-symmetry	92.43±1.37	99.68±0.32	61.81±1.63	77.42±1.84	66.57±1.73

Table 12. Full version of the Table 2 in the main paper. LDM1 (LS 64x64). Comparison of different techniques for various prompts: presented numbers are the percentage of batches containing images with at least 2 different dominant colors for the following parameters: $K = 1.1$, batch size = 50; the methods are sorted by the average percentage over the different prompts. 3s computational cost. ng_ is an average of all methods starting with ng_ and rs_ is an average of all methods starting with rs_, as detailed in Table 11. Both max-pooling and dispersion-with-big-conv perform significantly better than pure-random.

Mode	bird	butterfly	cat	horse	rose
max-pooling	95.08±1.27	99.67±0.33	45.86±1.59	73.11±2.10	53.59±1.53
max	88.89±1.90	99.19±0.57	37.87±1.94	63.08±2.18	41.73±1.80
covering	76.93±1.25	76.30±1.22	9.86±0.97	21.66±1.21	36.80±1.11
dispersion-with-conv	77.53±1.82	75.43±1.91	15.51±1.74	31.03±1.79	40.65±1.66
dispersion-with-big-conv	77.34±1.94	77.41±1.87	16.61±1.79	22.30±1.97	41.72±1.66
dispersion-with-mini-conv	71.30±1.77	73.75±1.87	10.30±1.57	30.48±1.95	39.94±1.54
greedy-dispersion-with-mini-conv	76.79±1.77	74.70±1.79	10.57±1.47	20.96±1.89	39.41±1.60
greedy-dispersion-with-big-conv	79.57±1.66	70.76±1.74	13.65±1.73	24.04±1.92	37.65±1.66
Riesz_	78.36±1.25	75.48±1.21	9.04±0.93	25.40±1.24	36.08±1.09
greedy-dispersion	80.95±1.79	74.83±1.90	9.87±1.54	24.05±1.82	38.68±1.68
lhs	79.74±1.83	78.03±1.85	8.49±1.55	22.89±1.92	34.63±1.70
ng_	76.15±1.15	76.34±1.13	9.24±0.86	21.56±1.13	36.17±1.02
rs_	77.26±0.63	73.12±0.63	9.23±0.50	23.92±0.64	35.48±0.58
pure-random	71.80±1.85	75.25±1.89	6.42±1.33	24.59±1.86	37.82±1.71
jittered	77.25±1.42	75.00±1.43	7.69±1.08	25.53±1.37	33.39±1.32
dispersion	75.81±1.35	72.53±1.23	9.12±0.96	23.99±1.28	34.56±1.22
antithetic_pm	77.22±1.94	75.66±1.76	9.27±1.52	22.60±1.90	35.71±1.72
block-symmetry	52.37±1.47	52.06±1.47	6.41±1.24	19.35±1.81	24.01±1.78
big-block-symmetry	52.07±1.53	58.31±1.58	7.28±1.38	12.70±1.64	17.21±1.70

Table 13. Full version of the Tab. 3 in the main paper. LDM1 with LS 64x64. Comparison of different techniques for various prompts: presented numbers are the percentage of batches containing images with at least 2 different dominant colors for the following parameters: $K = 1.2$, batch size = 50; the methods are sorted by the average percentage over the different prompts. 3s computational cost. Both recommended methods perform well for this 64x64 LS (512x512 images), with better results for max-pooling.

Mode	bird	butterfly	cat	horse	rose
max-pooling	100.00	100.00	99.05±0.54	100.00	97.87±0.78
max	100.00	100.00	100.00	100.00	97.82±0.86
dispersion-with-big-conv	99.28±0.50	98.01±0.79	89.70±1.57	95.54±1.20	96.36±1.04
covering	99.55±0.26	97.25±0.61	86.99±1.08	94.23±0.82	95.20±0.74
dispersion-with-conv	99.37±0.44	97.23±0.94	83.50±1.78	92.79±1.35	94.52±1.22
greedy-dispersion-with-mini-conv	99.40±0.42	97.56±0.83	85.71±1.60	97.25±0.94	97.06±0.89
dispersion-with-mini-conv	97.58±0.82	95.68±1.12	88.70±1.62	94.05±1.36	93.11±1.24
lhs	99.67±0.32	98.36±0.72	89.67±1.66	94.72±1.26	97.01±0.90
greedy-dispersion-with-big-conv	100.00	93.57±1.24	85.32±1.76	94.43±1.28	97.29±0.87
ng_	99.38±0.28	97.49±0.53	86.47±0.97	95.13±0.71	94.47±0.72
jittered	99.24±0.38	96.71±0.76	84.42±1.35	96.63±0.73	93.39±0.98
dispersion	98.97±0.41	97.67±0.56	89.47±1.09	95.17±0.80	95.08±0.81
rs_	99.44±0.14	96.82±0.33	88.89±0.53	95.95±0.37	94.14±0.42
pure-random	99.34±0.46	95.59±1.14	86.15±1.73	95.74±1.11	93.27±1.32
antithetic_pm	100.00	96.48±0.96	89.74±1.57	95.55±1.15	95.34±1.12
greedy-dispersion	99.68±0.32	96.60±1.02	86.51±1.70	95.89±1.07	92.77±1.35
Riesz_	99.85±0.15	98.21±0.48	89.68±0.97	95.50±0.75	96.25±0.65
big-block-symmetry	93.10±1.39	88.52±1.55	72.85±1.87	87.30±1.64	73.89±1.77
block-symmetry	90.22±1.51	84.13±1.73	74.34±1.75	86.77±1.67	77.81±1.78

