On the Theoretical Foundations of Score-Based Generative Models

Stochastic Control, Discrete Extensions, and Stability Guarantees

Antonio Ocello

March 3, 2026 — Séminaire Palaisien

If You Ever Get Bored During This Talk. . . Just Think About This

If You Ever Get Bored During This Talk... Just Think About This

Figure: AI-generated image from DALL-E 2. Image: OpenAI.

Generative Models

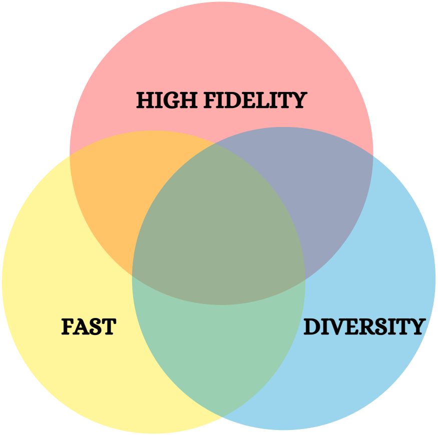

Generative Models

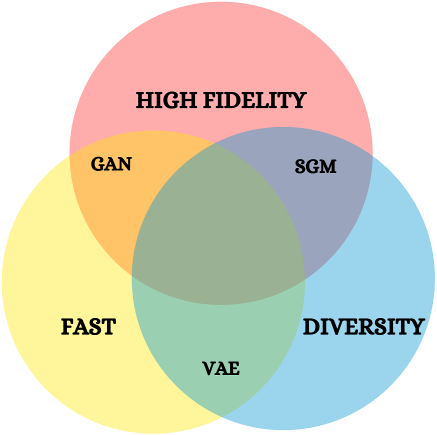

Generative Models

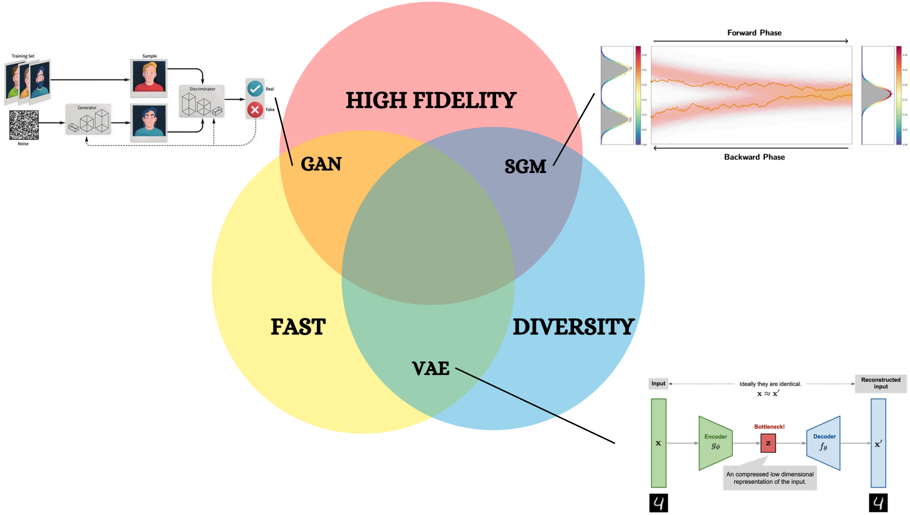

Score-Based Models: Natural Questions

A natural question arises:

Why do these models seem to work better?

A natural question arises:

Why do these models seem to work better?

1. **Why do these models work?**

Score-Based Models: Natural Questions

A natural question arises:

Why do these models seem to work better?

1. **Why do these models work?**
2. **Is the pipeline general?**

A natural question arises:

Why do these models seem to work better?

1. **Why do these models work?**
2. **Is the pipeline general?**
3. **Stochastic vs deterministic generation?**

Principles of Score-based Generative Models (SGMs)

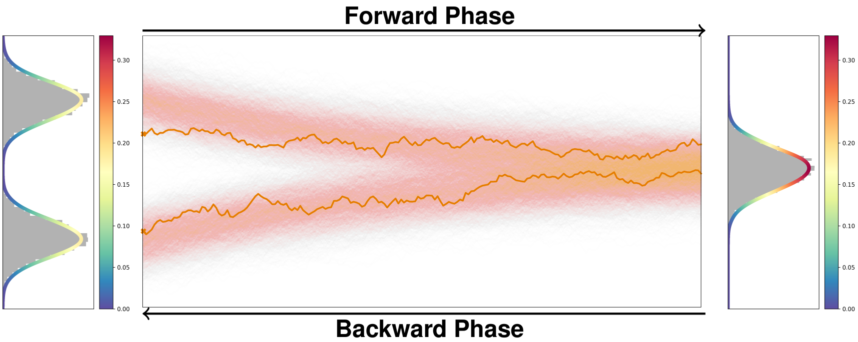

Interpolating between two distributions:

- The **data distribution**: $\pi_{\text{data}} \in \mathcal{P}(\mathbb{R}^d)$.
- The **easy-to-sample distribution**: $\pi_{\text{ref}} \in \mathcal{P}(\mathbb{R}^d)$ (typically, a *standard multivariate Gaussian*).

Principles of Score-based Generative Models (SGMs)

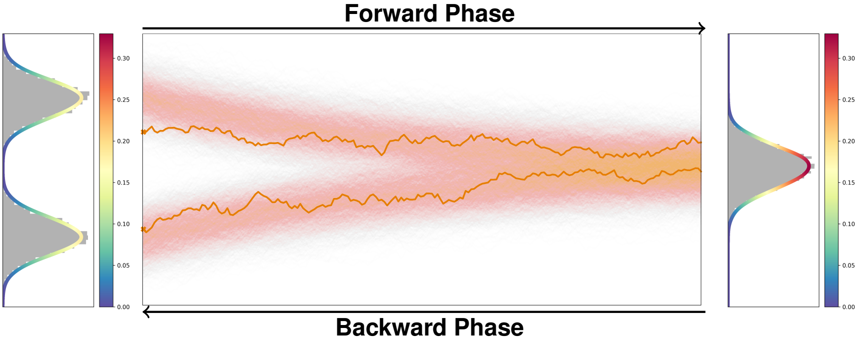

Two fundamental processes in SGMs:

- **Forward process**: Transforms π_{data} into π_{ref} by gradually adding noise.
- **Backward process**: Inverts the noising process to reconstruct samples from π_{data}.

