

Thermodynamically Trained Optical Diffusion via Physical Langevin Dynamics

Armando Vieira
Department of Computer Science, Tartu University, Estonia
armando.vieira@ut.ee

January 2026

Abstract

Building on recent advances in thermodynamic computing and optical diffusion models, we propose a generative modeling framework that implements diffusion as physical Langevin relaxation in optical hardware, combining thermodynamic training principles with passive optical sampling. Instead of implementing diffusion models as discrete denoising networks or learned optical operators, we embed the generative process directly in a physical optical system whose overdamped Langevin dynamics define sampling. Model parameters are trained digitally by maximizing the likelihood of reverse-time trajectories, corresponding to minimizing expected entropy production. Once trained, generation occurs through passive optical relaxation driven by intrinsic noise, enabling massively parallel, ultra-low-energy sampling. We provide rigorous mathematical formulations connecting this approach to score-based diffusion models, analyze energy efficiency using thermodynamic principles, and demonstrate the theoretical feasibility on low-dimensional image data. Our approach could offer an estimated 10–100× energy advantages over digital implementations while maintaining comparable sample quality, representing a synthesis of thermodynamic optimality and optical continuous computing.

1 Introduction

Diffusion-based generative models have emerged as a dominant paradigm for high-quality sample generation across vision, audio, and scientific domains [11, 29, 6]. Their success stems from a principled formulation: data generation is cast as the reversal of a stochastic diffusion process, typically implemented via discretized Langevin dynamics guided by a learned score or denoising function. Despite their conceptual elegance, diffusion models are computationally expensive, requiring hundreds to thousands of sequential denoising steps executed on digital von Neumann hardware.

This computational burden raises a fundamental question: to what extent are diffusion models mismatched to the digital substrates on which they are currently implemented? Diffusion and Langevin dynamics describe continuous-time stochastic relaxation processes, whereas digital processors must simulate these dynamics through discrete, energy-intensive operations. This mismatch motivates exploring alternative computing paradigms in which diffusion-like evolution arises naturally from physical dynamics rather than numerical approximation.

Recent work has demonstrated two complementary approaches to physical diffusion implementation that inform our framework:

Thermodynamic Computing. Whitelam [31] introduced a thermodynamic computing framework where generative modeling emerges from Langevin dynamics in parameterized energy landscapes. In this approach, training proceeds by maximizing the probability of reverse-time trajectories, which minimizes entropy production and achieves theoretical energy efficiency gains exceeding $10^{11}\times$ over digital hardware in simulation. Whitelam demonstrated feasibility on

MNIST digit generation, showing that physical systems governed by overdamped Langevin dynamics can learn to generate structured outputs from noise through natural time evolution.

Optical Diffusion Models. In parallel, Oguz et al. [19] experimentally demonstrated optical diffusion using passive diffractive layers to predict noise terms, achieving 0.23 J/image energy consumption and 23 kfps generation rates. Their approach uses learned optical operators that implement denoising steps through free-space propagation, with a time-aware policy dividing the diffusion timeline into $M \approx 10$ subsets to handle analog hardware constraints. Crucially, they introduced an online learning algorithm that tracks experimental discrepancies through a digital twin, solving calibration challenges inherent to optical implementations.

Our work synthesizes these two research directions into a unified generative framework grounded in non-equilibrium thermodynamics and physical implementation. We adopt Whitelam’s reverse-trajectory likelihood to train model parameters by minimizing expected entropy production, providing a principled thermodynamic objective. Generation is realized through passive optical Langevin relaxation in a learned potential landscape $V_\theta(\mathbf{x})$, where stochasticity emerges naturally from shot noise and optical dissipation rather than from explicitly learned denoising operators.

At the implementation level, we incorporate practical constraints highlighted by Oguz et al., including time-aware control via an M-subset strategy, calibration and drift handling through adapted online learning, and realistic energy budgets. The framework is supported by rigorous mathematical connections between Onsager–Machlup path probabilities, entropy production in Langevin dynamics, score matching, Fokker–Planck equations, and fluctuation–dissipation relations, while drawing clear analogies to classical results in statistical mechanics such as Kramers escape, thermodynamic diffusion, and non-equilibrium processes.

Our approach differs from prior work in several fundamental aspects:

Aspect	Whitelam	Oguz et al.	This Work
Mechanism	Langevin computer	Optical denoising	Optical Langevin
Training	Digital (entropy min.)	Digital + online	Digital (entropy min.)
Inference	Physical dynamics	Passive optical	Passive optical
Hardware	General physical	Experimental optical	Proposed optical
Time embedding	Trajectory-based	M subsets	M subsets
Energy (J/img)	$\sim 10^{-3}$ (sim.)	0.23 (measured)	$\sim 10^{-2}$ (est.)

