

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 SCORE-BASED GENERATIVE MODELING THROUGH ANISOTROPIC SPDEs

Anonymous authors

Paper under double-blind review

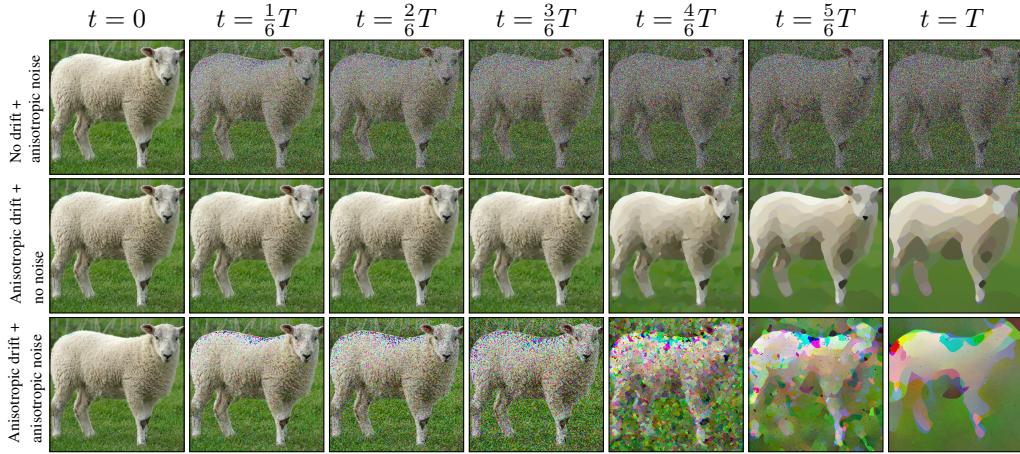


Figure 1: We visualize forward processes within our anisotropic diffusion framework. From top to bottom: (1) isotropic drift with anisotropic noise, (2) anisotropic drift **without noise**, and (3) both anisotropic drift and noise. Different anisotropy settings influence the preservation of geometric structures during the forward process, thereby affecting their reconstructability during generative sampling.

ABSTRACT

Score-based generative modeling (SBGM) has achieved state-of-the-art performance in image generation, with the quality of generated images highly dependent on the design of the forward (diffusion) process. Among these, models based on stochastic differential equations (SDEs) have proven particularly effective. While traditional methods aim to progressively destroy all image information to enable reconstruction from pure noise, we introduce a novel class of anisotropic stochastic partial differential equations (SPDEs) that preserve the geometric structure of the data throughout the transformation. These SPDEs consist of a drift term that enforces deterministic destruction via structured smoothing, and a diffusion coefficient that enables random destruction through noise injection. Both components are governed by anisotropy coefficients, enabling controlled, direction-dependent information degradation. This framework provides the theoretical foundation for a novel anisotropic SBGM. Due to geometry-aware degradation, the data generation process can exploit residual geometric cues, leading to improved fidelity in image reconstruction. We empirically validate this improvement in a proof-of-concept implementation on unconditional image generation, showing that anisotropic diffusion can achieve superior image quality metrics.

1 INTRODUCTION

Diffusion-based generative models (Song & Ermon, 2019; Ho et al., 2020) have gained significant traction due to their remarkable ability for high-quality image synthesis in both conditional and unconditional settings. Their impressive performance in modeling high-dimensional datasets of considerable size has led to their adoption across various other fields, such as video synthesis (Xing et al., 2024), speech and audio synthesis (Kong et al., 2020; Huang et al., 2022), medical imaging (Song et al., 2021a; Kazerouni et al., 2023), and molecular design (Weiss et al., 2023; Schneuing et al., 2024).

Broadly, diffusion-based generative models can be categorized into two classes. The first class comprises SBGMs, which are based on learning the gradient of the log-density (the *score*) of the data transformation.

054 These models are further divided into two subcategories: The first category, ordinary score matching (OSM)
 055 (Song & Ermon, 2019), directly estimates the score of the data distribution itself. The second category, Denoising
 056 Score Matching (DSM), learns instead the score of perturbed data, which includes methods such as Score
 057 Matching with Langevin Dynamics (SMLD) (Song & Ermon, 2019), Denoising Diffusion Probabilistic Models
 058 (DDPM) (Ho et al., 2020), and SDE-driven score-based diffusion models (SDE-driven SBGMs) (Song et al.,
 059 2021c). In this paradigm, a noise perturbation is applied to the data, and the model learns to approximate the
 060 score of this transformed distribution at each noise level. This learned score can then be used within various
 061 sampling algorithms to generate new data points.

062 The second class of diffusion-based generative models follows an alternative approach, as introduced by Sohl-
 063 Dickstein et al. (2015). Rather than explicitly estimating the score function, these models define a latent variable
 064 hierarchy and optimize a variational bound to maximize the likelihood of the data. While these methods share
 065 conceptual similarities with DDPM, their formulation does not rely on score estimation and instead focuses on
 066 parameterizing the forward and reverse processes explicitly.

067 In this work, we adopt the approach of the former class. More precisely, we build up on and extend the
 068 framework of SDE-driven SBGMs. These SDE-driven models provide a powerful mathematical framework
 069 with fine-grained control and flexibility when it comes to designing the generative process. The framework
 070 we introduce allows us to consider anisotropic diffusion processes that are aware of the geometrical structures
 071 present in the underlying data. Inspired by promising results from previous work (Yu et al., 2023; Vandersanden
 072 et al., 2024) on anisotropic diffusion processes, our hope is that our framework enables the design of better
 073 performing generative diffusion models.

074 Our main contributions are the following:

- 075 • **Anisotropic Diffusion Framework:** We introduce a novel anisotropic diffusion framework that re-
 076 respects geometric structures during data destruction, facilitating the resemblance of geometric features
 077 in the generative sampling process (Section 4).
- 078 • **Empirical Validation:** We demonstrate the potential of our theoretical framework through a numer-
 079 ical study of a proof-of-concept implementation on unconditional image generation (Section 5).

080 In addition, we provide a unifying formulation of SBGM that is a straightforward generalization of existing
 081 work (Lim et al., 2023; 2024) (Section 3), extending it to encompass our class of anisotropic diffusion processes
 082 as well. This serves primarily as a by-product: it allows us to describe training and generative sampling in a
 083 common framework that subsumes both existing approaches and our anisotropic diffusion framework, thereby
 084 avoiding a separate, framework-specific presentation.

085 2 RELATED WORK

086 **Conceptually related models** Song et al. (2021c) first used time SDEs for SBGMs and showed how
 087 existing diffusion models can be unified by an SDE framework. However, they only considered linear SDEs
 088 with spatially independent diffusion coefficients.

089 Rissanen et al. (2023) considered a stochastic heat equation with isotropic noise, which is effectively destroying
 090 the data by blurring up to complete dissipation. This is in contrast to earlier approaches that typically destroyed
 091 data into pure noise. Hoogeboom & Salimans (2022) extended this idea by introducing a temporally increasing
 092 isotropic noise term, further refining the blurring process over time.

093 **State-of-the-art models** Lipman et al. (2022) propose Flow Matching (FM), a simulation-free training
 094 method for continuous normalizing flows (CNFs) that regresses vector fields along predefined probability paths.
 095 FM enables training CNFs with more efficient paths such as optimal transport interpolations, yielding faster
 096 sampling and superior sample quality. We refer to Section B for a discussion on introducing anisotropy in FM.

097 Zhou et al. (2023) introduce Denoising Diffusion Bridge Models (DDBM), which generalize diffusion mod-
 098 els to map between arbitrary endpoint distributions using learned diffusion bridges. This framework unifies
 099 generative modeling paradigms and enables tasks like image translation, achieving strong performance while
 100 remaining competitive with state-of-the-art models in standard generation settings. Conditional generative
 101 modeling through anisotropic diffusion bridges is part of future work; we elaborate on that in Section C.

102 **Anisotropic models** Several studies (Voleti et al., 2022; Yu et al., 2023; Vandersanden et al., 2024) ex-
 103 plored the role of anisotropic noise in diffusion models. Vandersanden et al. (2024) propose a structure-aware
 104 anisotropic diffusion process that preserves edges longer, improving sample quality, particularly in shape-
 105 oriented generative tasks.

106 Our approach shares similarities with this work, as the anisotropic SPDE we introduce is also guided by struc-
 107 tural image content in both the drift and diffusion terms. However, our method differs in that it models a

108 genuinely anisotropic and nonlinear diffusion process, where both the drift and diffusion coefficients evolve
 109 dynamically based on the current state, rather than being fixed by the initial state.
 110

111 **SPDE-based models** Lim et al. (2023; 2024) also consider generative modeling using SPDEs. The
 112 parabolic SPDE studied in Lim et al. (2024) is restricted to spatially-independent diffusion coefficients. This
 113 limitation prevents the modeling of spatially varying or anisotropic effects in the forward process, and therefore
 114 represents a more constrained setting than our framework that allows for general anisotropic diffusion.

115 3 UNIFIED FORMULATION OF SBGM

118 In this section, we establish a unified formulation of SBGM that encompasses existing approaches on SBGM,
 119 including OSM, SMLD, DDPM and SDE-driven SBGMs, as well as our anisotropic diffusion framework,
 120 which we propose in Section 4. Our unifying formulation is a straightforward generalization of existing work
 (Lim et al., 2023; 2024), extending it to encompass our class of anisotropic diffusion processes as well.

121 This unifying perspective allows training and data generation to be described in a common, high-level fashion
 122 that applies uniformly across all existing SBGMs — including ours. In this way, our framework can be pre-
 123 sented within the same generic algorithmic structure as prior work, rather than through bespoke descriptions
 124 that might suggest a fundamental methodological departure.

126 **Overview** Generative modeling operates as a two-pass procedure: In the first (*forward*) pass, considered
 127 in Section 3.1, the information in the data is systematically destroyed to a certain extent. During this pass,
 128 the score (i.e. the log-density) of the forward process is learned (by a neural network) — as explained in
 129 Section 3.2.

130 In the *backward* pass, the destroyed data is stochastically reconstructed, resulting in new, previously unseen
 131 samples resembling the original data. Because of this dual perspective, the transformation $(U_t)_{t \in \bar{I}}$ is called the
 132 *forward process*, while its time reversal

$$\bar{U}_t := U_{T-t} \quad \text{for } t \in \bar{I}. \quad (1)$$

133 is referred to as the *backward process*. In Section 3.4, we explain how it can be used for generative sampling.

136 3.1 FORWARD PASS (DATA DESTRUCTION)

138 In generative modeling, the goal is to learn a data distribution μ and generate new samples that closely resemble
 139 the data. Conceptually, the data distribution μ is a probability measure on \mathbb{R}^D , where D is a finite (index) set.
 140 In practice, μ is unknown and only implicitly given by a dataset, where we assume that this dataset is an
 141 independent and identically μ -distributed sequence.

142 More specifically, in *score-based* generative modeling, the (complex) data is mapped by a (stochastic) transform
 143 to a simpler representation, and the *score* (i.e., the log-density) of this transformation is learned. Conceptually,
 144 the information in the data is progressively destroyed until fully degraded, after which the objective is to gen-
 145 erate new data resembling the original distribution. From an analytical perspective, the data is smoothed over
 146 time, progressively simplifying the learning task — a practice referred to as *regularization by noise*.

146 Historically, various types of transformations have been explored in SBGM. To describe all of them, including
 147 our approach proposed in Section 4, in a unified manner, let $\bar{I} \subseteq [0, \infty)$ denote the time index set of the trans-
 148 formation, and let $(U_t)_{t \in \bar{I}}$ represent the transformation itself. Formally, $(U_t)_{t \in \bar{I}}$ is an \mathbb{R}^D -valued (stochastic)
 149 process with initial distribution μ , i.e., it begins by sampling its initial state from the data distribution μ .

150 3.2 LEARNING THE SCORE FUNCTION

152 The score of the forward process provides the necessary information for reconstructing samples in the gener-
 153 ative sampling process. Without a score, the reconstruction would be purely deterministic, meaning that the
 154 generative sampling process would have to exactly invert the forward transformation. However, exact inver-
 155 sion is often impossible, as the forward process is designed to progressively degrade information in a way that
 156 cannot be deterministically undone.

157 By incorporating stochasticity into the forward dynamics, we ensure that the degradation is probabilistic, ren-
 158 dering the reverse process well-posed in a statistical sense. We approximate the score using a neural network
 159 trained during the forward pass (Section 3.1) and leverage it to guide generative sampling, the generative sam-
 160 pling process, as described in Section 3.4.