Forward and Backward Process in SGMs

Continuous-Time Formulation of SGMs: trasform π_{data} into an easy-to-sample distribution π_{ref} through an **SDE**.

Forward and Backward Process in SGMs

Continuous-Time Formulation of SGMs: trasform π_{data} into an easy-to-sample distribution π_{ref} through an **SDE**.

Forward Process: solution of an **Ornstein-Uhlenbeck Process**

$$d\overrightarrow{X}_t = -\overrightarrow{X}_t dt + \sqrt{2} dW_t, \quad \overrightarrow{X}_0 \sim \pi_{\text{data}},$$

- The drift term $-X_t$ ensures that the process *converges* to $\mathcal{N}(0, \mathbb{I})$ as $t \to \infty$.
- The diffusion term $\sqrt{2} dW_t$ introduces controlled noise, *smoothing the transition*.
- *Notation:* $p_t(x)\mathrm{d}x := \mathcal{L}(\overrightarrow{X}_t)$.

Forward and Backward Process in SGMs

Continuous-Time Formulation of SGMs: trasform π_{data} into an easy-to-sample distribution π_{ref} through an **SDE**.

Forward Process: $d\overrightarrow{X}_t = -\overrightarrow{X}_t dt + \sqrt{2}dW_t$, $\overrightarrow{X}_0 \sim \pi_{\text{data}}$, with $p_t(x)\mathrm{d}x := \mathcal{L}(\overrightarrow{X}_t)$.

Backward Process (Time-Reversed SDE): Under mild conditions the forward process admits a **time-reversed process** [And82; HP86; CCGL23], *i.e.*,

$$\mathcal{L}\left(\left(\overleftarrow{X}_t\right)_{t\in[0,T]}\right) = \mathcal{L}\left(\left(\overrightarrow{X}_{T-t}\right)_{t\in[0,T]}\right),$$

$$\text{with } d\overleftarrow{X}_t = \left[\overleftarrow{X}_t + 2\underbrace{\nabla \log p_{T-t}}_{\text{score function}}\left(\overleftarrow{X}_t\right)\right] dt + \sqrt{2}d\bar{W}_t, \quad \overleftarrow{X}_0 \sim p_T \approx \mathcal{N}(0, \mathbb{I}).$$

- The **score term** drives the equation in **regions of space of high probability**.
- **Backward process** is a generative model in SGMs.

SGMs in Practice I: Mixing Time

- **Problem 1:** Time-reversal holds when $\overleftarrow{X}_0 \sim p_T$.
- However, p_T depends π_{data}

$$p_t(x_t) = \int_{\mathbb{R}^d} \underbrace{p_t(x_t|x_0)}_{\mathcal{L}\left(\overrightarrow{X}_t \middle| \overrightarrow{X}_0\right)} \pi_{\mathrm{data}}(x_0)\mathrm{d}x_0 .$$

- Leveraging the ergodicity of the Ornstein-Uhlenbeck kernel:

$$X_t \overset{\mathcal{L}}{=} \mathrm{e}^{-t} X_0 + \sqrt{1 - \mathrm{e}^{-2t}} Z \,, \qquad \text{with } Z \sim N(0, \mathbb{I}), Z \perp X_0 \,.$$

- For large T enough, $p_T \approx \pi_{\mathrm{ref}} \sim \mathcal{N}(0, \mathbb{I})$.

Mixing time error: p_T is not exactly π_{ref}.

SGMs in Practice II: Learning the Score Function

- **Problem 2:** the backward process depends on the score function $\nabla \log p_{T-t}(x)$.
- **Idea:** Train a **deep neural network** $s_\theta : [0, T] \times \mathbb{R}^d \to \mathbb{R}^d$ to minimize

$$\mathcal{L}_{\text{explicit}}(\theta) = \mathbb{E}\left[\left\| s_\theta\left(\tau, \overrightarrow{X}_\tau\right) - \nabla \log p_\tau\left(\overrightarrow{X}_\tau\right)\right\|^2\right], \quad \text{with } \tau \sim \mathcal{U}(0, T), \tau \perp \overrightarrow{X}.$$

- Since $p_\tau(x)$ is unknown, use a **conditional score-matching objective** [HD05; Vin11]

$$\mathcal{L}_{\text{score}}(\theta) = \mathbb{E}\left[\left\| s_\theta\left(\tau, \overrightarrow{X}_\tau\right) - \nabla \log p_\tau\left(\overrightarrow{X}_\tau \middle| \overrightarrow{X}_0\right)\right\|^2\right],$$

as the conditional score is the score of a Gaussian kernel that writes as

$$\nabla \log p_\tau\left(\overrightarrow{X}_\tau \middle| \overrightarrow{X}_0\right) = \frac{\mathrm{e}^{-\tau} X_0 - \overrightarrow{X}_\tau}{1 - \mathrm{e}^{-2\tau}}.$$

Approximation error: s_θ is not exactly $\nabla \log p_{T-t}$.

SGMs in Practice III: Simulating from the Backward Kernel

- **Problem 3:** unlike the forward process, the backward process is **non-linear**.
- Consider a partition $0 = t_0 \leq t_1 \leq \cdots \leq t_N = T$ of $[0, T]$, and denote $h_k = t_{k+1} - t_k$.
- Use an **Euler-Maruyama scheme** to simulate the backward process

$$d\overleftarrow{X}_t^N = \left(\overleftarrow{X}_t^N + 2s_\theta\left(T - t_k, \overleftarrow{X}_{t_k}^N\right)\right) dt + \sqrt{2}dB_t.$$

Discretization error: Euler-Maruyama scheme is not exact.