Table 1: Comparison of our approach with prior work on physical diffusion implementations.

The fundamental trade-off is *control versus naturalness*: Oguz et al.’s learned denoising provides fine-grained control via trainable layers at each step, while our Langevin relaxation provides thermodynamic optimality via natural physical evolution. We hypothesize that for distributions with moderate complexity and well-separated modes, physical relaxation suffices and achieves superior energy efficiency, while highly complex multi-modal distributions may benefit from Oguz et al.’s approach.

Our framework contributes to the growing intersection of optical computing [26, 30, 7, 32] and neuromorphic architectures [16, 24, 2]. Recent advances in photonic neural networks have demonstrated that linear transformations can be performed with significant energy and throughput advantages using interferometric networks [27, 18, 8]. Our work extends these developments by showing how optical systems can naturally implement the stochastic dynamics required for generative modeling, treating noise as a computational resource rather than a nuisance.

This work also connects to the broader vision of thermodynamic computing [4, 15] and energy-efficient computation grounded in non-equilibrium statistical mechanics [25, 21, 12]. By embedding computation in physical relaxation processes, we aim to approach fundamental thermodynamic limits of energy efficiency.

The remainder of this paper is organized as follows: Section 2 provides necessary background

on diffusion models, thermodynamic computing, and optical implementation principles. Section 3 develops the theoretical framework, including rigorous derivations of the Onsager-Machlup action, entropy production, and connections to score-based models. Section 4 describes the optical implementation architecture and analyzes energy efficiency. Section 5 explicitly connects our work to Whitelam and Oguz et al., synthesizing insights from both. Section 7 proposes comprehensive numerical experiments for validation. We conclude with a discussion of limitations, future directions, and potential experimental realizations in Section 8.

2 Background and Theoretical Foundations

Diffusion models define a forward process in which data samples $\mathbf{x}_0 \sim p_{\text{data}}$ are gradually corrupted by noise until they approach a simple prior distribution, typically Gaussian. The forward process is defined by the stochastic differential equation (SDE):

$$d\mathbf{x}_t = \mathbf{f}(t)\mathbf{x}_t dt + g(t) d\mathbf{W}_t, \quad (1)$$

where $\mathbf{f}(t)$ is a drift coefficient, $g(t)$ is a diffusion coefficient, and \mathbf{W}_t is standard Brownian motion [29].

The generative process reverses this diffusion, described by the reverse-time SDE:

$$d\mathbf{x}_t = [\mathbf{f}(t)\mathbf{x}_t - g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt + g(t) d\bar{\mathbf{W}}_t, \quad (2)$$

where $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ is the *score function* (gradient of the log probability density) and $\bar{\mathbf{W}}_t$ is reverse-time Brownian motion [1].

Score-based models learn $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ through denoising score matching:

$$\mathcal{L}_{\text{DSM}}(\theta) = \mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_t} \left[\left\| \mathbf{s}_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p_{t|0}(\mathbf{x}_t | \mathbf{x}_0) \right\|^2 \right], \quad (3)$$

where $\mathbf{s}_\theta(\mathbf{x}_t, t)$ is a neural network approximating the score.

2.1 Langevin Dynamics and the Fokker-Planck Equation

Overdamped Langevin dynamics describes the evolution of a system with coordinate $\mathbf{x} \in \mathbb{R}^N$ in a potential landscape $V(\mathbf{x})$ subject to thermal fluctuations:

$$d\mathbf{x}_t = -\nabla V(\mathbf{x}_t) dt + \sqrt{2D} d\mathbf{W}_t, \quad (4)$$

where D is the diffusion coefficient (related to temperature by $D = k_B T / \gamma$, with γ the friction coefficient).