161 That is, for SBGM, we need to make sure that the score of the forward process actually exists. Consequentially,
 we assume that U_t has a positive differentiable density p_t with respect to the D -dimensional Lebesgue measure

162 for all $t \in I$. The *score* of the transformation at time $t \in I$ is now defined to be
 163

$$s(t, \cdot) := \nabla \ln p_t. \quad (2)$$

164 The goal of SBGM is now to train a (time-dependent) score-based model to find an approximation of s . To this
 165 end, we need a suitable metric measuring the distance between any given approximation \tilde{s} to the true score s .
 166 Existing approaches have consistently relied on an L^2 -norm for that purpose. We now derive a *loss measure*,
 167 unifying OSM, SMLD, DDPM and SDE-driven SBGMs, with respect to which this L^2 -norm is defined.

168 The time domain \bar{I} of the forward process can either be discrete or continuous. We assume that $0 \in \bar{I}$ and
 169

$$T := \sup \bar{I} \in [0, \infty) \setminus I \quad (3)$$

170 for a smaller subset $I \subset \bar{I}$, which is actually used during the learning process as we will describe now. Introducing a *loss measure*
 171

$$\eta(A \times B) := \int_A \mathcal{U}_I(dt) \zeta(t) \mathbb{P}[U_t \in B] \quad (4)$$

172 for measurable $(A, B) \subseteq I \times \mathbb{R}^d$, the typical *loss function* L is the $L^2(\eta)$ -distance to the actual score s ; i.e.
 173

$$L(\tilde{s}) := \|\tilde{s} - s\|_{L^2(\eta)}^2 = \int \mathcal{U}_I(dt) \zeta(t) \mathbb{E}[\|(\tilde{s} - s)(U_t)\|^2] \quad (5)$$

174 for measurable $\tilde{s} : I \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, where \mathcal{U}_I denotes the uniform distribution on I . In this definition, $\zeta : I \rightarrow$
 175 $[0, \infty)$ is a *weighting* function allowing us to put more importance on certain parts of the transformation.
 176

180 3.3 UNIFYING PRIOR WORK

181 As apparent from the definition, the loss is not considered over the entire time domain \bar{I} of the transformation,
 182 but only over the smaller subset I . With this description, we capture all of the previous works. They differ in
 183 how they handle the forward process and score estimation. With our unified formulation, we got

- 185 ○ $I = \{0\}$ and $\bar{I} = I \uplus \{1\} = \{0, 1\}$ in OSM;
- 186 ○ $I = \{1, \dots, k-1\}$ and $\bar{I} = \{0\} \uplus I \uplus \{k\} = \{0, \dots, k\}$
 for some $k \in \mathbb{N}$ with $k \geq 2$ in SMLD and DDPM; and
- 188 ○ $I = [t_0, T)$ for some $0 < t_0 < T < \infty$ and $\bar{I} = [t_0, T]$ in SDE-driven SBGMs.

190 That is, in OSM, the data from μ is not actually transformed, but remains unchanged. Instead, the score
 191 $s(0, \cdot)$ of the data distribution μ is tried to be learned directly. To generate samples, we initialize from a *prior*
 192 distribution, which corresponds to the distribution of U_T — the state of the forward process at the final time T .
 193 In this setting, the transformation by the forward process is effectively performed in a single step by drawing
 194 from the *prior*, without dependence on an initial dataset sample. Thus, the index set of the transformation is
 195 given by $I = \{0, 1\}$.

196 In contrast, SMLD, DDPM, and SDE-driven SBGMs do not aim to learn the score of the data distribution di-
 197 rectly. Instead, they transform dataset samples through a nontrivial forward process, producing representations
 198 at multiple (potentially infinitely many) *noise*, or — more generally — *diffusivity levels*, and not at the initial
 199 stage corresponding to the data distribution directly. These models then learn the score of the transformed
 200 distributions at each diffusivity level.

200 In SMLD and DDPM, the forward process operates at a discrete set of diffusivity levels, indexed by $\bar{I} =$
 201 $\{0, \dots, k\}$. In contrast, in SDE-driven SBGMs, the forward process is continuous, resulting in a continuous
 202 range of diffusivity levels with index set $\bar{I} = \{0\} \cup [t_0, \infty)$ for some $t_0 > 0$, where score learning begins at t_0 .
 203

204 **Prior sampling** The distribution of U_T is being referred to as the *prior* distribution of the backward, data
 205 generation pass. A sample from the prior distribution U_T is formally obtained by initializing the forward
 206 process $(U_t)_{t \in \bar{I}}$ with a random sample from the data distribution μ and simulating it up to the terminal time T .
 207 Depending on the complexity of the prior, this may be the only practical sampling procedure; for example, this
 208 is the approach taken in Rissanen et al. (2023).

209 However, in OSM, the prior distribution is completely decoupled from the initial data distribution μ from which
 210 U_0 is initialized. In fact, in traditional methods, the prior distribution is typically (approximately) Gaussian.
 211

212 3.4 BACKWARD PASS (DATA GENERATION)

213 Once we have trained a time-dependent score-based model to approximately minimize L , we can generate new,
 214 previously unseen data that resembles the training data distribution μ .

215 This generation process follows an iterative sampling scheme. It begins with a draw from the prior distribution.
 216 At each iteration, the scheme applies an optional *predictor* step, followed by an optional *corrector* step.

- 216 ○ The predictor step, if applied, propagates the sample backward in time by simulating a step of the
217 backward process. This step may incorporate information from previous samples, which is particu-
218 larly beneficial when the backward process has a Markov property.
- 219 ○ The corrector step, if applied, refines the sample using unadjusted Langevin algorithm (ULA), treat-
220 ing the sample from the previous iteration as the initial state and targeting the distribution at the
221 corresponding time step. This is feasible because ULA only requires the gradient of the log-density,
222 which can be estimated using the learned score function. However, using the *Metropolis-adjusted*
223 ULA would necessitate additional density estimation techniques.

224 The full procedure is summarized in Algorithm 3.1. Depending on the specific score-based generative frame-
225 work, only the predictor step (e.g., DDPM, SDE-driven SBGMs), only the corrector step (e.g., OSM, SMLD),
226 or both (as in SDE-driven SBGMs) may be used. However, at least one of these steps must be applied.

227 How sampling in the predictor step can be performed, completely depends on the complexity of the forward
228 process. If it has a Markov property or is given as the solution of an SDE, as it is the case for our forward
229 process introduced in Section 4, special sampling techniques are available. For details, we refer to Section D.

231 **Algorithm 3.1** Generative sampling process

232 **Output:** Sample u_0 from the data distribution μ
233 1: **if** I is continuous
234 2: Choose $k \in \mathbb{N}$ and (strictly) increasing $t_0, \dots, t_{k-1} \in I$;
235 3: **else**
236 4: Enumerate $I = \{t_0 < \dots < t_{k-1}\}$, where $k := |I|$;
237 5: Sample u_k from the prior distribution U_T ;
238 6: **for** $i = k, \dots, 1$
239 7: Sample u_{i-1} from $U_{t_{i-1}}$ given u_i, \dots, u_k ;
240 {Optional predictor step}
241 8: Correct u_{i-1} by applying ULA with initial state u_{i-1} and target density $p_{t_{i-1}}$;
242 {Optional corrector step}

243 With Algorithm 3.1, we capture all of the existing work. In OSM, Algorithm 3.1 employs only the corrector
244 step, using ULA iterations with the learned score and the prior sample as the starting point. In SMLD, Algo-
245 rithm 3.1 also applies only the corrector step, meaning ULA iterations are performed sequentially, descending
246 through diffusivity levels from an initial prior sample.

247 In contrast, in DDPM and SDE-driven SBGMs, the predictor step in Algorithm 3.1 is used, where samples are
248 generated by sequentially following the dynamics of the learned backward process, which is a Markov process
249 in those frameworks.

250 SDE-driven SBGMs can also incorporate a corrector step. Unlike DDPM, where exact sampling from the
251 backward process is feasible, SDE-driven SBGMs only allow approximate sampling due to continuous-time
252 dynamics. Consequently, ULA can be interleaved with predictor steps to correct bias introduced by approxi-
253 mation errors.

254 4 ANISOTROPIC DIFFUSION FRAMEWORK

257 In this section, we describe our anisotropic diffusion framework. We extend conventional SDE-based ap-
258 proaches by formulating the forward process as the solution of an SPDE. This formulation naturally incorpo-
259 rates spatial derivatives, enabling structured and anisotropic transformations of images.

260 Unlike conventional SBGMs that operate directly on the discrete pixel grid, we formally treat each color channel
261 of the image as a function $\Lambda \rightarrow \mathbb{R}$ evolving over a continuous spatial domain Λ .

262 Mathematically, we define Λ as a bounded open subset of \mathbb{R}^d , where $d \in \mathbb{N}$ represents the spatial dimension.
263 In our application, images are naturally two-dimensional, so we have $d = 2$. A common choice is to model
264 the image space Λ as $(0, 1)^2$ for normalized coordinates or as $\Lambda = (0, \text{width}) \times (0, \text{height})$ in the physical
265 pixel space, where width and height denote the dimensions of the image in pixels.

266 4.1 ANISOTROPIC FORWARD PROCESS: THE THEORETICAL MODEL

268 We now define the specific SPDE our forward process will satisfy. The equation describes how information
269 is progressively diffused in an anisotropic manner, ensuring a structured degradation of images and enhanced
270 reconstruction capability of geometric features in the generative sampling process (Algorithm 3.1).

270 To this end, we propose to model the forward process $(U_t)_{t \in \bar{I}}$ as the *formal* solution of
 271

$$272 \quad dU_t = b(t, U_t) dt + \sigma(t, U_t) dW_t \quad \text{for all } t \in \bar{I}, \quad (6)$$

273 where

$$274 \quad b(t, u) := \nabla \cdot g_1(t, \nabla u) \nabla u \quad (7)$$

275 for $(t, u) \in \bar{I} \times H^2(\Lambda)$,

$$276 \quad \sigma(t, u)v := g_2(t, \nabla u)v \quad (8)$$

277 for $(t, u) \in \bar{I} \times H^1(\Lambda)$ and $v \in Q^{1/2}L^2(\Lambda)$,

$$279 \quad g_i(t, x) := \frac{\alpha_i(t)}{\sqrt{1 + \left\| \frac{x}{\lambda_i(t)} \right\|^2}} \quad (9)$$

283 for $(t, x) \in I \times \mathbb{R}^d$, $\alpha_i : \bar{I} \rightarrow [0, \infty)$ and $\lambda_i : \bar{I} \rightarrow (0, \infty]$ are continuous and nondecreasing and $(W_t)_{t \in \bar{I}}$ is
 284 a Q -Wiener process (Da Prato & Zabczyk, 2014, Definition 4.2) with

$$285 \quad (Qf)(x) := \int_{\Lambda} q(x, y)f(y) dy \quad \text{for } x \in \Lambda \quad (10)$$

287 for $f \in L^2(\Lambda)$ and, for some $\ell \in [0, \infty)$,

$$289 \quad q(x, y) := \exp\left(-\frac{\|x - y\|}{\ell}\right) \quad \text{for } x, y \in \mathbb{R}^d. \quad (11)$$

291 As usual, $H^r(\Lambda)$ denotes the Hilbert Sobolev space (Renardy & Rogers, 2004, Chapter 7) of order $r \in \mathbb{N}_0$.

292 The reader, which is interested in, but unfamiliar with, stochastic analysis can just think about an Q -Wiener
 293 process as a generalization of a standard Wiener process (or Brownian motion) whose covariance operator is
 294 given by Q and refer to (Da Prato & Zabczyk, 2014; Lord et al., 2014) for details. Especially, integral operators
 295 of the form (10) are considered in (Lord et al., 2014, Definition 1.64). A discussion on the *Cameron-Martin*
 296 space $Q^{1/2}L^2(\Lambda)$ can be found in Section F.

297 (6) is the natural stochastic generalization of the deterministic Perona—Malik diffusion Perona & Malik (1990).
 298 In Section A, we provide design guidelines by detailing how the individual ingredients — the diffusivity coeffi-
 299 cients α_k (Section A.1.1), intensity coefficient (Section A.1.2), anisotropy coefficient λ_i (Section A.1.3) and
 300 *correlation length* ℓ (Section E.2) — of our anisotropic diffusion framework control the image transformations
 301 described by our forward process. A classification of the SPDE types arising from different parameter choices
 302 is provided in Section G.

303 4.2 PRACTICAL FORWARD AND BACKWARD PROCESSES

305 From a strictly mathematical perspective, ensuring correctness of the generative modeling methodology re-
 306 quires that the backward process is the *exact* time-reversal of the forward process and that both forward and
 307 backward processes are *exact* simulatable. However, for a complex SPDE like (6), exact simulation is impossi-
 308 ble. Both spatial and temporal discretization must be performed to obtain a practically simulatable process.