SGM Pipeline: Training vs Generation

Training

Inputs: dataset $\{x^{(i)}\}$ (i.i.d. $\sim \pi_{\mathrm{data}}$), horizon T, partition $0 = t_0 \leq \cdots \leq t_N = T$.
Initialize network parameters θ.
for each training iteration **do**
  Sample a mini-batch $x_0 \sim \pi_{\mathrm{data}}$.
  Sample $\tau \sim \mathcal{U}(0, T), z \sim \mathcal{N}(0, I)$.
  Simulate x_τ: $\quad x_\tau = \mathrm{e}^{-\tau} x_0 + \sqrt{1 - \mathrm{e}^{-2\tau}}\, z$.
  Compute the **conditional score** target:

$$\nabla \log p_\tau(x_\tau \mid x_0) = \frac{\mathrm{e}^{-\tau} x_0 - x_\tau}{1 - \mathrm{e}^{-2\tau}}.$$

  Update θ by SGD on

$$\mathcal{L}(\theta) = \mathbb{E}\left[\| s_\theta(\tau, x_\tau) - \nabla \log p_\tau(x_\tau \mid x_0)\|^2\right].$$

end for
Output: trained score network s_θ.

Generation

Inputs: trained s_θ, T, partition $0 = t_0 \leq \cdots \leq t_N = T$.
Initialize $X_0^\star \sim \pi_{\mathrm{ref}} = \mathcal{N}(0, I)$.
for $k \in \{0, \ldots, N-1\}$ **do**
  Sample $Z_k \sim \mathcal{N}(0, I)$.
  Euler–Maruyama update:

$$X^\star_{t_{k+1}} = X^\star_{t_k} + (t_{k+1} - t_k)(X^\star_{t_k} + 2 s_\theta(T - t_k, X^\star_{t_k})) + \sqrt{2(t_{k+1} - t_k)}\, Z_k.$$

end for
Output: sample $X_T^\star \approx \pi_{\mathrm{data}}$.

Time Reversal as an Optimal Control Problem

Theorem (Time Reversal under Finite Entropy [CCGL23])

Let $(X_t)_{t\in[0,T]}$ solve a stationary ergodic SDE

$$dX_t = b(X_t)\,dt + \sigma\,dB_t, \qquad X_0 \sim \mu_0, \qquad \text{with invariant measure } \mu_\star.$$

Assume the **finite entropy condition:** $\mathrm{KL}(\mu_0 \mid \mu_\star) < \infty$. Denote p_t denotes the density of X_t. Then, the time-reversed process $(\overleftarrow{X}_t)_{t\in[0,T]}$ satisfies

$$d\overleftarrow{X}_t = \left(b(\overleftarrow{X}_t) + \sigma^2 \nabla \log p_{T-t}(\overleftarrow{X}_t)\right) dt + \sigma\,dB_t,$$

and the reversed drift is the unique minimizer of the following **stochastic control problem**:

$$\inf_u \mathbb{E}\left[\frac{1}{2}\int_0^T \|u_t\|^2 dt + \log \frac{d\mu_\star}{d\mu_0}(X_T^u)\right]$$
$$\text{s.t.} \quad dX_t^u = b(X_t^u)dt + u_t dt + \sigma dB_t, \quad X_0^u \sim p_T.$$

Log Density Solves an HJB Equation

Let $\tilde{p}_t = \frac{p_t}{\mu_\star}$ denote the Radon-Nykodum density of p_t with respect to the invariant measure $\mu_\star$. Then, the log density $\Phi_t = \log \tilde{p}_t$ solves a **Hamilton–Jacobi–Bellman (HJB) equation**:

$$\partial_t \Phi + b^\top \Phi + \frac{\sigma^2}{2} \Delta^2 \Phi + \frac{\sigma^2}{2} \|\nabla \Phi\|^2 = 0.$$

Thus:

- The **score** is the **optimal control**, i.e., $u_t^\star = \nabla \Phi_{T-t} = \nabla \log \tilde{p}_{T-t}$.
- The **backward SDE** is the **controlled diffusion**.
- Fisher information $\mathcal{I}(\mu_0 \mid \mu_\star) = \int \left\| \nabla \log \left(\frac{\mathrm{d}p_0}{\mathrm{d}\mu_\star} \right) \right\|^2 \mathrm{d}p_0$ corresponds to control energy.

KL Divergence as Control Energy

Girsanov identity: $\mathrm{KL}(\mu_0 \mid p_T^\theta) = \frac{1}{2}\mathbb{E}\int_0^T \|u_t^\star - \hat{u}_t\|^2 dt$ + initialization term.

Theorem (KL Convergence Bound — the OU case [CDS25])

Assume there exist $\varepsilon^2 > 0$ and $\theta^\star \in \Theta$ such that

$$\frac{1}{T}\sum_{k=0}^{N-1} h_{k+1}\,\mathbb{E}\left[\left\|\tilde{s}_{\theta^\star}(T-t_k, \overrightarrow{X}_{T-t_k}) - 2\nabla \log \tilde{p}_{T-t_k}(\overrightarrow{X}_{T-t_k})\right\|^2\right] \le \varepsilon^2.$$

Then, $\mathrm{KL} = E_1 + E_2 + E_3$ where

$E_1 := e^{-2T}\mathrm{KL}(\mu_0|\mu_\star)$ **(initialization)**

$E_2 := T\varepsilon^2$ **(score training error)**

$E_3 := h\,\mathcal{I}(\mu_0|\mu_\star)$ **(Euler discretization error)**.

Key insight: Euler error is controlled directly by Fisher information.

Beyond Log-Concavity and Score Regularity: Improved Convergence Bounds for Score-Based Generative Models in W2-distance

Marta Gentiloni-Silveri[1], Antonio Ocello[2]

[1]CMAP, École Polytechnique, Institut Polytechnique de Paris, France.
[2]CREST, Groupe ENSAE-ENSAI, ENSAE Paris, Institut Polytechnique de Paris, France.

The Forty-Second International Conference on Machine Learning (ICML).