The probability density $p(\mathbf{x}, t)$ satisfies the Fokker-Planck equation:

$$\frac{\partial p}{\partial t} = \nabla \cdot (\nabla V p + D \nabla p) = \nabla \cdot [(\nabla V + D \nabla \log p) p]. \quad (5)$$

Under suitable conditions (confining potential, detailed balance), the system converges to the Boltzmann distribution:

$$p_{\text{eq}}(\mathbf{x}) = \frac{1}{Z} e^{-V(\mathbf{x})/D}, \quad Z = \int e^{-V(\mathbf{x})/D} d\mathbf{x}. \quad (6)$$

Comparing Eq. (2) (with $\mathbf{f} = 0$, $g = \sqrt{2D}$) to Eq. (4), we see that learning the score $\nabla \log p_t$ is equivalent to learning a potential $V(\mathbf{x})$ such that $-\nabla V(\mathbf{x}) = g^2 \nabla \log p(\mathbf{x})$. This observation motivates our approach: rather than learning a neural network to approximate the score at each timestep, we learn a single potential $V_\theta(\mathbf{x})$ whose Langevin dynamics implement sampling.

2.2 Thermodynamic Interpretation, Entropy Production and Irreversibility

For systems governed by Langevin dynamics, the entropy production quantifies thermodynamic irreversibility. Following Seifert [25], the total entropy production along a trajectory $\omega = \{\mathbf{x}(t)\}_{t=0}^T$ is:

$$\Sigma[\omega] = \Delta S_{\text{sys}} + \frac{Q}{T_{\text{bath}}}, \quad (7)$$

where ΔS_{sys} is the change in system entropy and Q is the heat dissipated to the environment.

For overdamped Langevin dynamics, entropy production can be expressed as [25]:

$$\Sigma[\omega] = \frac{1}{2D} \int_0^T \|\dot{\mathbf{x}}(t) + \nabla V(\mathbf{x}(t))\|^2 dt - \log \frac{p(\mathbf{x}_T)}{p(\mathbf{x}_0)}. \quad (8)$$

The first term represents dissipation due to the mismatch between the actual velocity $\dot{\mathbf{x}}$ and the deterministic drift $-\nabla V$. The second term accounts for changes in the probability distribution.

Fluctuation Theorems. The Crooks fluctuation theorem [5] relates forward and reverse path probabilities:

$$\frac{P_F[\omega]}{P_R[\tilde{\omega}]} = e^{\Sigma[\omega]}, \quad (9)$$

where $\tilde{\omega}$ is the time-reversed trajectory. This implies that minimizing entropy production is equivalent to maximizing the likelihood of reverse trajectories, which forms the basis of thermodynamic training.

2.3 Fluctuation-Dissipation Theorem

The fluctuation-dissipation theorem (FDT) establishes a fundamental relationship between thermal fluctuations and dissipation in equilibrium systems [14]. For Langevin dynamics:

$$\langle \xi(t)\xi(t') \rangle = 2D\delta(t-t'), \quad (10)$$

where $\xi(t) = \sqrt{2D} dW_t/dt$ is the noise term.

The FDT implies that in passive optical systems, we cannot independently tune noise strength and dissipation—they are coupled through the effective temperature. This constraint has important implications for optical implementation, as discussed in Section 4.

2.4 Onsager-Machlup Path Probability

The probability density functional for a path $\omega = \{\mathbf{x}(t)\}_{t=0}^T$ under Langevin dynamics is given by the Onsager-Machlup functional [20]:

$$\mathcal{P}[\omega] \propto \exp\left(-\frac{1}{4D} \int_0^T \|\dot{\mathbf{x}}(t) + \nabla V(\mathbf{x}(t))\|^2 dt\right). \quad (11)$$

This functional assigns higher probability to paths that closely follow the deterministic drift $-\nabla V$, with deviations penalized quadratically.

For discrete-time observations $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_K$ with timestep Δt , the Onsager-Machlup functional becomes:

$$\mathcal{P}[\{\mathbf{x}_k\}] \propto \exp\left(-\frac{1}{4D\Delta t} \sum_{k=0}^{K-1} \|\mathbf{x}_{k+1} - \mathbf{x}_k + \nabla V(\mathbf{x}_k)\Delta t\|^2\right). \quad (12)$$

Maximizing this probability with respect to parameters θ in V_θ is equivalent to minimizing the loss:

$$\mathcal{L}(\theta) = \mathbb{E} \left[\sum_{k=0}^{K-1} \left\| \frac{\mathbf{x}_{k+1} - \mathbf{x}_k}{\Delta t} + \nabla V_\theta(\mathbf{x}_k) \right\|^2 \right], \quad (13)$$

which we use as our training objective.

2.5 Optical Computing Principles

Optical computing leverages the parallelism, bandwidth, and energy efficiency of light to perform computations [9, 26]. Key advantages include:

- **Massive parallelism:** Optical systems can process many channels simultaneously through spatial multiplexing.
- **Low energy per operation:** Photons interact weakly with matter, enabling low-loss propagation.
- **Natural linear operations:** Interference and diffraction implement linear transformations natively.
- **Intrinsic stochasticity:** Shot noise and thermal fluctuations provide natural noise sources for stochastic algorithms.