309 Spatial discretization, whether by Galerkin methods (as in (Lim et al., 2024)) or finite differences (as in our
 310 numerical scheme), inevitably leads to a projection onto a finite-dimensional SDE of the form

$$311 \quad d\tilde{U}_t = \tilde{b}\left(t, \tilde{U}_t\right) dt + \tilde{\sigma}\left(t, \tilde{U}_t\right) d\tilde{W}_t \quad \text{for all } t \in \bar{I}, \quad (12)$$

313 where $\tilde{b} : \mathbb{R}^D \rightarrow \mathbb{R}^D$, $\tilde{\sigma} : \mathbb{R}^D \rightarrow \mathbb{R}^{D \times D}$, and $(\tilde{W}_t)_{t \geq 0}$ is a D -dimensional Brownian motion for some $D \in \mathbb{N}$.

314 This procedure unavoidably introduces approximation error, which is further compounded by the subsequent
 315 temporal discretization. In special cases — such as the parabolic SPDEs with additive noise considered in Lim
 316 et al. (2024) — this may be acceptable. Nevertheless, if one uses the time-reversal of the infinite-dimensional
 317 SPDE as the backward process, a mismatch arises between the processes that can actually be simulated and the
 318 theoretical foundation of their use, and exact correctness of the backward process is no longer guaranteed.

319 In our framework, we consider a substantially more complex SPDE — with gradient-dependent nonlinear drift
 320 and gradient-dependent multiplicative noise — than in previous SPDE-based SBGMs (e.g., Lim et al. (2024)).
 321 In our case, no fixed eigenbasis diagonalizes the drift operator or the noise, and modal decoupling is unavailable.

322 We therefore adopt a different strategy: we first apply spatial discretization, thereby defining the actual forward
 323 process directly as the finite-dimensional SDE (12). The backward process is then taken as the time-reversal of
 this SDE, ensuring that both forward and backward dynamics are defined at the same level of approximation.

324 As a result, the only simulation error arises from temporal discretization, rather than from a combination of
 325 spatial *and* temporal discretizations.

326 The concrete numerical scheme leading to the SDE (12) used in our experiments in Section 5 is described in
 327 Section H. Theoretical existence of a solution to (12) is verified by Pascucci (2011, Theorem 9.11). That the
 328 corresponding backward process also satisfies an SDE follows from classical results (Haussmann & Pardoux,
 329 1986; Anderson, 1982). For the explicit form of this SDE we refer to Section D.

331 4.3 RESIDUAL DEPENDENCE ON THE INITIAL STATE

333 Introducing anisotropy inherently induces a residual dependence on the initial state. Our framework contains
 334 many user-definable parameters. For meaningful use in a generative modeling context, the modeling process
 335 that determines these parameters should ensure that the information contained in the initial state is almost en-
 336 tirely destroyed by the terminal time T . Conceptually, anisotropy should only serve to prolong the preservation
 337 of certain structures (such as edges) to facilitate the reconstruction of geometric features during data generation.

338 For specific parameter choices in our framework, the forward process reduces to an Ornstein–Uhlenbeck pro-
 339 cess (as in the instance described in Section 5.3), or it can be designed such that the distribution at the terminal
 340 time T is close to a known Gaussian distribution (as in the instance described in Section 5.2). In such cases,
 341 the prior distribution can be replaced by this closed-form distribution when performing Algorithm 3.1.

342 Whether the prior distribution admits a closed form (or a tractable approximation) depends on the chosen
 343 parameters. If not, a sample from the prior can be generated by simulating the forward process up to the
 344 terminal time. In either case, mathematical correctness is ensured; see our discussion in Section 3.3.

345 5 NUMERICAL STUDY

347 Intuitively, the anisotropy introduced by our anisotropic diffusion framework helps preserve structural infor-
 348 mation over longer time scales during the forward transformation. This, in turn, facilitates learning and recon-
 349 struction of geometric features during data generation.

350 In this section, we empirically validate this intuition through a numerical study on unconditional image gener-
 351 ation. We compare our framework against three baselines: Rissanen et al. (2023), the *variance exploding SDE*
 352 (*VESDE*) from Song et al. (2021c), and the *flow matching / optimal transport* method of Lipman et al. (2022),
 353 which is state of the art at the time of writing.

354 We focus on two specific instances of our framework: an isotropic version (*Ours (isotropic)* described in Sec-
 355 tion 5.3) and an anisotropic version (*Ours (anisotropic)* described in Section 5.2) of a stochastic heat equation.
 356 The isotropic variant is included primarily for educational purposes, as it demonstrates that the model equation
 357 of Rissanen et al. (2023) arises as a special case of our framework.

358 Both Rissanen et al. (2023) and Song et al. (2021c) provide especially relevant baselines, since they are likewise
 359 based on S(P)DEs. In contrast to our approach, however, their drift and diffusion coefficients are isotropic. Con-
 360 ceptually, the only difference between their methods and *Ours (anisotropic)* is the introduction of anisotropy.
 361 This design choice allows us to attribute observed improvements directly to anisotropy.

362 We now describe the specific instances of our framework considered in the numerical study (Sections 5.1 to 5.3),
 363 followed by a detailed account of the experiments in Section 5.4.

365 5.1 PURE ISOTROPIC NOISE (SONG ET AL. (2021C))

$$367 dU_t = \alpha_2(t) dW_t \quad \text{for all } t \in \bar{I}. \quad (13)$$

368 The SPDE formulated in (13) defines a Gaussian process, of which the *VESDE* considered in Song et al. (2021c)
 369 is a specific instance. It has a vanishing drift term, $b = 0$, and an isotropic diffusion coefficient, σ . Intuitively,
 370 this corresponds to a process in which an increasing amount of noise is added to the data over time.

372 5.2 ANISOTROPIC STOCHASTIC HEAT EQUATION WITH ISOTROPIC NOISE (OURS 373 (ANISOTROPIC))

$$375 dU_t = \nabla \cdot \frac{\alpha_1(t)}{\sqrt{1 + \left\| \frac{\nabla U_t}{\lambda_1} \right\|^2}} \nabla U_t dt + \alpha_2(t) dW_t \quad \text{for all } t \in \bar{I} \quad (14)$$

$$376$$

$$377$$

378 The SPDE in (14) is a genuinely anisotropic instance of our general anisotropic diffusion framework (6), where
 379 the drift gradually transitions from anisotropy to isotropy while the diffusion coefficient remains isotropic. We
 380 consider *geometric* transitions of the form

$$382 \quad \alpha_i(t) := \alpha_i^{\min} \left(\frac{\alpha_i^{\max}}{\alpha_i^{\min}} \right)^{\frac{t}{T}} \quad \text{for } t \in \bar{I} \quad (15)$$

384 for $0 < \alpha_i^{\min} < \alpha_i^{\max}$ (other common transitions are shown in Figure 5). Specifically, α_1 increases geometrically
 385 from $\alpha_1^{\min} = 0.5$ to $\alpha_1^{\max} = 2 \cdot \text{image_size}$, while the anisotropy coefficient λ_1 ensures a slow transition
 386 from anisotropy to isotropy via

$$387 \quad \lambda_1(t) := \lambda_1^{\min} \frac{e^{kT} - 1}{e^{k(T-t)} - 1} \quad \text{for } t \in \bar{I} \quad (16)$$

390 with $\lambda_1^{\min} = 0.025$ and $k = 1/2$ (see Figure 5 for a visualization). The intensity coefficient α_2 also in-
 391 creases geometrically, from $\alpha_2^{\min} = 0.01$ to $\alpha_2^{\max} = 2.0$. With $\lambda_2 \equiv \infty$, the diffusion coefficient remains
 392 spatially isotropic throughout. We set the noise correlation length to $\ell = 0$, leading to a *cylindrical* Wiener
 393 process $(W_t)_{t \in \bar{I}}$ and hence spatially white noise. The corresponding forward and backward processes are
 394 visualized in Figure 8.

395 Since $\lambda_1(t) \rightarrow \infty$ as $t \rightarrow T$, the numerical simulation (12) of the SPDE (14) is, at least approximatively,
 396 conditionally Gaussian given the initial state. Hence, prior sampling can be performed from a closed-form
 397 (Gaussian) distribution in the implementation.

398 5.3 STOCHASTIC HEAT EQUATION WITH ISOTROPIC NOISE 399 (RISSANEN ET AL. (2023) AND OURS (ISOTROPIC))

$$401 \quad dU_t = \alpha_1(t) \Delta U_t dt + \alpha_2(t) dW_t \quad \text{for all } t \in \bar{I}. \quad (17)$$

402 The SPDE in (17) has an isotropic drift b and a *small-scale* isotropic diffusion coefficient σ . It closely resembles
 403 the forward process considered by Rissanen et al. (2023). Intuitively, the data is smoothed over time, and a small
 404 amount of noise is injected to make the forward process stochastic, which is required to ensure that the reverse
 405 process is well-posed (see Section 3.2).

406 For α_1 , we again use a geometric transition (15) with $\alpha_1^{\min} = 0.5$ and $\alpha_1^{\max} = 2 \cdot \text{image_size}$. A minor
 407 difference from Rissanen et al. (2023) is that we do not keep the intensity coefficient α_2 constant; instead, it
 408 increases slightly over time (while remaining small) under a geometric transition equation 15 with $\alpha_2^{\min} = 0.01$
 409 and $\alpha_2^{\max} = 0.5$. Formally, the anisotropy coefficients are fixed as $\lambda_1 = \lambda_2 \equiv \infty$. The noise correlation length
 410 is again set to $\ell = 0$, so we also work with a cylindrical Wiener process $(W_t)_{t \in \bar{I}}$ here.

411 With the parameter choices described above, the numerical simulation (12) of the SPDE (17) is conditionally
 412 Gaussian given the initial state. Consequently, prior sampling can be carried out from a closed-form (Gaussian)
 413 distribution in the implementation.

414 5.4 EXPERIMENTS

416 In our experiments, our methods *Ours (anisotropic)* and *Ours (isotropic)* used both the predictor and corrector
 417 steps of Algorithm 3.1. Each corrector step consists of a single ULA iteration. We simulated up to a final time
 418 of $T = 2$ with a numerical step size of 0.001 in the predictor step, resulting in 1999 score function evaluations.

419 **Limitations** The flexibility of our framework opens the door to exploring a wide range of anisotropic for-
 420 ward processes beyond *Ours (anisotropic)*, which may further enhance quality. However, due to hardware
 421 resource limitations, this study is limited to the single parameter configuration given by *Ours (anisotropic)*, a
 422 small number of datasets listed below, and the restricted set of baseline methods mentioned above.

424 **Test datasets** We trained all generative models on the CIFAR-10 (Krizhevsky et al., 2009), CelebA (Liu
 425 et al., 2015), ImageNet2012 (Russakovsky et al., 2015), and LSUN/church_outdoor (Yu et al., 2015) datasets.

427 **Evaluation** We assess the quality of generated samples using standard metrics in generative image mod-
 428eling: the Inception Score (IS) (Salimans et al., 2016), Fréchet Inception Distance (FID) (Heusel et al., 2017),
 429 and Kernel Inception Distance (KID) (Błaszczyk et al., 2021). To ensure consistent evaluation, we used the
 430 implementation (Song et al., 2021b) provided by Song et al. (2021c) and regenerated all samples — including
 431 those for baseline methods — to compute these metrics. Because metric implementations vary slightly across
 toolsets, our reported values may differ from those originally published. This makes it particularly important
 that all models are evaluated under identical conditions.

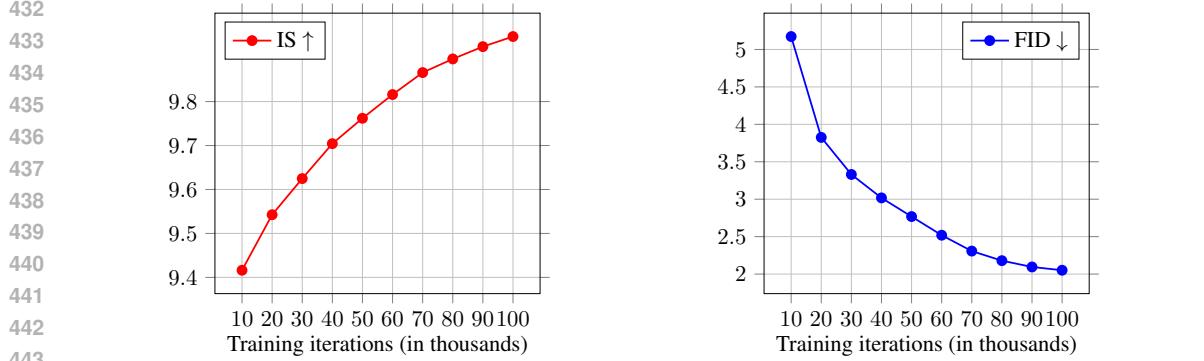


Figure 2: Evaluation metrics on CIFAR-10 between 10k and 100k training iterations with *Ours (anisotropic)*. We employed a fine-tuning strategy, initializing training from the checkpoint of Song et al. (2021c). For reference, applying our sampling procedure directly to the Song et al. (2021c) checkpoint — without additional training — results in an IS of 1, FID of 678.258, and KID of 0.872. After 100k training iterations, these values improve substantially to an IS of 9.978, FID of 2.041, and KID of 6.411.