Evolution of Concavity and Smoothness in Time

Propagation of Strong Concavity and Lipschitzness

Assume that π_{data} **is C_0-strongly log-concave** and $\nabla \log \pi_{\mathrm{data}}$ **is L_0-smooth** with $\min\{C_0, L_0\} > 1/2$. Then, for any $t \in (0, T]$, the score function $p_t(x)$ **is C_t-strongly log-concave** with

$$C_t = \frac{1}{\mathrm{e}^{-t}/C_0 + 2\left(1 - \mathrm{e}^{-2t}\right)} \,,$$

and $\nabla \log p_t$ **is L_t-Lipschitz** with

$$L_t = \min\left\{\frac{1}{2\left(1 - \mathrm{e}^{-2t}\right)}; L_0 \mathrm{e}^{2t}\right\} \,.$$

Moreover, $C_t \leq C_0$ and $L_t \leq L_0$, for $t \geq 0$.

Key insight: The first result is a consequence of the **Prékopa-Leindler inequality**, a fundamental result in convex analysis [Pré71; BL76; SW14].

Beyond Concavity: Weak Concavity

Convexity Profile: For a differentiable vector field γ, its *weak convexity profile* is

$$\kappa_\gamma(r) = \inf_{x,y\in\mathbb{R}^d : \|x-y\|=r} \left\{ \frac{(\nabla\gamma(x) - \nabla\gamma(y))^\top (x-y)}{\|x-y\|^2} \right\} .$$

Weak convexity goes beyond $\kappa_\gamma \geq 0$ by allowing non-uniform bounds.

Definition (Weak Concavity):

A vector field γ is **weakly convex** if there exist $\alpha, M > 0$ such that its convexity profile satisfies

$$\kappa_\gamma(r) \geq \alpha - \frac{1}{r} f_M(r) ,$$

with $f_M(r) := 2\sqrt{M}\tanh\left(r\sqrt{M}/2\right)$. A vector field γ is **weakly concave** if $-\gamma$ is weakly convex.

New Model Assumptions

Assumption 1:

The data distribution π_{data} satisfies $\pi_{\mathrm{data}}(\mathrm{d}x) = \exp(-U(x))\mathrm{d}x$, with

1. **(Weak Convexity)** U is weakly convex, i.e., its profile κ_U satisfies

$$\kappa_U(r) \geq \alpha - \frac{1}{r} f_M(r), \quad \text{for some } \alpha, M > 0.$$

2. **(Lipschitz Condition)** ∇U is L_U-Lipschitz.

Assumption 2:

There exist $\varepsilon \geq 0$ and $\theta^\star \in \Theta$ such that

$$\sup_{k \in [N]} \left\| \nabla \log p_{T-t_k}\left(X_{t_k}^{\infty}\right) - s_\theta\left(T - t_k, X_{t_k}^{\infty}\right) \right\|_{L^2} \leq \varepsilon.$$

Example

Which kind of distributions satisfy Assumption 1?

Example

Which kind of distributions satisfy Assumption 1?

Proposition [GSO25]:

Let p_n be a Gaussian mixture in $\mathbb{R}^d$, *i.e.*,

$$p_n(x) = \sum_{i=1}^{n} \beta_i \frac{1}{(2\pi\sigma^2)^{d/2}} \exp\left(-\frac{|x-\mu_i|^2}{2\sigma^2}\right) ,$$

with $\sigma > 0$, $\mu_i \in \mathbb{R}^d$ and $\beta_i \in [0,1]$, such that $\sum_{i=1}^n \beta_i = 1$. Then, $-\log p_n$ is **weakly convex** with coefficients

$$\alpha_{p_n} = \frac{1}{\sigma^2} , \qquad \sqrt{M_{p_n}} = \frac{2n}{\sigma^2} \sum_{i=1}^{n} \|\mu_i\| .$$

Moreover, we have that $\nabla \log p_n$ is **$(\beta_{p_n} + \sqrt{M_{p_n}})$-Lipschitz**, with $\beta_{p_n} = \frac{1}{\sigma^2}$.

Key Results: Convexity of the Score Function

Proposition:

Assume that U is **(α, M)-weakly convex**. Then, the function $x \mapsto -\log \tilde{p}_{T-t}(x)$ is weakly convex with weak convexity profile:

$$\kappa_{-\log \tilde{p}_{T-t}}(r) \geq \frac{\alpha}{\alpha + (1-\alpha)\mathrm{e}^{-2(T-t)}} - 1 - \frac{\mathrm{e}^{-(T-t)}}{\alpha + (1-\alpha)\mathrm{e}^{-2(T-t)}} \frac{1}{r} f_M\left(\frac{\mathrm{e}^{-(T-t)}}{\alpha + (1-\alpha)\mathrm{e}^{-2(T-t)}} r\right).$$

Key Results: Convexity of the Score Function

Proposition:

Assume that U is (α, M)**-weakly convex**. Then, the function $x \mapsto -\log \tilde{p}_{T-t}(x)$ is weakly convex with weak convexity profile:

$$\kappa_{-\log \tilde{p}_{T-t}}(r) \geq \frac{\alpha}{\alpha + (1-\alpha)\mathrm{e}^{-2(T-t)}} - 1 - \frac{\mathrm{e}^{-(T-t)}}{\alpha + (1-\alpha)\mathrm{e}^{-2(T-t)}} \frac{1}{r} f_M\left(\frac{\mathrm{e}^{-(T-t)}}{\alpha + (1-\alpha)\mathrm{e}^{-2(T-t)}} r\right).$$

In particular, the modified score function satisfies

$$\left(\nabla \log \tilde{p}_{T-t}(x) - \nabla \log \tilde{p}_{T-t}(y)\right)^\top (x-y) \leq -C_t \|x-y\|^2, \quad \forall x, y \in \mathbb{R}^d,$$

with $C_t = \dfrac{\alpha}{\alpha + (1-\alpha)\mathrm{e}^{-2(T-t)}} - \dfrac{\mathrm{e}^{-2(T-t)}}{(\alpha + (1-\alpha)\mathrm{e}^{-2(T-t)})^2} M - 1.$

Key Results: Smoothness of the Score Function

Proposition:

Assume that ∇U is L_U**-Lipschitz**. Then, the modified score function is L_t-Lipschitz:

$$\|\nabla \log \tilde{p}_{T-t}(x) - \nabla \log \tilde{p}_{T-t}(y)\| \leq L_t \|x - y\| , \quad \forall x, y \in \mathbb{R}^d.$$

with

$$L_t = \min\left\{\frac{1}{1 - \mathrm{e}^{-2(T-t)}} ; \mathrm{e}^{2(T-t)} L_U\right\} + 1.$$