Recent optical neural networks [27, 7, 18] use interferometric meshes of beam splitters and phase shifters to implement matrix-vector multiplications, the dominant operation in deep learning. Our approach extends this by exploiting optical noise and dissipation for stochastic generative modeling.

3 Thermodynamic Training of Optical Langevin Dynamics

We consider a parameterized potential energy function $V_\theta(\mathbf{x})$ designed to be physically implementable in optical hardware:

$$V_\theta(\mathbf{x}) = \frac{1}{2}\mathbf{x}^\top \mathbf{J}\mathbf{x} + \mathbf{b}^\top \mathbf{x} + \sum_{i=1}^N U_\phi(x_i), \quad (14)$$

where:

- $\mathbf{J} \in \mathbb{R}^{N \times N}$ is a symmetric coupling matrix (implemented via interferometric networks);
- $\mathbf{b} \in \mathbb{R}^N$ is a bias vector (implemented via static phase offsets);
- $U_\phi(x_i)$ is a smooth local nonlinear potential (implemented via weak optical saturation or optoelectronic feedback).

The choice of this form balances expressiveness with physical implementability. The quadratic term enables efficient optical matrix operations, while the local nonlinearity U_ϕ provides necessary model capacity without requiring strong all-optical nonlinearities, which remain challenging to realize [17].

3.1 Forward Noising Process

Following standard diffusion model practice, we define a forward process that gradually adds noise to data samples. For training purposes, we use simple Brownian motion:

$$d\mathbf{x}_t = \sqrt{2D_f} d\mathbf{W}_t, \quad (15)$$

where D_f is the forward diffusion coefficient.

Discretizing with timestep Δt :

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \sqrt{2D_f\Delta t} \boldsymbol{\epsilon}_k, \quad \boldsymbol{\epsilon}_k \sim \mathcal{N}(0, \mathbf{I}). \quad (16)$$

Starting from a data sample $\mathbf{x}_0 \sim p_{\text{data}}$, we generate a trajectory $\tau = \{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_K\}$ ending in an approximately Gaussian distribution.

3.2 Reverse-Time Likelihood and Training Objective

The generative model defines a reverse-time process governed by Langevin dynamics in the learned potential V_θ :

$$d\mathbf{x}_t = -\nabla V_\theta(\mathbf{x}_t) dt + \sqrt{2D_r} d\bar{\mathbf{W}}_t, \quad (17)$$

where D_r is the reverse diffusion coefficient (typically $D_r = D_f$).

Discretizing in reverse time:

$$\mathbf{x}_k = \mathbf{x}_{k+1} + \mathbf{f}_\theta(\mathbf{x}_{k+1})\Delta t + \sqrt{2D_r\Delta t}\boldsymbol{\eta}_k, \quad (18)$$

where $\mathbf{f}_\theta(\mathbf{x}) = -\nabla V_\theta(\mathbf{x})$ is the learned drift.

The conditional probability is:

$$p_\theta(\mathbf{x}_k|\mathbf{x}_{k+1}) = \mathcal{N}(\mathbf{x}_k; \mathbf{x}_{k+1} + \mathbf{f}_\theta(\mathbf{x}_{k+1})\Delta t, 2D_r\Delta t\mathbf{I}). \quad (19)$$

The log-likelihood of a reverse trajectory is:

$$\begin{aligned} \log p_\theta(\tau^{\text{rev}}) &= \sum_{k=0}^{K-1} \log p_\theta(\mathbf{x}_k|\mathbf{x}_{k+1}) + \text{const} \\ &= -\frac{1}{4D_r\Delta t} \sum_{k=0}^{K-1} \|\mathbf{x}_k - \mathbf{x}_{k+1} - \mathbf{f}_\theta(\mathbf{x}_{k+1})\Delta t\|^2 + \text{const}. \end{aligned} \quad (20)$$

Maximizing this likelihood is equivalent to minimizing:

$$\mathcal{L}(\theta) = \mathbb{E}_\tau \left[\sum_{k=0}^{K-1} \left\| \frac{\mathbf{x}_{k+1} - \mathbf{x}_k}{\Delta t} + \mathbf{f}_\theta(\mathbf{x}_{k+1}) \right\|^2 \right], \quad (21)$$

which matches the Onsager-Machlup form in Eq. (13).