Quantitative comparison In Table 1, we report the evaluation metrics of samples generated by the different methods on the test datasets. All values reported for *Ours (isotropic)* and *Ours (anisotropic)* are based on training from scratch for 200,000 steps. Since no checkpoints are provided by the authors, we retrained the method of Lipman et al. (2022) on all datasets using the official configuration from their codebase for the same number of steps. The same procedure was applied to Rissanen et al. (2023), although we restricted evaluation to CIFAR-10, given the method’s already non-competitive performance on this dataset. For Song et al. (2021c), we trained CelebA and ImageNet2012 from scratch, while for CIFAR-10 and LSUN/church_outdoor we relied on the official checkpoints released by the authors.

Figure 2 illustrates how quickly (i.e., after how few training iterations) *Ours (anisotropic)* improves the evaluation metrics when initialized from a model pretrained with Song et al. (2021c). Notably, according to the original authors, continuing training with their own method did not yield further metric improvements.

	IS \uparrow	FID \downarrow	KID \downarrow
Ours (anisotropic)	10.2 / 3.1 / 13.5 / 3.6	2.0 / 2.4 / 19.1 / 5.9	6.1e-4 / 1.7e-3 / 1.9e-2 / 4.6e-3
Lipman et al. (2022)	9.2 / 2.4 / 10.5 / -	2.0 / 2.3 / 26.8 / -	7.1e-4 / 1.4e-3 / 3.4e-2 / -
Song et al. (2021c)	9.8 / 2.5 / 12.3 / 2.5	7.1 / 3.7 / 24.0 / 16.7	6.6e-4 / 2.6e-3 / 2.5e-2 / 1.2e-2
Vandersanden et al. (2024)	7.1 / 2.7 / - / 3.4	28.7 / 12.0 / - / 49.1	2.2e-2 / 8.4e-3 / - / 4.3e-2

Table 1: Each column within the IS, FID, and KID metrics corresponds, in order, to results on CIFAR-10 (32x32), CelebA (64x64), ImageNet2012 (64x64), and LSUN/church_outdoor (256x256). Higher IS is better, while lower FID and KID are better. On CIFAR-10, we also compared *Ours (isotropic)* with Rissanen et al. (2023). *Ours (isotropic)* achieved metrics of 8.8 / 19.6 / 1.6e-02, while Rissanen et al. (2023) produced the worse metrics 5.9 / 84.3 / 7.2e-2. *Ours (isotropic)* and *Ours (anisotropic)* refer to the isotropic and anisotropic stochastic heat equation with isotropic noise described in Section 5.3 and Section 5.2, respectively.

Hardware resources All experiments were conducted on a server equipped with 8 \times NVIDIA Tesla H100 NVL GPUs (94 GB HBM3 each, PCIe 5.0) and 2 \times AMD EPYC 9554 CPUs (64 cores / 128 threads each, 3.1–3.75 GHz, Genoa microarchitecture, 256 MB L3 cache).

Hyperparameters and architecture For the implementation of our framework, we adopted the NCSN++ (continuous) network architecture from Song et al. (2021c). The training parameters — in particular, a learning rate of 2e-4 using the Adam optimizer (Kingma, 2014) — are taken from Song et al. (2021c) for both CIFAR-10, CelebA and LSUN/church_outdoor. For ImageNet2012, we reuse the parameter settings for CelebA.

6 CONCLUSION

This work demonstrated that introducing anisotropy into SBGM can be practically superior to traditional isotropic approaches. We extended SBGM by proposing a novel class of anisotropic diffusion processes theoretically founded on SPDEs. These processes generalize the conventional isotropic setting by enabling geometry-aware transformations that align the generative sampling dynamics more closely with intrinsic geometric structures in the data. Beyond the theoretical model, we presented a proof-of-concept implementation showing that this anisotropic framework can preserve fine-grained structural information over longer time scales and achieve competitive generative performance, supporting the underlying intuition. Exploring the broader parameter space of our anisotropic diffusion framework therefore constitutes a promising direction for future work. Together, these contributions broaden the design space of SBGMs and indicate that anisotropic transformations have the potential to further improve sample quality and convergence.

486 **7 REPRODUCIBILITY STATEMENT**

488 The theoretical framework underlying our approach is presented in Section 4. The specific instances used
 489 in our experiments, *Ours (isotropic)* and *Ours (anisotropic)*, are described in (17) and Section 5.2, including
 490 the corresponding parameter choices. The numerical simulation scheme for both forward and backward pro-
 491 cesses is detailed in Section H. Finally, the training hyperparameters and network architecture are specified in
 492 Section 5.4.

493 **8 ETHICS STATEMENT**

496 We acknowledge that diffusion-based generative models carry potential risks. In particular, they can be misused
 497 for generating deepfakes, which may facilitate misinformation, deception, or harassment, and they pose privacy
 498 risks if applied to generate images of individuals without consent. Our method is not intended for such purposes,
 499 and we ensured that our experiments relied solely on publicly available benchmark datasets without personally
 500 identifiable information. Moreover, our models are trained in a way that avoids replication or memorization of
 specific training images.

501 At the same time, diffusion models also offer positive contributions, including applications in medical imaging,
 502 scientific visualization, and artistic content creation. We believe that, with responsible use, the potential benefits
 503 of this line of research outweigh its risks.

505 **REFERENCES**

507 Brian D.O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12
 508 (3):313–326, 1982. ISSN 0304-4149. doi: [https://doi.org/10.1016/0304-4149\(82\)90051-5](https://doi.org/10.1016/0304-4149(82)90051-5). URL <https://www.sciencedirect.com/science/article/pii/0304414982900515>.

510 Mikołaj Bińkowski, Danica J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans, 2021.
 511 URL <https://arxiv.org/abs/1801.01401>.

512 Giuseppe Da Prato and Jerzy Zabczyk. *Stochastic Equations in Infinite Dimensions*. Encyclopedia of Mathe-
 513 matics and its Applications. Cambridge University Press, 2 edition, 2014.

514 H. Föllmer and A. Wakolbinger. Time reversal of infinite-dimensional diffusions. *Stochastic Pro-
 515 cesses and their Applications*, 22(1):59–77, 1986. ISSN 0304-4149. doi: [https://doi.org/10.1016/0304-4149\(86\)90114-6](https://doi.org/10.1016/0304-4149(86)90114-6). URL <https://www.sciencedirect.com/science/article/pii/0304414986901146>.

518 U. G. Haussmann and E. Pardoux. Time Reversal of Diffusions. *The Annals of Probability*, 14(4):1188 – 1205,
 519 1986. doi: 10.1214/aop/1176992362. URL <https://doi.org/10.1214/aop/1176992362>.

521 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by
 522 a two time-scale update rule converge to a local nash equilibrium. In I. Guyon, U. Von Luxburg, S. Bengio,
 523 H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing
 524 Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/8a1d694707eb0fefef65871369074926d-Paper.pdf.

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

528 Emiel Hoogeboom and Tim Salimans. Blurring diffusion models. *arXiv preprint arXiv:2209.05557*, 2022.

530 Rongjie Huang, Max WY Lam, Jun Wang, Dan Su, Dong Yu, Yi Ren, and Zhou Zhao. Fastdiff: A fast
 531 conditional diffusion model for high-quality speech synthesis. *arXiv preprint arXiv:2204.09934*, 2022.

532 Martin Hutzenthaler and Arnulf Jentzen. Numerical approximations of stochastic differential equations with
 533 non-globally lipschitz continuous coefficients. *Memoirs of the American Mathematical Society*, 236(1112):
 534 0–0, July 2015. ISSN 1947-6221. doi: 10.1090/memo/1112. URL <http://dx.doi.org/10.1090/memo/1112>.

536 Amirhossein Kazerouni, Ehsan Khodapanah Aghdam, Moein Heidari, Reza Azad, Mohsen Fayyaz, Ilker Haci-
 537 haliloglu, and Dorit Merhof. Diffusion models in medical imaging: A comprehensive survey. *Medical Image
 538 Analysis*, 88:102846, 2023.

539 Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

540 Peter E. Kloeden and Eckhard Platen. *Numerical Solution of Stochastic Differential Equations*. Springer Berlin,
 541 Heidelberg, 1 edition, 1992. URL <https://doi.org/10.1007/978-3-662-12616-5>.

542

543 Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model
 544 for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020.

545

546 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

547

548 Jae Hyun Lim, Nikola B Kovachki, Ricardo Baptista, Christopher Beckham, Kamyar Azizzadenesheli, Jean
 549 Kossaifi, Vikram Voleti, Jiaming Song, Karsten Kreis, Jan Kautz, et al. Score-based diffusion models in
 function space. *arXiv preprint arXiv:2302.07400*, 2023.

550

551 Sungbin Lim, Eun Bi Yoon, Taehyun Byun, Taewon Kang, Seungwoo Kim, Kyungjae Lee, and Sungjoon
 552 Choi. Score-based generative modeling through stochastic evolution equations in hilbert spaces. *Advances
 553 in Neural Information Processing Systems*, 36, 2024.

554

555 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for genera-
 556 tive modeling. *arXiv preprint arXiv:2210.02747*, 2022.

557

558 Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaou Tang. Deep learning face attributes in the wild. In
 559 *ICCV*, pp. 3730–3738. IEEE Computer Society, 2015. ISBN 978-1-4673-8391-2. URL <http://dblp.uni-trier.de/db/conf/iccv/iccv2015.html#LiuLWT15>.

560

561 Gabriel J. Lord, Catherine E. Powell, and Tony Shardlow. *An Introduction to Computational Stochastic PDEs*.
 562 Cambridge Texts in Applied Mathematics. Cambridge University Press, 2014.

563

564 Annie Millet, David Nualart, and Marta Sanz. Time reversal for infinite-dimensional diffusions. *Probability
 565 Theory and Related Fields*, 82(3):315–347, August 1989. ISSN 1432-2064. doi: 10.1007/BF00339991.
 566 URL <https://doi.org/10.1007/BF00339991>.

567

568 Andrea Pascucci. *PDE and Martingale Methods in Option Pricing*. Springer, 2011.

569

570 Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transac-
 571 tions on pattern analysis and machine intelligence*, 12(7):629–639, 1990.

572

573 Michael Renardy and Robert C. Rogers. *An Introduction to Partial Differential Equations*. Springer New York,
 574 NY, 2 edition, 2004.

575

576 Severi Rissanen, Markus Heinonen, and Arno Solin. Generative modelling with inverse heat dissipation. In *The
 577 Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=4PJUBT9f201>.

578

579 Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej
 580 Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale
 581 Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
 582 doi: 10.1007/s11263-015-0816-y.

583

584 Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen.
 585 Improved techniques for training gans. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and
 586 R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 29. Curran Associates,
 587 Inc., 2016. URL https://proceedings.neurips.cc/paper_files/paper/2016/file/8a3363abe792db2d8761d6403605aeb7-Paper.pdf.

588

589 Arne Schneuing, Charles Harris, Yuanqi Du, Kieran Didi, Arian Jamasb, Ilia Igashov, Weitao Du, Carla Gomes,
 590 Tom L Blundell, Pietro Lio, et al. Structure-based drug design with equivariant diffusion models. *Nature
 591 Computational Science*, 4(12):899–909, 2024.

592

593 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning
 594 using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265.
 595 PMLR, 2015.

596

597 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances
 598 in neural information processing systems*, 32, 2019.

599

600 Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with
 601 score-based generative models. *arXiv preprint arXiv:2111.08005*, 2021a.

602

603 Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole.
 604 Score-based generative modeling through stochastic differential equations. https://github.com/yang-song/score_sde, 2021b. Accessed: 2025-05-14.

594 Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole.
 595 Score-based generative modeling through stochastic differential equations. In *International Conference on*
 596 *Learning Representations*, 2021c. URL <https://openreview.net/forum?id=PxTIG12RRHS>.