Regime switching

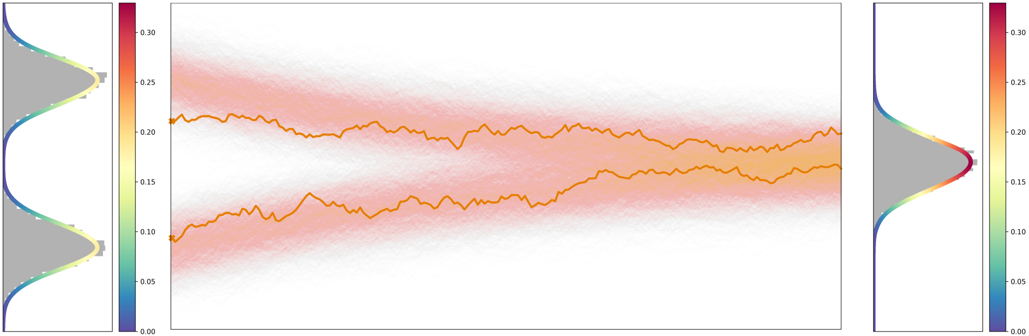

Regime switching

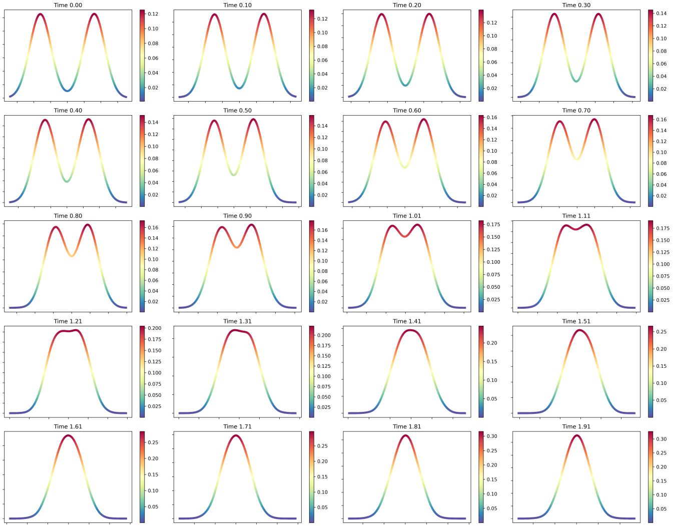

Regime switching

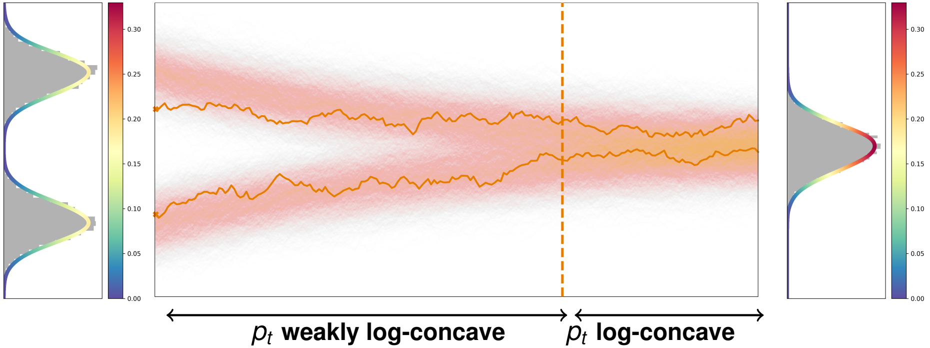

Regime switching

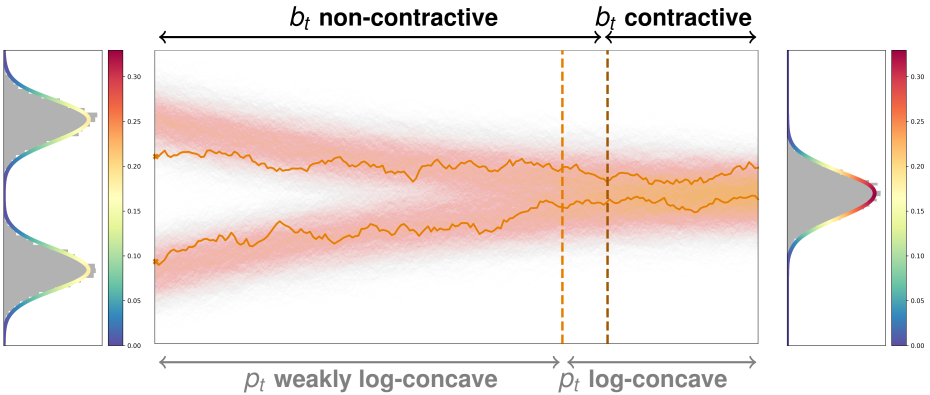

Regime switching

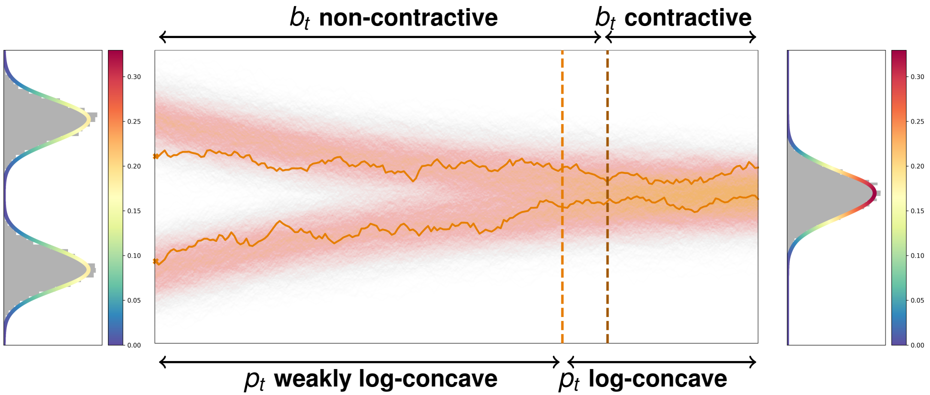

Key Insight: Contractivity foster efficiency.