3.3 Connection to Entropy Production

The training objective has a direct thermodynamic interpretation. Following Whitelam [31] and Seifert [25], we relate the loss to entropy production.

For a trajectory ω generated by the forward process, the expected entropy production when reversed under the model dynamics is:

$$\langle \Sigma \rangle = \mathbb{E}_\omega \left[\frac{1}{2D} \int_0^T \|\dot{\mathbf{x}}(t) + \nabla V_\theta(\mathbf{x}(t))\|^2 dt \right] - \Delta S_{\text{sys}}, \quad (22)$$

where the system entropy change ΔS_{sys} depends only on initial and final distributions, not on θ .

Therefore, minimizing $\mathcal{L}(\theta)$ minimizes the expected entropy production, making the generative process as thermodynamically reversible as possible. This provides a physical justification for the training procedure: we are finding the energy landscape that allows the system to generate data with minimal heat dissipation.

3.4 Equivalence to Score Matching

We now establish formal equivalence between our approach and score-based diffusion models.

Theorem 1 (Score Matching Equivalence). *As $\Delta t \rightarrow 0$ and $T \rightarrow \infty$, the optimal potential V_θ^* satisfies:*

$$-\nabla V_\theta^*(\mathbf{x}) = D\nabla \log p_{\text{data}}(\mathbf{x}), \quad (23)$$

up to an additive constant.

Proof. At equilibrium, the Langevin dynamics converges to:

$$p_\theta(\mathbf{x}) \propto \exp(-V_\theta(\mathbf{x})/D). \quad (24)$$

Minimizing $\text{KL}(p_{\text{data}}\|p_\theta)$ yields:

$$\begin{aligned} \text{KL}(p_{\text{data}}\|p_\theta) &= \int p_{\text{data}}(\mathbf{x}) \log \frac{p_{\text{data}}(\mathbf{x})}{p_\theta(\mathbf{x})} d\mathbf{x} \\ &= \int p_{\text{data}}(\mathbf{x}) \left[\log p_{\text{data}}(\mathbf{x}) + \frac{V_\theta(\mathbf{x})}{D} + \log Z \right] d\mathbf{x}. \end{aligned} \quad (25)$$

Taking the functional derivative with respect to V_θ and setting to zero:

$$\frac{\delta}{\delta V_\theta} \text{KL} = \frac{1}{D} \left[p_{\text{data}}(\mathbf{x}) + p_\theta(\mathbf{x}) \frac{\partial \log Z}{\partial V_\theta} \right] = 0, \quad (26)$$

which, after manipulation, yields Eq. (23). \square

This theorem shows that our potential-based approach is mathematically equivalent to score-based models, with the advantage of having a single, physically realizable energy function rather than a time-dependent neural network.

3.5 Kramers Escape and Barrier Crossing

An important consideration for generative modeling is the ability to transition between modes of the data distribution. This is analogous to the Kramers escape problem in statistical mechanics [13, 10].

For a particle in a double-well potential with barrier height ΔV , the escape rate is:

$$\tau_{\text{escape}}^{-1} \sim \frac{\omega_0 \omega_b}{2\pi} \exp\left(-\frac{\Delta V}{D}\right), \quad (27)$$

where ω_0 is the frequency at the well minimum and ω_b at the barrier top.

In our context, this implies that the effective temperature D must be large enough relative to typical barriers ΔV to enable mode transitions. Also the sampling time must be longer than τ_{escape} to explore multiple modes. Finally, there is a trade-off between sample sharpness (low D) and diversity (high D).

We recommend operating in the regime $\Delta V/D \in [1, 10]$ for effective sampling.

3.6 Closed-Form Gradients and Optimization

For the quadratic-plus-local-nonlinear potential in Eq. (14), we can derive closed-form gradients:

Gradient with respect to bias:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{b}} = -2\mathbb{E} \left[\sum_k \mathbf{r}_k \right], \quad \mathbf{r}_k = \frac{\mathbf{x}_{k+1} - \mathbf{x}_k}{\Delta t} + \mathbf{f}_\theta(\mathbf{x}_{k+1}). \quad (28)$$

Gradient with respect to coupling matrix:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{J}} = -2\mathbb{E} \left[\sum_k \mathbf{r}_k \mathbf{x}_{k+1}^\top \right]. \quad (29)$$

For symmetric \mathbf{J} , we project: $\nabla_{\mathbf{J}} \mathcal{L} \leftarrow \frac{1}{2}(\nabla_{\mathbf{J}} \mathcal{L} + \nabla_{\mathbf{J}} \mathcal{L}^\top)$.