597 Jente Vandersanden, Sascha Holl, Xingchang Huang, and Gurprit Singh. Edge-preserving noise for diffusion
 598 models. *arXiv preprint arXiv:2410.01540*, 2024.

599

600 Vikram Voleti, Christopher Pal, and Adam Oberman. Score-based denoising diffusion with non-isotropic gaus-
 601 sian noise models. *arXiv preprint arXiv:2210.12254*, 2022.

602 Tomer Weiss, Eduardo Mayo Yanes, Sabyasachi Chakraborty, Luca Cosmo, Alex M Bronstein, and Renana
 603 Gershoni-Poranne. Guided diffusion for inverse molecular design. *Nature Computational Science*, 3(10):
 604 873–882, 2023.

605 Zhen Xing, Qijun Feng, Haoran Chen, Qi Dai, Han Hu, Hang Xu, Zuxuan Wu, and Yu-Gang Jiang. A survey
 606 on video diffusion models. *ACM Computing Surveys*, 57(2):1–42, 2024.

607

608 Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Lsun: Construction of a large-scale image
 609 dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.

610 Xi Yu, Xiang Gu, Haozhi Liu, and Jian Sun. Constructing non-isotropic gaussian diffusion model using
 611 isotropic gaussian diffusion model for image editing. *Advances in Neural Information Processing Systems*,
 612 36:76839–76851, 2023.

613 Linqi Zhou, Aaron Lou, Samar Khanna, and Stefano Ermon. Denoising diffusion bridge models, 2023. URL
 614 <https://arxiv.org/abs/2309.16948>.

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 **A DESIGN GUIDELINE: CONTROLLING IMAGE TRANSFORMATIONS IN OUR**
 649 **FRAMEWORK**
 650

651 In this section, we provide design guidelines by detailing how the individual ingredients of our anisotropic
 652 diffusion framework control the image transformations described by our forward process.
 653

654 **A.1 HOW DRIFT AND DIFFUSION COEFFICIENT CONTRIBUTE TO DEGRADING INFORMATION**
 655

656 The drift b (7) introduces deterministic smoothing that respects the anisotropy defined by the *anisotropy coefficient* λ_1 . This smoothing is controlled by the *diffusivity coefficient* α_1 , with its strength and direction determined
 657 by the local image structure (captured by the gradients ∇U_t).
 658

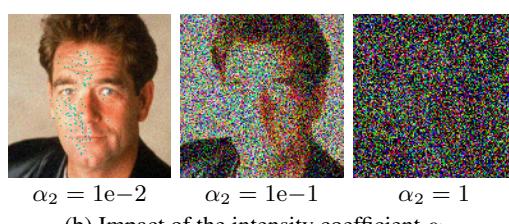
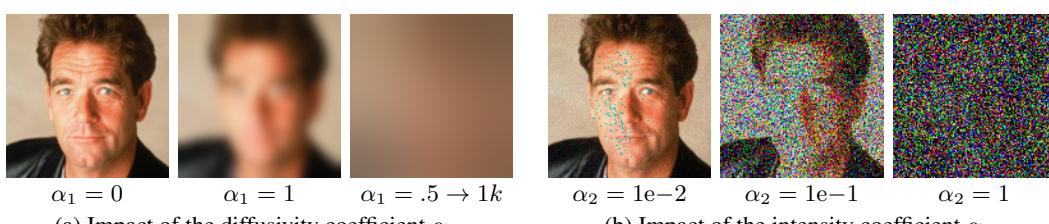
659 The diffusion coefficient σ (8) injects random noise into the image, which is also modulated anisotropically by
 660 λ_2 . The intensity of the noise is modulated by the *intensity coefficient* α_2 and destroys fine-grained details in a
 661 controlled manner, complementing the drift's smoothing effect.
 662

663 While the drift b (7) focuses on smoothing (deterministic destruction), the diffusion coefficient σ introduces
 664 stochasticity (random destruction). Together they degrade information in a controlled manner, providing an
 665 excellent framework for SBGM.
 666

667 To enable (stochastic) reconstruction, we need to make sure that the forward process destroys the data suffi-
 668 ciently. Anisotropy can be seen as preserving information to some extent. This preserved data can be destroyed
 669 in various ways. For example, if the anisotropy is in the drift, we can either the destroy the information over
 670 time by spreading a lot of isotropic noise at later time points or by transitioning to isotropy in the drift. The
 671 latter can be achieved by letting $\lambda_1(t) \rightarrow \infty$ as $t \rightarrow T$.
 672

673 Subsequently, we will give details on the effect of the diffusivity coefficients α_k (Section A.1.1), intensity
 674 coefficient (Section A.1.2), anisotropy coefficient λ_i (Section A.1.3) and *correlation length* ℓ (Section E.2).
 675

676 **A.1.1 DIFFUSIVITY COEFFICIENT α_1**
 677



682 The diffusivity coefficient α_1 crucially controls the rate of the (an-)isotropic smoothing in the drift (7) in the
 683 diffusion coefficient (8) of the SPDE (6), respectively.
 684

685 In the drift (7), a larger α_1 results in stronger smoothing, leading to a faster elimination of high-frequency
 686 details (e.g., geometric structure, like edges and corners, and textures) in the image. Conversely, smaller values
 687 of α_1 preserve more of the fine-grained details, slowing down the destruction of information.
 688

689 In Figure (a) we visualized the impact of α_1 , disabling the effect of the diffusion coefficient σ by setting
 690 $\alpha_2 = 0$. The third column is using a *geometric* transition $\alpha_1(t) = \alpha_1^{\min} (\alpha_1^{\max} / \alpha_1^{\min})^{t/T}$ with $\alpha_1^{\min} = .5$ and
 691 $\alpha_1^{\max} = 1k$. In Figure 5 (b) we show more choices for the diffusivity/intensity coefficients. The original image
 692 had a resolution of 128x128 pixels and the forward process was simulated up to time $T = 8$.
 693

694 **A.1.2 INTENSITY COEFFICIENT α_2**
 695

696 The intensity coefficient α_2 determines the intensity of the injected noise in the diffusion coefficient (8) of the
 697 SPDE (6).
 698

699 In the diffusion coefficient (8), larger α_2 increases the randomness in the image transformation, introducing
 700 more noise and accelerating the destruction of structured information. On the other hand, smaller α_2 reduces
 701 the randomness, preserving some of the original structure while still degrading information.
 702

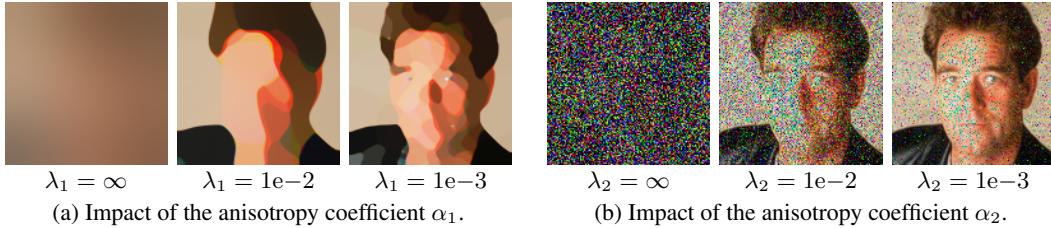
703 In Figure 3b we visualized the impact of α_2 , disabling the effect of the drift b by setting $\alpha_1 = 0$.
 704

705 **Impact on generative modeling** The progression of the α_i over time determines how quickly the image
 706 is degraded by the forward process.
 707

702 By modulating α_i , we can design a destruction process that aligns with the objective of generative modeling:
 703 creating a sequence of progressively smoother / less noise images while maintaining enough structure for the
 704 model to learn the reconstruction of meaningful samples with a high resemblance to the images from the dataset.

705 A suitable balance between α_1 in the drift (7) and α_2 in the diffusion diffusion coefficient (8) ensures that
 706 the degradation is smooth but irreversible, providing a structured data destruction trajectory for training the
 707 generative model.

708
 709 **A.1.3 ANISOTROPY COEFFICIENT λ_i**



719 The anisotropy coefficients α_i control the directional sensitivity of both the drift and diffusion terms by mod-
 720 ulating the degree of anisotropy in the transformation. The drift term (7) depends on $g_1(t, \nabla U_t)$, which intro-
 721 duces directional weighting to the smoothing property of the drift.

722 When λ_1 is small, the smoothing from the drift term (7) is highly anisotropic, meaning the process prefers
 723 certain directions for smoothing while preserving others (e.g., along edges in the image). This ensures that
 724 geometric structures are degraded in a structured manner. As λ_1 grows larger, the smoothing becomes more
 725 isotropic, uniformly degrading all directions and gradually eliminating all structural features.

726 The diffusion coefficient term (8), modulated by $g_2(t, \nabla U_t)$, introduces anisotropic noise. For smaller λ_2 , the
 727 noise is injected along specific directions, preserving certain patterns while destroying others. As λ_2 increases,
 728 the noise becomes isotropic, introducing randomness uniformly across the image and further accelerating struc-
 729 tured information loss. In the limit, when $\lambda_2 = \infty$, noise is injected isotropically.

730 In Figure 4a and Figure 4b we visualized the impact of λ_1 and λ_2 , disabling the effect of the diffusion coefficient
 731 σ and the drift b by setting $\alpha_2 = 0$ and $\alpha_1 = 0$, respectively. In Figure 5 (a) we depicted common choices for
 732 these anisotropy coefficients.

733 **Impact on generative modeling** The anisotropy coefficients λ_i allow for a structured destruction of in-
 734 formation. For example, by preserving geometric structures, like edges or corners, for longer time, the forward
 735 process provides richer intermediate representations, which can enhance the model’s ability to reconstruct these
 736 structures during sampling.

737 Structured anisotropic degradation may lead to better generative sampling, as the score-based model learns to
 738 reverse transformations that align with natural image statistics (e.g., edge preservation and texture destruction).

739 In contrast, overly isotropic processes (corresponding to large up to infinite λ_i) degrade the images uniformly,
 740 which may simplify the forward process but could result in reduced resemblance of the dataset images, espe-
 741 cially if they admit significant geometric patterns.

756 **B INCORPORATING ANISOTROPY IN FLOW MATCHING (FM)**
757758 Flow matching methods (Lipman et al., 2022) offer a flexible framework for generative modeling by con-
759 structing probability paths between distributions without requiring explicit SPDE formulations. In particular,
760 they support structured perturbations via non-isotropic noise, as demonstrated in edge-aware extensions such
761 as (Vandersanden et al., 2024). However, these perturbations remain spatially uniform and are typically con-
762 ditioned only on the initial dataset sample. This limits their ability to adapt to the evolving geometry of the
763 sample during generation.764 In contrast, our SPDE-based framework allows for spatially dependent, anisotropic diffusion that reacts dy-
765 namically to the evolving image gradients. As a result, structural features such as edges are preserved not only
766 based on their presence in the initial image but also as they emerge, weaken, or shift throughout the transfor-
767 mation. For example, if an initially weak edge becomes stronger over time, our formulation naturally reduces
768 diffusion across it. Conversely, if a previously strong edge fades, the diffusion increases, enabling appropriate
769 smoothing. In flow matching, by contrast, structural information is fixed at the start of the transformation and
770 cannot adapt to changes during generation. This may result in the preservation of features that should fade
771 away or in the undesired blurring of structures that only become salient later in the process.772 **C CONDITIONAL GENERATIVE MODELING THROUGH ANISOTROPIC
773 DIFFUSION BRIDGES**774 DDBM (Zhou et al., 2023) address a fundamentally different task from ours. Rather than generating images
775 unconditionally from noise, DDBM learn mappings between two given image distributions by modeling the
776 bridge dynamics connecting them. While their method can, in principle, be applied to unconditional generation
777 by choosing noise as the source distribution, their framework is primarily designed for conditional tasks such as
778 image-to-image translation or editing. In this first application of our framework, however, we focus exclusively
779 on unconditional image generation and therefore do not include a direct comparison.780 Extending our framework to enable bridging between arbitrary source and target distributions, akin to DDBMs
781 is part of future work.782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

810 D SAMPLING FROM THE BACKWARD PROCESS

812 If $(\bar{U}_t)_{t \in \bar{I}}$ is a Markov process, sampling in the predictor step is done by following its transition dynamics.
 813 This is the case, for example, in DDPM and in the situation we encounter in our anisotropic diffusion framework
 814 from Section 4 as well.