Non-Asymptotic Guarantees for SGMs

Theorem [GSO25]:

Consider a uniform grid $\{kh\}_k$ for discretizing time. Suppose, Assumptions 1 and 2 hold. Then, there exists $C > 0$ such that

$$\mathcal{W}_2\left(\pi_{\mathrm{data}}, \mathcal{L}\left(X_T^\star\right)\right) \lesssim \underbrace{\mathrm{e}^{-T}\mathcal{W}_2\left(\pi_{\mathrm{data}}, \pi_{\mathrm{ref}}\right)}_{\text{Mixing time error}} + \underbrace{T\varepsilon}_{\text{Approx. error}} + \underbrace{\sqrt{dh}T}_{\text{Discr. error}} .$$

Discrete Markov Probabilistic Models

Le-Tuyet-Nhi PHAM[1], Dario Shariatian[2], Antonio Ocello[3], Giovanni Conforti[4], Alain Oliviero Durmus[1]

[1]CMAP, École Polytechnique, Institut Polytechnique de Paris, France.
[2]INRIA - Department of Computer Science, PSL Research University, Paris, France.
[3]CREST, Groupe ENSAE-ENSAI, ENSAE Paris, Institut Polytechnique de Paris, France.
[4]Dipartimento di Matematica, Università degli Studi di Padova, Padova, Italy.





The Forty-Second International Conference on Machine Learning (ICML).

CTMCs on Discrete Spaces: Generator and Sampling

A CTMC $(X_t)_{t\in[0,T]}$ on a discrete space $\mathbf{X}$ is characterized by a (possibly time-inhomogeneous) **rate matrix** $(q_t)_{t\in[0,T]}$ such that

$$\sum_{y\in\mathbf{X}} q_t(x,y) = 0, \qquad \forall x \in \mathbf{X},$$

and, as $h \downarrow 0$,

$$\mathbb{P}(X_{t+h} = y \mid X_t = x) = \delta_x(y) + h\, q_t(x,y) + o(h).$$

Sampling [Gil07]: Define the jump rate and jump kernel

$$\lambda_t(x) = \sum_{y\in\mathbf{X}} q_t(x,y), \qquad k_t(x,y) = \mathbf{1}_{x\neq y}\frac{q_t(x,y)}{\lambda_t(x)}.$$

Jump times are generated via an exponential clock and the integral hazard

$$\Delta T_{i+1} = \inf\left\{t \geq 0 : \int_0^t \lambda_{T_i+r}(X_{T_i})\, dr \geq E_i\right\},$$

then sample the next state using $\mathrm{Cate}(k_{T_{i+1}}(X_{T_i},\cdot))$.

Forward Process on $\{0,1\}$: Closed-Form Transition

State space $\mathbf{X} = \{0,1\}$, homogeneous bit-flip CTMC:

$$\overrightarrow{q}_1(x,y) = \begin{cases} \lambda, & y \neq x, \\ -\lambda, & y = x. \end{cases}$$

Closed form transition matrix. For $0 \leq t \leq T$, $\overrightarrow{p}_t^{\,1} = \exp(t\,\overrightarrow{q}_1)$ and

$$\overrightarrow{p}_t^{\,1}(x,y) = \begin{cases} \frac{1}{2} + \frac{1}{2}e^{-2\lambda t}, & x = y, \\ \frac{1}{2} - \frac{1}{2}e^{-2\lambda t}, & x \neq y. \end{cases}$$

Equivalently,

$$\overrightarrow{p}_t^{\,1} = \frac{1}{2}\begin{pmatrix} 1 + e^{-2\lambda t} & 1 - e^{-2\lambda t} \\ 1 - e^{-2\lambda t} & 1 + e^{-2\lambda t} \end{pmatrix}.$$

Exponential mixing to uniform. As $t \to \infty$, $e^{-2\lambda t} \to 0$ and $\overrightarrow{p}_t^{\,1}(x,\cdot) \to (1/2, 1/2)$.

Backward Process on $\{0,1\}$: Time Reversal Formula

Define the reversed process $\overleftarrow{X}_t = \overrightarrow{X}_{T-t}$. Then it is again a (time-inhomogeneous) CTMC with generator $\overleftarrow{q}_t$ satisfying the **time-reversal identity**:

$$\mu_{T-t}(x)\,\overleftarrow{q}_t(x,y) = \mu_{T-t}(y)\,\overrightarrow{q}_1(y,x), \qquad x \neq y,$$

where $\mu_t(x) = \mathbb{P}(\overrightarrow{X}_t = x)$.

Since $\overrightarrow{q}_1$ is symmetric,

$$\overleftarrow{q}_t(x,y) = \overrightarrow{q}_1(y,x)\frac{\mu_{T-t}(y)}{\mu_{T-t}(x)}.$$

Discrete Score on $\{0,1\}$ and Conditional Expectation

Discrete score function (discrete analogue of $\nabla \log p_t$):

$$s_t(x) = \frac{\mu_{T-t}(x) - \mu_{T-t}(1-x)}{\mu_{T-t}(x)}.$$

Backward generator expressed via the score:

$$\overleftarrow{q}_t(x,y) = \begin{cases} \lambda\big(1 - s_t(x)\big), & y \neq x, \\ -\lambda\big(1 - s_t(x)\big), & y = x. \end{cases}$$

Explicit conditional expectation representation.
Let $\alpha_t = e^{-2\lambda t}$. Then

$$s_t(x) = \mathbb{E}\left[\frac{2\alpha_{T-t}\,\mathbf{1}_{\{\overrightarrow{X}_0 = \overrightarrow{X}_{T-t}\}} - (1+\alpha_{T-t})}{1-\alpha_{T-t}^2} \,\middle|\, \overrightarrow{X}_{T-t} = x \right].$$

Key insight: The discrete score is a regression target obtained from forward trajectories.