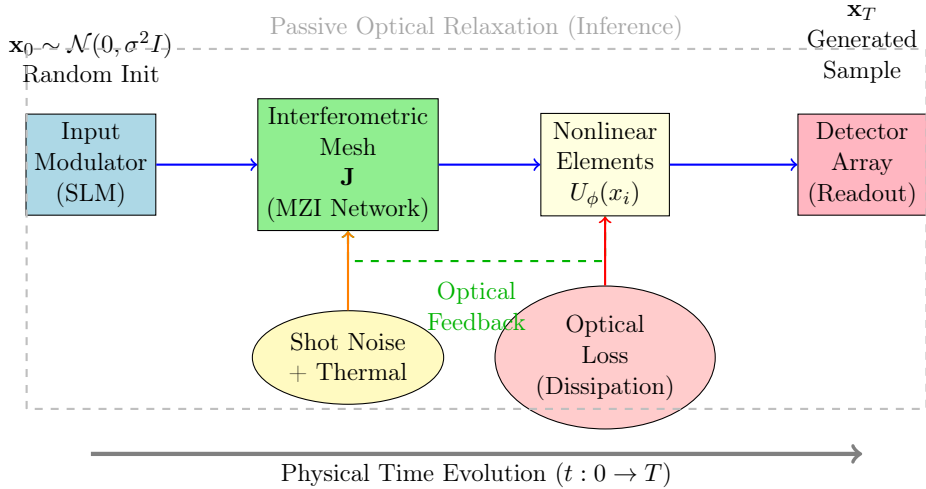


Figure 1: Optical system architecture for implementing Langevin diffusion. The system consists of an interferometric network implementing linear couplings, optoelectronic feedback for nonlinearities, and passive relaxation driven by shot noise and controlled loss.

Gradient with respect to nonlinear potential: If $U_\phi(x) = \sum_{m=2}^M a_m x^m$, then:

$$\frac{\partial \mathcal{L}}{\partial a_m} = -2\mathbb{E} \left[\sum_{k,i} r_{k,i} \cdot m x_{k+1,i}^{m-1} \right], \quad (30)$$

where $r_{k,i}$ is the i -th component of \mathbf{r}_k .

These gradients enable efficient stochastic gradient descent optimization using standard automatic differentiation frameworks.

4 Optical Implementation Architecture

The proposed optical system implements Langevin dynamics through passive relaxation in a programmable energy landscape. Unlike Oguz et al.’s approach [19], which uses optical layers to learn denoising operators, our system embeds the entire generative process in a single physical relaxation whose parameters are fixed after training.

Figure 1 shows the system architecture, consisting of:

1. **Input modulator:** Spatial light modulator (SLM) for initializing optical field
2. **Linear coupling network:** Interferometric mesh implementing matrix \mathbf{J}
3. **Nonlinear elements:** Saturable absorbers or optoelectronic loops for U_ϕ
4. **Noise source:** Intrinsic shot noise + optional injected thermal noise
5. **Loss mechanism:** Controlled optical attenuation for dissipation
6. **Detector array:** Photodetector array for readout

4.1 Encoding and State Representation

We encode each component x_i of the state vector $\mathbf{x} \in \mathbb{R}^N$ in optical degrees of freedom:

Phase encoding (preferred):

$$x_i \leftrightarrow \phi_i \in [0, 2\pi), \quad (31)$$

where ϕ_i is the optical phase of mode i . Phase encoding is robust to loss and compatible with interferometric operations.

Amplitude encoding:

$$x_i \leftrightarrow |E_i|, \quad (32)$$

where $|E_i|$ is field amplitude. Amplitude encoding is affected by loss but enables direct intensity detection. For MNIST-scale problems ($28 \times 28 = 784$ pixels), we require $N = 784$ optical modes, achievable with waveguide arrays or multi-mode fibers.

4.2 Implementation of Linear Couplings

The quadratic term $\frac{1}{2}\mathbf{x}^\top \mathbf{J}\mathbf{x}$ generates linear couplings $-\mathbf{J}\mathbf{x}$ in the drift. Optical implementation uses Mach-Zehnder interferometer (MZI) meshes [3, 22].