815 In SDE-driven SBGMs, we are in the special situation, where the forward process is the solution to an SDE of
 816 the form

$$818 \quad dU_t = b(t, U_t) dt + \sigma(t, U_t) dW_t \quad \text{for all } t \in \bar{I} \quad (18)$$

819 for some Q -Wiener process. If a suitable regularity condition (Haussmann & Pardoux, 1986) is in place, then

$$821 \quad d\bar{U}_t = \bar{b}(t, \bar{U}_t) dt + \bar{\sigma}(t, \bar{U}_t) dW_t \quad \text{for all } t \in \bar{I}, \quad (19)$$

823 where

$$825 \quad \bar{b}(t, u) := \text{tr } D_u \Sigma(T - t, u) + \Sigma(T - t, u) s(T - t, u) - b(T - t, u); \quad (20)$$

$$826 \quad \bar{\sigma}(t, u) := \sigma(T - t, u) \quad (21)$$

828 for $(t, x) \in \bar{I} \times \mathbb{R}^D$ and

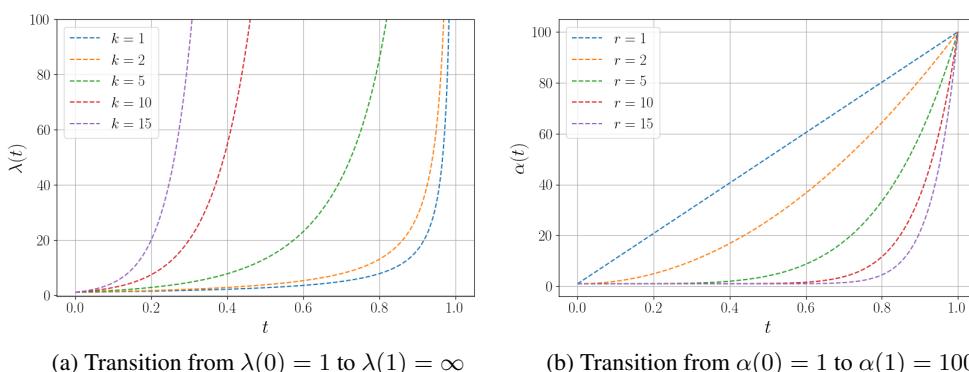
$$830 \quad \Sigma := \left(\sigma Q^{\frac{1}{2}} \right) \left(\sigma Q^{\frac{1}{2}} \right)^*. \quad (22)$$

832 In this case, the sampling in the predictor step is performed using a method for the numerical solution of a SDE,
 833 with the Euler-Maruyama method (Kloeden & Platen, 1992) being the simplest approach.

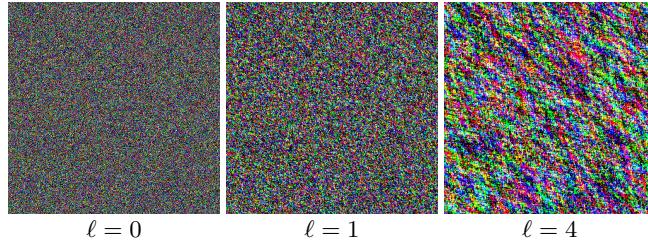
835 **Discussion** We emphasize that, in the practical application of our framework, it is the finite-dimensional
 836 SDE — obtained via the numerical scheme simulating our SPDE (6) described in Section H — that must be
 837 reversed in time, not the SPDE (6) itself. While — under a suitable set of assumptions — time-reversal of the
 838 SPDE (6) is theoretically possible (Föllmer & Wakolbinger, 1986; Millet et al., 1989), the results presented in
 839 Haussmann & Pardoux (1986); Anderson (1982) are sufficient for our purposes, as they apply directly to the
 840 finite-dimensional SDE setting.

842 E FURTHER VISUALIZATIONS OF THE 843 PARAMETERS OF OUR ANISOTROPIC DIFFUSION FRAMEWORK

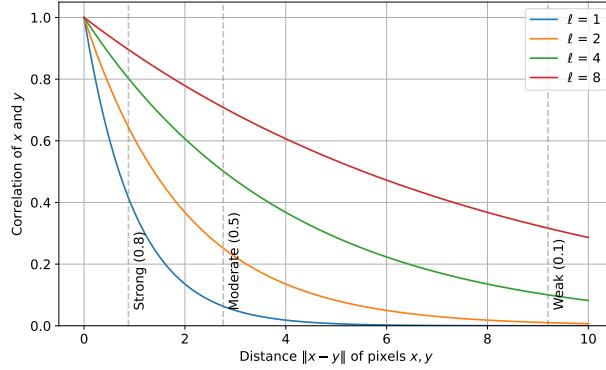
845 E.1 DIFFUSIVITY, INTENSITY AND ANISOTROPY COEFFICIENTS α_1 , α_2 AND λ_i



860 Figure 5: (a) Visualization of the common choice $\lambda(t) := \lambda^{\min} \frac{e^{kT} - 1}{e^{k(T-t)} - 1}$ for the anisotropy coefficients (Sec-
 861 tion A.1.3). (b) Visualization of the common choice $\alpha(t) := \alpha^{\min} + (\alpha^{\max} - \alpha^{\min}) \left(\frac{t}{T} \right)^r$ for the diffusivity
 862 (Section A.1.1) and intensity coefficients (Section A.1.2).

864 E.2 WIENER PROCESS COVARIANCE CORRELATION LENGTH ℓ
865873 Figure 6: Noise spread by the Wiener process W for correlation lengths $\ell = 0, 1, 4$.
874

875 With the *correlation length* ℓ we effectively set the spatial extent of correlation among the spread noise. A
876 large ℓ means noise at a given pixel is correlated with noise in a broader region around that pixel. In contrast,
877 a small ℓ means the correlation is localized, and noise at a pixel primarily affects nearby pixels. A given pixel
878 is roughly strongly (correlation $> .8$), moderately (correlation $> .5$) and weakly (correlation $> .1$ correlated
879 to all pixels in a radius of $.22\ell$, $.69\ell$ and 2.3ℓ , respectively (cf. Figure 7). For our experiments in Section 5,
880 we only considered the border case $\ell = 0$ in which (11) formally reduces to the Dirac delta $\delta_x(y)$ and hence
881 $Q = \text{id}_{L^2(\Lambda)}$ and hence $(W_t)_{t \in \bar{I}}$ is actually cylindrical. In Figure 6 we compared different choices for ℓ .
882

894 Figure 7: Spatial exponential decay of correlations in the Wiener process W for correlation lengths $\ell =$
895 1, 2, 4, 8. For $\ell = 0$, the noise spread between pixels is independent.
896897 F THE CAMERON-MARTIN SPACE $Q^{\frac{1}{2}}L^2(\Lambda)$
898

900 The space $Q^{\frac{1}{2}}L^2(\Lambda)$ is usually called a *Cameron-Martin space*; see (Da Prato & Zabczyk, 2014, Chapter I.4)
901 or (Lord et al., 2014, Definition 10.15). To understand and apply our paper, it is only important to be aware
902 of the set-theoretic definition $Q^{\frac{1}{2}}L^2(\Lambda) := \{Q^{\frac{1}{2}}u : u \in L^2(\Lambda)\}$. In plain English, it is the space of all
903 transformations of $L^2(\Lambda)$ -functions under the operator $Q^{\frac{1}{2}}$. For more theoretical considerations, one important
904 aspect is that it inherits a Hilbert space structure from $L^2(\Lambda)$.
905

906 G SPDE CLASSIFICATION
907

908 In general, (6) is a quasilinear parabolic SPDE with multiplicative noise.

909 If g_1 does not depend on the second argument, (6) is a semilinear parabolic SPDE with multiplicative noise:

$$910 \quad dU_t = \alpha_1(t) \Delta U_t dt + \sigma(t, U_t) dW_t \quad \text{for all } t \in \bar{I}. \quad (23)$$

912 If g_2 does not depend on the second argument, (6) is a quasilinear parabolic SPDE with additive noise:

$$914 \quad dU_t = b(t, U_t) dt + \alpha_2(t) dW_t \quad \text{for all } t \in \bar{I}. \quad (24)$$

915 Finally, if g_1 and g_2 both do not depend on the second argument, (6) is a linear parabolic SPDE with additive
916 noise:

$$917 \quad dU_t = \alpha_1(t) \Delta U_t dt + \alpha_2(t) dW_t \quad \text{for all } t \in \bar{I}. \quad (25)$$

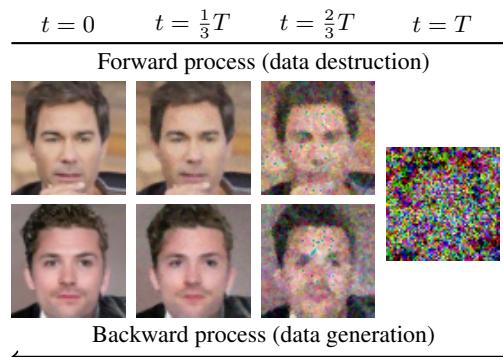


Figure 8: Visualization of the forward and backward processes corresponding to our anisotropic stochastic heat equation with isotropic noise (Section 5.2).

H NUMERICAL SIMULATION

For the numerical simulation of the forward and backward processes, (6) and (1), we modeled the image space Λ as $\Lambda = (0, d_1) \times (0, d_2)$ and decomposed the boundary $\partial\Lambda$ according to

$$\partial_L \Lambda := \{0\} \times [0, d_2]; \quad (26)$$

$$\partial_T \Lambda := [0, d_1] \times \{d_2\}; \quad (27)$$

$$\partial_R \Lambda := \{d_1\} \times (0, d_2]; \quad (28)$$

$$\partial_B \Lambda := (0, d_1] \times \{0\} \quad (29)$$

into its left, top, right and bottom part. We discretized the derivatives using a mixture of forward, backward and central finite differences, respecting Neumann boundary conditions.

H.1 DOMAIN DISCRETIZATION

After discretization, we decomposed the discretized domain $D = \{0, \dots, d_1\} \times \{0, \dots, d_2\}$ in the same spirit into its interior, left, top, right and bottom part:

$$D^\circ := \{1, \dots, d_1 - 2\} \times \{1, \dots, d_2 - 2\}; \quad (30)$$

$$\partial_L D := \{0\} \times \{0, \dots, d_2 - 2\}; \quad (31)$$

$$\partial_T D := \{0, \dots, d_2 - 2\} \times \{d_2 - 1\}; \quad (32)$$

$$\partial_R D := \{d_1 - 1\} \times \{1, \dots, d_2 - 1\}; \quad (33)$$

$$\partial_B D := \{1, \dots, d_1 - 1\} \times \{0\}. \quad (34)$$

972
973

H.2 SPATIAL DISCRETIZATION

974
975

For the specific finite difference approximation we have chosen, we ended up with the discretized drift being given by