Extension to $\{0,1\}^d$: Forward Process and Closed Form

State space $\mathbf{X} = \{0,1\}^d$. The forward CTMC flips one coordinate at each jump: sample jump times from a Poisson process and at each jump choose $\ell_i \sim \mathrm{Unif}\{1,\ldots,d\}$, then flip the ℓ_i-th bit.

The generator decomposes as a sum of coordinate generators:

$$\overrightarrow{q} = \sum_{\ell=1}^{d} \overrightarrow{q}_\ell, \quad \overrightarrow{q}_\ell(x,y) = \begin{cases} \lambda, & y = \phi^{(\ell)}(x) \text{ (flip bit } \ell\text{)}, \\ -\lambda, & y = x, \\ 0, & \text{otherwise.} \end{cases}$$

Closed form / factorization. Denoting $\overrightarrow{p}_t^{\,1}$ the $\{0,1\}$ transition, known in closed form, the transition kernel factorizes over dimensions, yielding

$$\mu_t(x) = \sum_{z\in\mathbf{X}} \mu_0(z) \prod_{i=1}^{d} \overrightarrow{p}_t^{\,1}(z_i, x_i).$$

Exponential mixing to uniform. As each coordinate mixes at rate $e^{-2\lambda t}$, the chain converges exponentially fast to the uniform measure γ_d on $\{0,1\}^d$.

Backward Process on $\{0,1\}^d$ and Vector Score

The time-reversed process $\overleftarrow{X}_t = \overrightarrow{X}_{T-t}$ is a non-homogeneous CTMC with generator $\overleftarrow{q}_t$ satisfying the same time-reversal formula.

Define the **vector discrete score** $s_t(x) = (s_t^{(1)}(x), \dots, s_t^{(d)}(x))$ by

$$s_t^{(\ell)}(x) := \frac{\mu_{T-t}(x) - \mu_{T-t}(\phi^{(\ell)}(x))}{\mu_{T-t}(x)}, \qquad \ell = 1, \dots, d,$$

where $\phi^{(\ell)}$ flips the ℓ-th bit.

Then the backward jump rates only connect Hamming-neighbors:

$$\overleftarrow{q}_t(x,y) = \sum_{\ell=1}^{d} \lambda\big(1 - s_t^{(\ell)}(x)\big)\mathbf{1}_{\{y=\phi^{(\ell)}(x)\}}.$$

Key insight: Each component of the discrete score admits a conditional expectation representation $\Longrightarrow$ learning s_t reduces to a regression problem based on forward trajectories (*discrete analogue of score matching*).

Time Reversal as an Optimal Control Problem

Following [CL22], the time-reversal of the forward CTMC can be interpreted as an **optimal control problem** minimizing the relative entropy between path measures.

The optimal control $u_t(x, y)$ (for $x \neq y$) admits the explicit form

$$u_t(x, y) = \exp\big(V(t, x) - V(t, y)\big), \qquad (t, x, y) \in [0, T] \times \mathbf{X}^2,$$

with $V(t, x) = -\log \tilde{\mu}_{T-t}(x)$, where $\tilde{\mu}_t$ denotes the relative density with respect to the invariant measure.

Proposition (HJB equation [PSOCD25])

The value function V satisfies the discrete Hamilton–Jacobi–Bellman equation:

$$\partial_t V(t, x) = \lambda \sum_{\ell=1}^{d} \left[e^{V(t,x) - V\left(t, \phi^{(\ell)}(x)\right)} - 1 \right],$$

$$V(T, x) = g(x) = -\log \frac{d\mu_0}{d\mu_\star}(x), \qquad (t, x) \in [0, T) \times \mathbf{X}.$$

Theorem (Sharp KL Convergence Bound [PSOCD25])

Consider the discrete forward CTMC on $\{0,1\}^d$ with horizon T, and assume:

(A1) Finite initial relative entropy: $\mathrm{KL}(\mu_0 \mid \mu_\star) < \infty$.

(A2) Score approximation error bounded by ε^2.

Then the learned reverse-time process satisfies

$$\mathrm{KL}(\mu_0 \,\|\, \hat{\mu}_T) \;\le\; E_1 + E_2 + E_3,$$

where

$$E_1 = e^{-2\lambda T}\mathrm{KL}(\mu_0 \mid \mu_\star) \qquad \textbf{(mixing / initialization)},$$

$$E_2 = T\,\varepsilon^2 \qquad \textbf{(score learning error)},$$

$$E_3 = \lambda h \beta_{\gamma_d}(\mu_\star), \qquad \textbf{(discretisation / sampling error)},$$

with $\beta_{\gamma_d}(\mu_0)$ a Fisher-like information of μ_0 relative to the uniform distribution γ_d, i.e.,

$$\beta_{\gamma_d}(\mu_0) = \mathbb{E}\left[\sum_{\ell=1}^{d} \mathfrak{h}\left(e^{g(\overrightarrow{X}_0) - g(\phi^{(\ell)}(\overrightarrow{X}_0))}\right)\right], \quad g := -\log\left(\frac{d\mu_0}{d\gamma_d}\right), \quad \mathfrak{h}(a) = a\log a - a + 1.$$

On Forgetting and Stability of Score-based Generative models

Stanislas Strasman[1], Gabriel Victorino Cardoso[2], Sylvain Le Corff[1], Vincent Lemaire[1], Antonio Ocello[3]

[1]Sorbonne Université and Université Paris Cité, CNRS, LPSM, France.
[2]Geostatistics team, Centre for geosciences and geoengineering, Mines Paris, PSL University, France.
[3]CREST, Groupe ENSAE-ENSAI, ENSAE Paris, Institut Polytechnique de Paris, France.




arxiv:2601.21868

Stability I: Forgetting Property of the Backward Chain

Setting. The reverse-time SGM defines a (possibly inhomogeneous) Markov chain

$$X_{k+1} \sim P_k(X_k, \cdot),$$

whose transition kernel depends on the learned score.

Main structural result of the paper: The backward Markov chain enjoys a **forgetting (contractivity) property**: for two initializations x, y,

$$\mathrm{TV}\big(\mathrm{Law}(X_k^x), \mathrm{Law}(X_k^y)\big) \;\le\; C\,\rho^k, \qquad \rho < 1.$$

Key insight:

- The chain progressively forgets its initial condition.
- Errors at early stages do not accumulate indefinitely.
- Sampling becomes robust to initialization mismatch.