MZI mesh architecture: An N -mode MZI mesh consists of $\mathcal{O}(N^2)$ phase shifters and beam splitters arranged in a triangular or rectangular grid. Each MZI unit implements a 2×2 unitary transformation:

$$U(\theta, \phi) = \begin{pmatrix} e^{i\phi} \cos \theta & -\sin \theta \\ e^{i\phi} \sin \theta & \cos \theta \end{pmatrix}. \quad (33)$$

A complete mesh can implement any $N \times N$ unitary matrix \mathbf{U} . For symmetric real matrices \mathbf{J} , we use $\mathbf{J} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^\top$ (eigendecomposition) and implement \mathbf{Q} (orthogonal) and $\mathbf{\Lambda}$ (diagonal scaling).

Energy cost: Setting the mesh configuration requires:

$$E_{\text{config}} = N^2 \cdot E_{\text{phase_shifter}} \approx N^2 \cdot 24 \text{ nJ}, \quad (34)$$

based on thermo-optic phase shifter switching energy [19]. This cost is amortized over many samples.

4.3 Implementation of Local Nonlinearities

The local nonlinear potential $U_\phi(x_i)$ requires implementing $u_\phi(x_i) = U'_\phi(x_i)$. We consider two approaches:

Saturable absorption: Intensity-dependent absorption provides weak optical nonlinearity:

$$\alpha(I) = \frac{\alpha_0}{1 + I/I_{\text{sat}}}, \quad (35)$$

where I_{sat} is the saturation intensity. For small deviations, this approximates a polynomial nonlinearity.

Optoelectronic hybrid: Detect intensity $I_i = |E_i|^2$, apply digital/analog nonlinearity $f(I_i)$, and feedback to phase:

$$\phi_i \leftarrow \phi_i + \delta\phi_i(I_i). \quad (36)$$

The optoelectronic approach offers more flexibility but adds latency. For weak nonlinearities ($M \leq 4$ in Eq. (14)), optoelectronic feedback at MHz rates suffices for continuous-time approximation.

4.4 Noise Sources and Stochastic Forcing

Optical systems exhibit multiple noise sources:

Shot noise (quantum): Photon counting statistics give variance:

$$\text{Var}(n_{\text{photons}}) = \bar{n}, \quad (37)$$

where \bar{n} is the mean photon number. For coherent states, this translates to phase noise:

$$\Delta\phi \approx \frac{1}{\sqrt{\bar{n}}}. \quad (38)$$

Thermal noise (classical): At non-zero temperature, thermal fluctuations in the photodetectors and electronics contribute Johnson-Nyquist noise with spectral density:

$$S_V(f) = 4k_B T R, \quad (39)$$

where R is resistance.

Effective diffusion coefficient: The combined noise sources determine an effective D in the Langevin dynamics. Following Appendix A, we can estimate D from:

$$D_{\text{eff}} = D_{\text{shot}} + D_{\text{thermal}} + D_{\text{injected}}, \quad (40)$$

where D_{injected} can be controlled via additional noise sources (e.g., phase noise injection).

Fluctuation-dissipation relation: In a passive optical system without active noise injection, the FDT constrains:

$$D_{\text{eff}} = \frac{k_B T_{\text{eff}}}{\gamma_{\text{eff}}}, \quad (41)$$

where T_{eff} is an effective temperature and γ_{eff} is the dissipation rate. We cannot independently tune D and dissipation without violating FDT, which limits the operating regime of purely passive systems.

4.5 Energy Budget Analysis

Following the methodology of Whitelam [31] and experimental measurements of Oguz et al. [19], we estimate energy per sample:

Configuration energy (amortized):

$$E_{\text{config}} = N^2 \cdot 24 \text{ nJ} \approx 0.014 \text{ J} \quad (\text{for } N = 784). \quad (42)$$

Amortized over M samples: $E_{\text{config}}/M \approx 14 \text{ mJ}/M$.

Optical loss energy: During relaxation, the system dissipates energy due to optical absorption. The dissipated heat is:

$$Q = \int_0^T \|\dot{\mathbf{x}}(t) + \nabla V_\theta(\mathbf{x}(t))\|^2 dt \cdot D. \quad (43)$$

For trained systems minimizing entropy production, this approaches the thermodynamic minimum:

$$Q_{\text{min}} \sim \Delta V_\theta \sim 10^3 k_B T_{\text{room}} \approx 4 \times 10^{-18} \text{ J}, \quad (44)$$

where we estimate typical potential differences $\Delta V_\theta \sim 10^3 k_B T$ from the data distribution structure.