976

$$\tilde{b}(t, u)_i := \begin{cases}
 g_1 \left(t, \begin{pmatrix} u_{i_1+2 \wedge d_1-1, i_2} - u_i \\ u_{i_1+1, i_2+1} - u_{i_1+1, i_2-1} \end{pmatrix} \right) (u_{i_1+1, i_2} - u_i) \\
 - g_1 \left(t, \begin{pmatrix} u_i - u_{i_1-2 \vee 0, i_2} \\ u_{i_1-1, i_2+1} - u_{i_1-1, i_2-1} \end{pmatrix} \right) (u_i - u_{i_1-1, i_2}) \\
 + g_1 \left(t, \begin{pmatrix} u_{i_1+1, i_2+1} - u_{i_1-1, i_2+1} \\ u_{i_1, i_2+2 \wedge d_2-1} - u_i \end{pmatrix} \right) (u_{i_1, i_2+1} - u_i) \\
 - g_1 \left(t, \begin{pmatrix} u_{i_1+1, i_2-1} - u_{i_1-1, i_2-1} \\ u_i - u_{i_1, i_2-2 \vee 0} \end{pmatrix} \right) (u_i - u_{i_1, i_2-1}) & , \text{ if } i \in D^\circ; \\
 \left(g_1 \left(t, \begin{pmatrix} u_{i_1+2 \wedge d_1-1, i_2} - u_i \\ 0 \end{pmatrix} \right) + g_1(t, 0) \right) (u_{i_1+1, i_2} - u_i) \\
 + \left(g_1 \left(t, \begin{pmatrix} 0 \\ u_{i_1, i_2+2 \wedge d_2-1} - u_i \end{pmatrix} \right) + g_1(t, 0) \right) (u_{i_1, i_2+1} - u_i) & , \text{ if } i \in \partial_L D \text{ with } i_2 = 0; \\
 \left(g_1 \left(t, \begin{pmatrix} u_{i_1+2 \wedge d_1-1, i_2} - u_i \\ u_{i_1+1, i_2+1} - u_{i_1+1, i_2-1} \end{pmatrix} \right) \right. \\
 \left. + g_1 \left(t, \begin{pmatrix} 0 \\ u_{i_1+1, i_2+1} - u_{i_1+1, i_2-1} \end{pmatrix} \right) \right) (u_{i_1+1, i_2} - u_i) \\
 + g_1 \left(t, \begin{pmatrix} 0 \\ u_{i_1, i_2+2 \wedge d_2-1} - u_i \end{pmatrix} \right) (u_{i_1, i_2+1} - u_i) \\
 - g_1 \left(t, \begin{pmatrix} 0 \\ u_i - u_{i_1, i_2-2 \vee 0} \end{pmatrix} \right) (u_i - u_{i_1, i_2-1}) & , \text{ if } i \in \partial_L D \text{ with } i_2 > 0; \\
 \left(g_1 \left(t, \begin{pmatrix} u_{i_1+2 \wedge d_1-1, i_2} - u_i \\ 0 \end{pmatrix} \right) + g_1(t, 0) \right) (u_{i_1+1, i_2} - u_i) \\
 + \left(g_1 \left(t, \begin{pmatrix} 0 \\ u_i - u_{i_1, i_2 \vee 0} \end{pmatrix} \right) + g_1(t, 0) \right) (u_{i_1, i_2-1} - u_i) & , \text{ if } i \in \partial_T D \text{ with } i_1 = 0; \\
 g_1 \left(t, \begin{pmatrix} u_{i_1+2 \wedge d_1-1, i_2} - u_i \\ 0 \end{pmatrix} \right) \\
 - g_1 \left(t, \begin{pmatrix} u_i - u_{i_1-2 \vee 0, i_2} \\ 0 \end{pmatrix} \right) (u_i - u_{i_1-1, i_2}) \\
 + \left(g_1 \left(t, \begin{pmatrix} u_{i_1+1, i_2-1} - u_{i_1-1, i_2-1} \\ 0 \end{pmatrix} \right) \right. \\
 \left. + g_1 \left(\begin{pmatrix} u_{i_1+1, i_2-1} - u_{i_1-1, i_2-1} \\ u_i - u_{i_1, i_2-2 \vee 0} \end{pmatrix} \right) \right) (u_{i_1, i_2-1} - u_i) & , \text{ if } i \in \partial_T D \text{ with } i_1 > 0
 \end{cases} \quad (35)$$

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026 and
1027

$$\tilde{b}(t, u)_i := \begin{cases}
\left(g_1 \left(t, \begin{pmatrix} u_i - u_{i_1-2 \vee 0, i_2} \\ 0 \end{pmatrix} \right) + g_1(t, 0) \right) (u_{i_1-1, i_2+1} - u_i) \\
+ \left(g_1 \left(t, \begin{pmatrix} 0 \\ u_i - u_{i_1, i_2-2 \vee 0} \end{pmatrix} \right) + g_1(t, 0) \right) (u_{i_1, i_2-1} - u_i) & , \text{ if } i \in \partial_R D \text{ with } i_2 = d_2 - 1; \\
\left(g_1 \left(t, \begin{pmatrix} 0 \\ u_{i_1-1, i_2+1} - u_{i_1-1, i_2-1} \end{pmatrix} \right) \right. \\
+ g_1 \left(t, \begin{pmatrix} u_i - u_{i_1-2 \vee 0, i_2} \\ u_{i_1-1, i_2+1} - u_{i_1-1, i_2-1} \end{pmatrix} \right) \left. \right) (u_{i_1-1, i_2} - u_i) \\
+ g_1 \left(t, \begin{pmatrix} 0 \\ u_{i_1, i_2+2 \wedge d_2-1} - u_i \end{pmatrix} \right) (u_{i_1, i_2+1} - u_i) \\
- g_1 \left(t, \begin{pmatrix} 0 \\ u_i - u_{i_1, i_2-2 \vee 0} \end{pmatrix} \right) (u_i - u_{i_1, i_2-1}) & , \text{ if } i \in \partial_R D \text{ with } i_2 < d_2 - 1; \\
\left(g_1 \left(t, \begin{pmatrix} u_i - u_{i_1-2 \vee 0, i_2} \\ 0 \end{pmatrix} \right) + g_1(t, 0) \right) (u_{i_1-1, i_2} - u_i) \\
+ \left(g_1 \left(t, \begin{pmatrix} 0 \\ u_{i_1, i_2+2 \wedge d_2-1} - u_i \end{pmatrix} \right) + g_1(t, 0) \right) (u_{i_1, i_2+1} - u_i) & , \text{ if } i \in \partial_B D \text{ with } i_1 = d_1 - 1; \\
g_1 \left(t, \begin{pmatrix} u_{i_1+2 \wedge d_1-1, i_2} - u_i \\ 0 \end{pmatrix} \right) (u_{i_1+1, i_2} - u_i) \\
- g_1 \left(t, \begin{pmatrix} u_i - u_{i_1-2 \vee 0, i_2} \\ 0 \end{pmatrix} \right) (u_i - u_{i_1-1, i_2}) \\
+ \left(g_1 \left(t, \begin{pmatrix} u_{i_1+1, i_2+1} - u_{i_1-1, i_2+1} \\ u_{i_1, i_2+2 \wedge d_2-1} - u_i \end{pmatrix} \right) \right. \\
+ g_1 \left(t, \begin{pmatrix} u_{i_1+1, i_2+1} - u_{i_1-1, i_2+1} \\ 0 \end{pmatrix} \right) \left. \right) (u_{i_1, i_2+1} - u_i) & , \text{ if } i \in \partial_B D \text{ with } i_1 < d_1 - 1
\end{cases} \quad (36)$$

1057 for $(t, u) \in \bar{I} \times \mathbb{R}^D$ and $i \in D$ and the discretized diffusion coefficient being given by
1058

$$(\tilde{\sigma}(t, u)v)_i := \begin{cases}
g_2 \left(t, \begin{pmatrix} u_{i_1+1, i_2} - u_{i_1-1, i_2} \\ u_{i_1, i_2+1} - u_{i_1, i_2-1} \end{pmatrix} \right) & , \text{ if } i \in D^\circ; \\
g_2 \left(t, \begin{pmatrix} 0 \\ u_{i_1, i_2+1} - u_{i_1, i_2-1} \end{pmatrix} \right) & , \text{ if } i \in \partial_L D \text{ with } i_2 > 0 \text{ or} \\
& i \in \partial_R D \text{ with } i_2 < d_2 - 1; \\
g_2 \left(t, \begin{pmatrix} u_{i_1+1, i_2} - u_{i_1-1, i_2} \\ 0 \end{pmatrix} \right) & , \text{ if } i \in \partial_T D \text{ with } i_1 > 0 \text{ or} \\
& i \in \partial_B D \text{ with } i_1 < d_1 - 1; \\
g_2(t, 0) & , \text{ if } i \in \partial_L D \text{ with } i_2 = 0 \text{ or} \\
& i \in \partial_T D \text{ with } i_1 = 0 \text{ or} \\
& i \in \partial_R D \text{ with } i_2 = d_2 - 1 \text{ or} \\
& i \in \partial_B D \text{ with } i_1 = d_1 - 1
\end{cases} \quad v_i \quad (37)$$

1071 for $v \in \mathbb{R}^D$, $(t, u) \in \bar{I} \times \mathbb{R}^D$ and $i \in D$.
1072

1073 H.3 TEMPORAL DISCRETIZATION

1074 For temporal discretization, we used a *drift-termed* (explicit) Euler-Maruyama scheme (Hutzenthaler & Jentzen,
1075 2015), where given a generic SDE of the form

$$d\tilde{U}_t = \tilde{b}(t, \tilde{U}_t) dt + \tilde{\sigma}(t, \tilde{U}_t) d\tilde{W}_t \quad \text{for all } t \in \bar{I}, \quad (38)$$

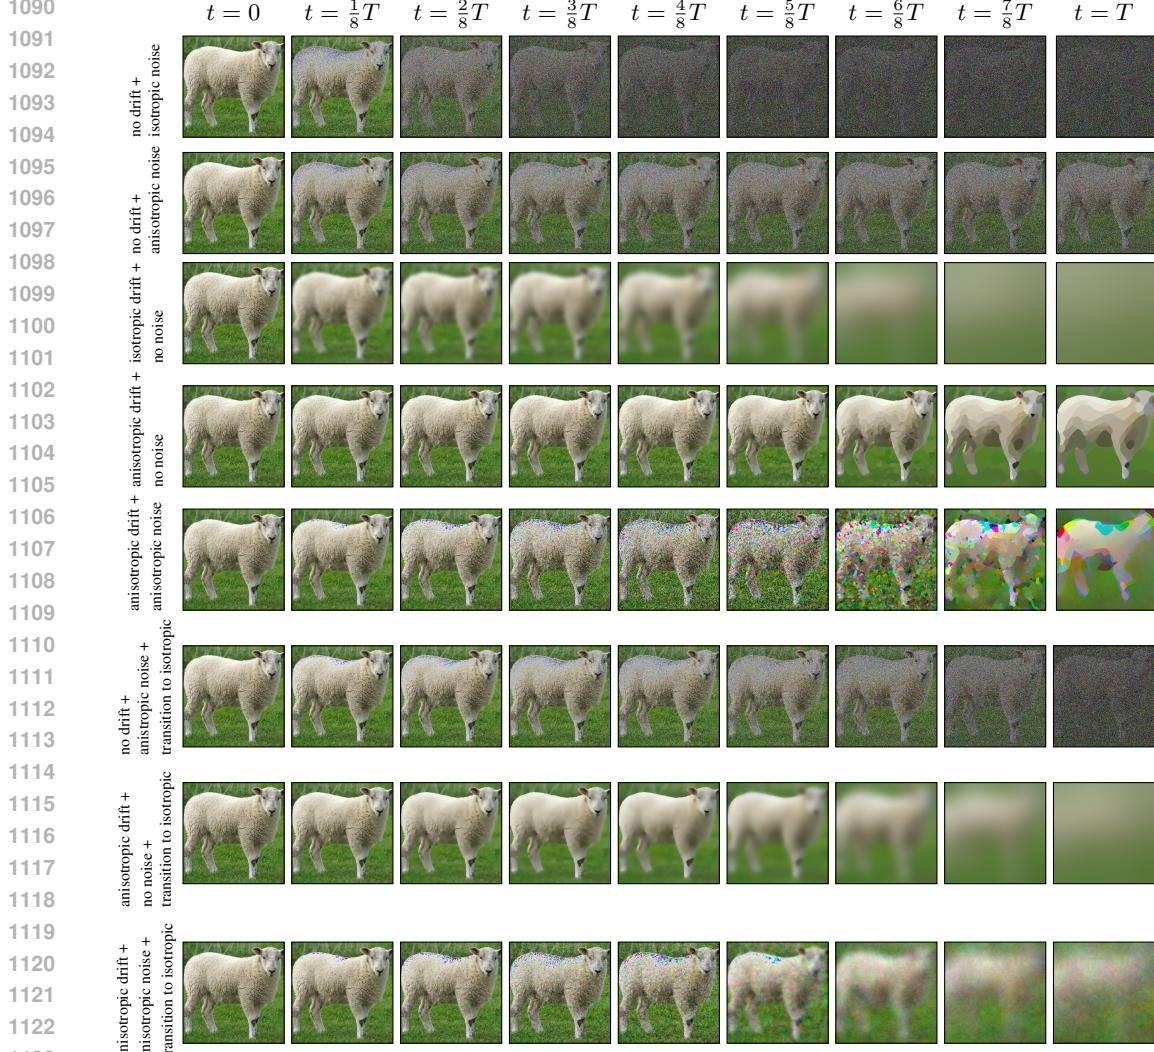
1076
1077
1078
1079

1080 the time stepping is given by
 1081

$$\tilde{U}_{t+\Delta t} = \tilde{U}_t + \frac{b(t, \tilde{U}_t)}{1 + \Delta t \|b(t, \tilde{U}_t)\|^\gamma} + \tilde{\sigma}(t, \tilde{U}_t)(\tilde{W}_{t+\Delta t} - \tilde{W}_t) \quad (39)$$

1085 for all $t, \Delta t \geq 0$ with $t + \Delta t \in \bar{I}$, where γ is a *taming coefficient* usually chosen to be 1.
 1086

1087 I SPDE TRAJECTORY VISUALIZATIONS



1124 Figure 9: We visualize the core ingredients of our generalized anisotropic SPDE diffusion framework. The diffusion process is governed by two fundamental components: the drift term, driven by the drift coefficient b and
 1125 the diffusion term, driven by the diffusion coefficient σ . Both terms can take on isotropic or anisotropic forms,
 1126 and their combinations open the door to a vast spectrum of processes. These processes destroy (and regenerate,
 1127 for the reverse process) the signal’s information in ways that range from subtle to profoundly distinct. The
 1128 interplay between these terms offers the designer of the generative process fine control over how information is
 1129 destroyed.
 1130
 1131
 1132
 1133