Implication for SGMs: Backward generation is *stable* even under imperfect score approximation.

Stability II: Lyapunov Propagation

Key technical tool: Lyapunov function propagation.
There exists a function $V : \mathbf{X} \to \mathbb{R}_+$ such that

$$\mathbb{E}[V(X_{k+1}) \mid X_k = x] \leq (1-\alpha)V(x) + b, \qquad \alpha > 0.$$

Consequences.

- Uniform moment control along the backward trajectory.
- No explosion under score perturbations.
- Tightness and stability of the generated distribution.

For SGMs:

- Prevents instability from large score gradients.
- Guarantees bounded propagation of discretization error.
- Ensures robustness of sampling dynamics.

Key insight: Lyapunov control $\Rightarrow$ stable generation.

Stability of Score-Based Generation

Combining forgetting + Lyapunov propagation yields:

1. Controlled error accumulation

$\mathrm{KL} = E_{\mathrm{init}} + E_{\mathrm{score}} + E_{\mathrm{disc}}$, with no exponential blow-up.

2. Stability w.r.t. score perturbation

Small regression error in the score $\|s_\theta - s^\star\|^2$ induces only linear degradation in KL.

3. Practical implication

SGM generation is:

- robust to imperfect training,
- robust to discretization,
- insensitive to initialization at large horizon.

Key insight: The backward SGM dynamics behaves as a stable controlled Markov chain, not as an unstable reverse-time diffusion.

Conclusion

1. **Why do SGMs work?**
   - Solve a **sequential problem**: bridging sampling from simple reference to complex data.
   - **Log-concavity** helps ensure contractivity and stability of the backward dynamics.
   - Explicit non-asymptotic convergence bounds in different metrics

$$\text{Mixing} + \text{Score Approximation} + \text{Discretization}.$$

   - Stability ensured by **forgetting** and **Lyapunov propagation**.
2. **Is the pipeline general?**
   - **Time-reversal = optimal control problem**. Link with stochastic control.
   - Discrete state spaces: $\{0, 1\}$ and $\{0, 1\}^d$.
   - Control-theoretic formulation suggests natural ways to encode priors.
3. **Stochastic vs deterministic generation?**
   - Stochastic backward dynamics enjoys forgetting and stability.
   - Robustness to initialization and score perturbations.

Perspectives

Future directions.

- Convergence guarantees for **conditional SGMs**.
- Incorporating **heavy-tailed priors** or non-Gaussian references.
- Beyond OU: structured forward dynamics.
- Sharper Wasserstein / path-space stability bounds.

SGMs are controlled stochastic systems — not just denoisers.

[Pré71] András Prékopa. "Logarithmic concave measures with application to stochastic programming". In: *Acta Scientiarum Mathematicarum* 32 (1971), pp. 301–316.

[BL76] Herm Jan Brascamp and Elliott H Lieb. "On extensions of the Brunn-Minkowski and Prékopa-Leindler theorems, including inequalities for log concave functions, and with an application to the diffusion equation". In: *Journal of functional analysis* 22.4 (1976), pp. 366–389.

[And82] Brian DO Anderson. "Reverse-time diffusion equation models". In: *Stochastic Processes and their Applications* 12.3 (1982), pp. 313–326.

[HP86] U. G. Haussmann and E. Pardoux. "Time reversal of diffusions". In: *The Annals of Probability* 14.4 (1986), pp. 1188–1205.

[HD05] Aapo Hyvärinen and Peter Dayan. "Estimation of non-normalized statistical models by score matching.". In: *Journal of Machine Learning Research* 6.4 (2005).

[Gil07] Daniel T Gillespie. "Stochastic simulation of chemical kinetics". In: *Annu. Rev. Phys. Chem.* 58.1 (2007), pp. 35–55.

[Vin11] Pascal Vincent. "A connection between score matching and denoising autoencoders". In: *Neural computation* 23.7 (2011), pp. 1661–1674.

[SW14] Adrien Saumard and Jon A Wellner. "Log-concavity and strong log-concavity: a review". In: *Statistics surveys* 8 (2014), p. 45.

[CL22] Giovanni Conforti and Christian Léonard. "Time reversal of Markov processes with jumps under a finite entropy condition". In: *Stochastic Processes and their Applications* 144 (2022), pp. 85–124.

[CCGL23] Patrick Cattiaux, Giovanni Conforti, Ivan Gentil, and Christian Léonard. “Time reversal of diffusion processes under a finite entropy condition”. In: *Annales de l’Institut Henri Poincaré (B) Probabilités et Statistiques*. Vol. 59. 4. Institut Henri Poincaré. 2023, pp. 1844–1881.

[CDS25] Giovanni Conforti, Alain Durmus, and Marta Gentiloni Silveri. “KL convergence guarantees for score diffusion models under minimal data assumptions”. In: *SIAM Journal on Mathematics of Data Science* 7.1 (2025), pp. 86–109.

[GSO25] Marta Gentiloni Silveri and Antonio Ocello. “Beyond log-concavity and score regularity: Improved convergence bounds for score-based generative models in W2-distance”. In: *Forty-second International Conference on Machine Learning*. 2025.

[PSOCD25] Le-Tuyet-Nhi Pham, Dario Shariatian, Antonio Ocello, Giovanni Conforti, and Alain Oliviero Durmus. “Discrete markov probabilistic models: An improved discrete score-based framework with sharp convergence bounds under minimal assumptions”. In: *Forty-second International Conference on Machine Learning*. 2025.

[SCCLO26] Stanislas Strasman, Gabriel Cardoso, Sylvain Le Corff, Vincent Lemaire, and Antonio Ocello. “On Forgetting and Stability of Score-based Generative models”. In: *arXiv preprint arXiv:2601.21868* (2026).

THANK YOU FOR YOUR ATTENTION

Center for Research in Economics and Statistics
CREST UMR 9194 - 5 avenue Henry Le Chatelier
91120 Palaiseau