Table 14. LDM1 with LS 64x64x4. Comparison of different techniques for various prompts: the presented numbers are the percentage of batches containing images with at least 2 different dominant colors for the following parameters: $K = 1$, batch size = 50; the methods are sorted by the average percentage over the different prompts. 3s computational cost.

M. List of prompts

M.1. Prompts for Figure 4

“a beautiful manga character killing animals”, “a man and a woman dancing together”, “a mma fight between two pharaohs”, “an incredible image”, “a cyberpunk superhero smiling”, “a landscape from a planet in outer space”.

Prompt	Naive random	Our tool
Tool = Dispersion With Big Conv		
A cow	0.952823	0.937371
A falcon	0.904841	0.880119
A landscape from a planet in outer space	0.877	0.869
A panda	0.879	0.884
A panda (CFG 2)	0.877	0.911
A unicorn and a dragon dancing in a magical garden under a rainbow	0.889	0.877

Table 15. Brisque scores for different prompts, in the vanilla (pure random) case and for our tool. All batches have size 8. We also tested CFG scale 2.

N. Quality estimation by Brisque

Fig. 15 presents the Brisque scores of our method on top of Large-LDM, compared to the vanilla method.

O. Diversity estimation by Lpips

Fig. 16 shows that our method increases the Lpips diversity score. Note that we compare our batches of images to a shuffling of our images into random batches: we compare numbers for the same batch size and the same folder of images, so that we just validate the fact that the batches as created by our methods are more diversified than the batches obtained by randomly grouping our images into batches. Fig. 17 compares our method to vanilla random. We also check that it works when CFG scale is modified.

Prompt	Our tool	Random
Tool = Dispersion With Big Conv		
a beautiful manga character killing animals	0.619378*	0.617485
a beautiful woman	0.616295	0.614785
a cute pocket monster	0.677759	0.672921
a cyberpunk superhero smiling	0.642695	0.638868
a gothic witch laughing and playing with a bazooka in the middle of hell	0.624919*	0.619987
a handsome man	0.58878	0.587172
a landscape from a planet in outer space	0.553805	0.550507
a man and a woman dancing together	0.669772	0.668891
a mma fight between two pharaohs	0.603717	0.601804
a unicolor image	0.44821	0.42293
a unicorn and a dragon dancing in a magical garden under a rainbow	0.637551	0.636992
an incredible image	0.602923*	0.597309
a falcon	0.617723	0.58933
a cow	0.632388	0.613173
Tool = Max-Pooling		
a beautiful manga character killing animals	0.618118	0.617284
a gothic witch laughing and playing with a bazooka in the middle of hell	0.614835	0.615309
a man and a woman dancing together	0.670249*	0.670072
a mma fight between two pharaohs	0.605242*	0.600826
an incredible image	0.590427	0.588236

Table 16. Lpips score for batches of images created by our tool and when randomly shuffling batches. This shows that the diversity is increased. All batches have size 8. Our proposed tool is always improving diversity, and Max-Pooling is most often ok. For the 5 prompts also run with Max-Pooling, the stars denote the best of random/max-pooling/dispersion-with-big-conv (3 times Dispersion-with-big-conv, twice Max-Pooling, and never random).

CFG scale	Lpips vanilla	Lpips Disp. Big Conv.
1	0.62	0.64
2	0.66	0.68

Table 17. Average Lpips diversity for the prompt “A panda”, with CFG scale 1 or 2. We compare Vanilla Large-LDM and our version improved by Dispersion with Big Conv. 100 batches of 8 images in each case.