J UNCURATED GENERATED SAMPLES ON CIFAR-10

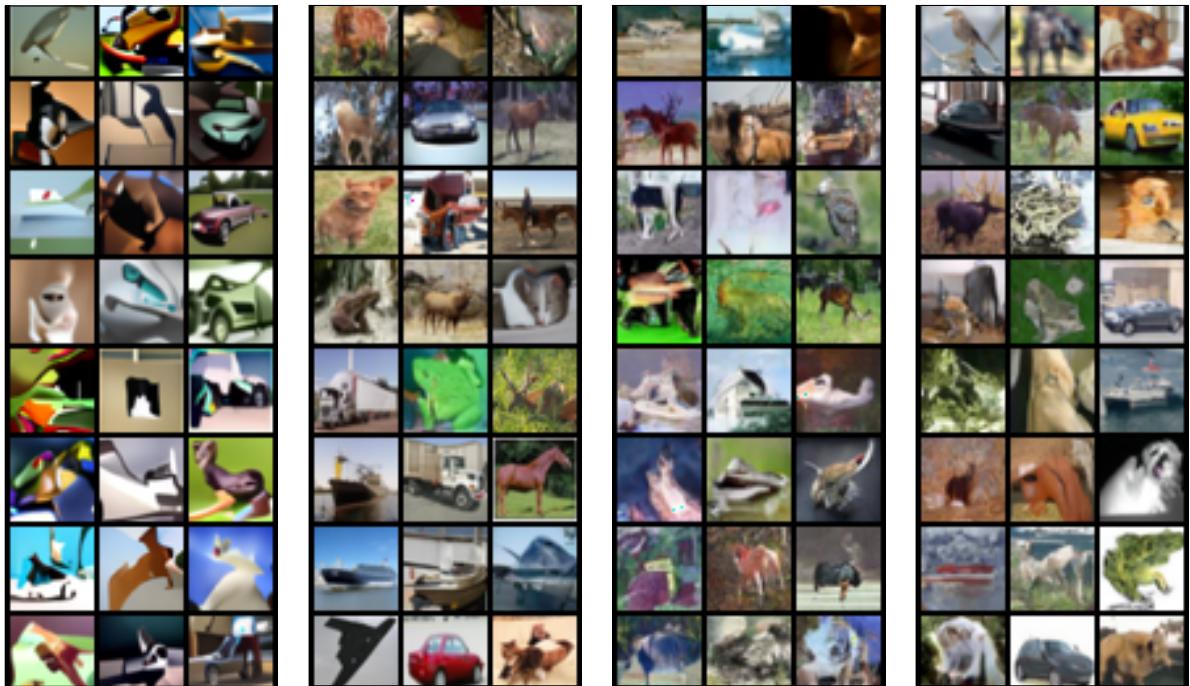


Figure 10: Uncurated samples for the baselines Rissanen et al. (2023), Song et al. (2021c) and two of our SPDEs: one with isotropic drift b and isotropic diffusion coefficient σ (see Section 5.3) and one with anisotropic b and isotropic σ (see Section 5.2). Note that Rissanen et al. (2023) and *Ours (isotropic)* essentially represent the same SPDE, but our score-based approach performs significantly better.

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
K UNCURATED GENERATED SAMPLES ON CELEBA

Figure 11: Uncurated samples for *Ours (anisotropic)* (see Section 5.2). The generated images are produced by a model trained from scratch — without initialization from a pre-trained network — for 100,000 iterations.

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295



Figure 12: Uncurated samples for Song et al. (2021c). The generated images are produced from the checkpoint provided by the authors.

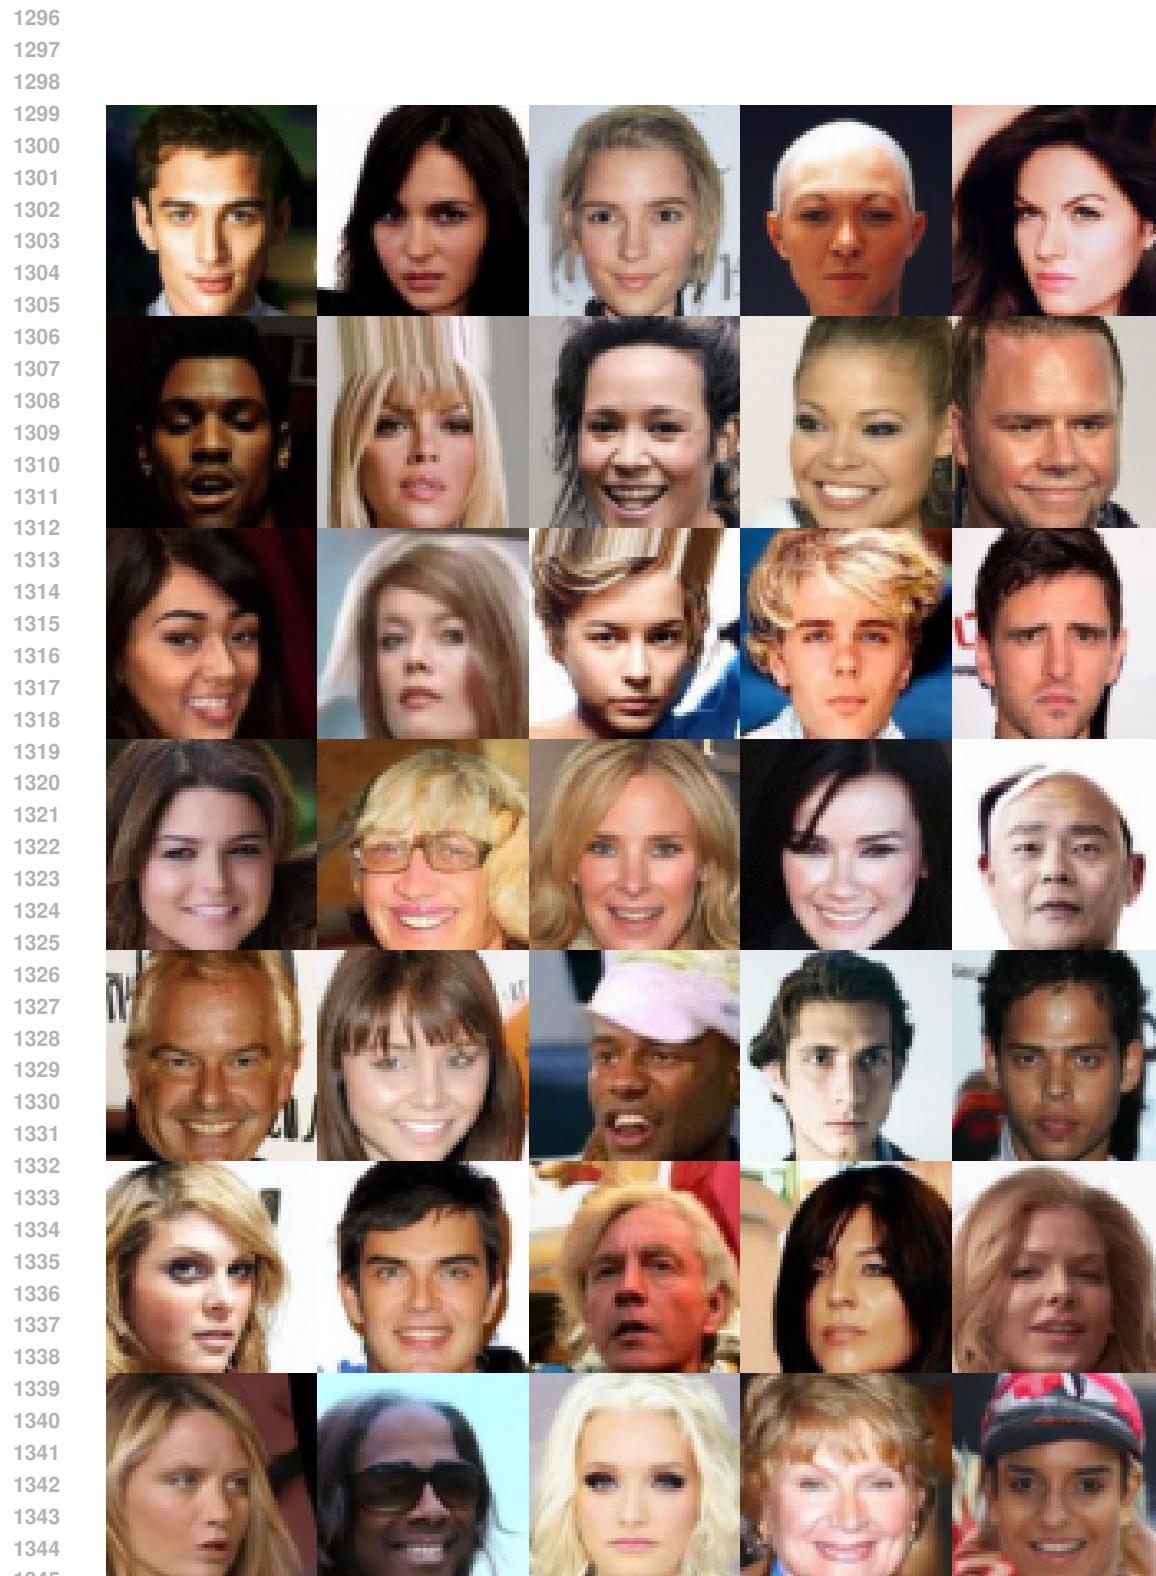


Figure 13: Uncurated samples for Lipman et al. (2022). The generated images are produced by a model trained from scratch — without initialization from a pre-trained network — for 100,000 iterations.

1348
1349

1350

1351

L UNCURATED GENERATED SAMPLES ON IMAGENET2012



Figure 14: Uncurated samples for *Ours (anisotropic)* (see Section 5.2). The generated images are produced by a model trained from scratch — without initialization from a pre-trained network — for 100,000 iterations.

1401
1402
1403

1512 M UNCURATED GENERATED SAMPLES ON LSUN/CHURCH_OUTDOOR
15131514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
Figure 17: Uncurated samples for *Ours (anisotropic)* (see Section 5.2). The generated images are produced by a model trained from scratch — without initialization from a pre-trained network — for 100,000 iterations.

1566

1567

1568

1569

1570

1571

1572

1573

1574

1575

1576

1577

1578

1579

1580

1581

1582

1583

1584

1585

1586

1587

1588

1589

1590

1591

1592

1593

1594

1595

1596

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619



Figure 18: Uncurated samples for Song et al. (2021c). The generated images are produced from the checkpoint provided by the authors.

1620
1621

N COMPUTATIONAL COSTS

1622 To contextualize our computational costs relative to Song et al. (2021c), we report the normalized training
1623 times per 10k steps on CelebA (64×64):1624
1625
1626
1627

Model	Time / 10k steps (s)
Song et al. (2021c)	2506.55
Ours (anisotropic)	4658.63

1628

Additionally, the following inference costs arise during sample generation:

1629
1630
1631
1632

Model	Time / generated image (s)
Song et al. (2021c)	0.5726
Ours (anisotropic)	0.1649

1633
1634
1635

As these timings indicate, our method achieves significantly improved inference times compared to Song et al. (2021c), due to our numerical implementation explained in Appendix H — despite having a theoretically more demanding drift and diffusion coefficient.

1636

O LATENT SPACE EXPERIMENT

1637
1638
1639
1640
1641We conducted a latent space experiment using *Ours (anisotropic)* on the LSUN/church_outdoor dataset. For the latent representation, we employed the pretrained variational autoencoder from stabilityai/sd-vae-ft-mse. The resulting generative performance metrics are:1642
1643
1644
1645

$$\begin{aligned} \text{Inception Score} &= 3.880973 \\ \text{FID} &= 3.936322 \\ \text{KID} &= 2.387085 \times 10^{-3} \end{aligned}$$

1646
1647

These results demonstrate that our anisotropic SPDE framework can be successfully applied in latent space as well.

1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

1674
 1675 **P UNCURATED GENERATED SAMPLES ON LSUN/CHURCH_OUTDOOR FROM**
 1676 **OUR LATENT SPACE EXPERIMENT**



1725 Figure 19: Uncurated samples for *Ours (anisotropic)* (see Section 5.2) from our latent space experiment. The
 1726 generated images are produced by a model trained from scratch — without initialization from a pre-trained
 1727 network — for 100,000 iterations.