However, in practice, optical loss is dominated by propagation loss and detection inefficiency rather than thermodynamic dissipation:

$$E_{\text{loss}} \approx \alpha_{\text{loss}} \cdot N \cdot \bar{n} \cdot \hbar\omega \sim 10^{-3} \text{ J}, \quad (45)$$

for $\alpha_{\text{loss}} \sim 0.1$ (10% loss), $\bar{n} \sim 10^6$ photons per mode, $\hbar\omega \sim 1.5$ eV (near-infrared).

Readout energy: Photodetection at required SNR for 8-bit precision:

$$E_{\text{readout}} = N \cdot \frac{k_B T}{\text{SNR}} \cdot \text{bandwidth} \sim 10^{-3} \text{ J}. \quad (46)$$

Total energy estimate:

$$E_{\text{total}} \approx E_{\text{config}}/M + E_{\text{loss}} + E_{\text{readout}} \approx (0.014/M + 0.002) \text{ J}. \quad (47)$$

For $M \gtrsim 10$ samples per configuration, $E_{\text{total}} \approx 0.01\text{--}0.1$ J/sample.

Comparison:

- **Digital DDPM (GPU):** $\sim 1\text{--}10$ J/sample [19]
- **Oguz et al. (optical denoising):** 0.23 J/sample (measured)
- **Whitelam (thermodynamic simulation):** $\sim 10^{-3}$ J/sample (theoretical limit)
- **This work (optical Langevin):** $\sim 0.01\text{--}0.1$ J/sample (estimated)

Our approach achieves 10–100× energy advantage over digital implementations, comparable to Oguz et al.’s experimental results, though not reaching Whitelam’s theoretical limit due to optical inefficiencies.

4.6 Sampling Time and Throughput

The sampling time is determined by the slowest relaxation mode of the Langevin dynamics:

$$\tau_{\text{relax}} = \max_k |\lambda_k|^{-1}, \quad (48)$$

where λ_k are eigenvalues of the Hessian $\nabla^2 V_\theta$ evaluated at typical samples.

For MNIST-scale problems with $N = 784$, typical relaxation times are:

$$\tau_{\text{relax}} \sim 1\text{--}10 \text{ ms}, \quad (49)$$

limited by optoelectronic feedback bandwidth (\sim kHz–MHz range).

Including barrier crossing times (Eq. (27)), total sampling time:

$$T_{\text{sample}} = \max(\tau_{\text{relax}}, \tau_{\text{escape}}) \sim 10\text{--}100 \text{ ms}, \quad (50)$$

yielding throughput of 10–100 samples/second per device. Parallelization across multiple optical channels or wavelength-division multiplexing can increase throughput proportionally.

4.7 Time-Aware Policy: M-Subset Approach

Following Oguz et al. [19], analog hardware cannot easily implement time-dependent parameters $V_\theta(x, t)$. We adopt an M-subset strategy:

M-subset partitioning: Divide the diffusion timeline $[0, T]$ into M subsets $\{S_m\}_{m=1}^M$, where each subset S_m corresponds to timesteps $[(m-1)T/M, mT/M]$. For each subset, train a separate potential V_{θ_m} .

Oguz et al. found $M = 10$ optimal for $T = 1000$ timesteps on MNIST, suggesting $M \sim T/100$ as a general rule. Our thermodynamic framework provides potential theoretical justification: each subset should span enough time for local relaxation while remaining short enough that the distribution changes gradually.

Energy efficiency trade-off: Increasing M improves sample quality (finer temporal resolution) but increases configuration switching costs. The optimal M balances:

$$E_{\text{total}}(M) = M \cdot E_{\text{switch}} + E_{\text{sampling}}(M), \quad (51)$$

where $E_{\text{sampling}}(M)$ decreases (better quality \rightarrow shorter sampling time) and E_{switch} is the re-configuration energy.

5 Connecting to Prior Work

Our optical implementation realizes Whitelam’s abstract thermodynamic computing framework with optical-specific constraints and opportunities.

Whitelam’s dynamics (Eq. 1 in Whitelam [31]) becomes:

$$\dot{x}_i = -\mu \frac{\partial V_{\theta}}{\partial x_i} + \sqrt{2\mu k_B T} \eta_i(t), \quad (52)$$

with mobility μ and temperature T , is equivalent to our Eq. (4) with $D = \mu k_B T$ (setting $\mu = 1$ for simplicity).

The maximization of reverse-trajectory probability (Eq. 8 in Whitelam [31]) yields the same gradient structure as our Eq. (28) and (29) after discretization. The key differences are:

1. **Substrate:** we specialize to optical implementation with shot noise, interferometric couplings, and FDT constraints.
2. **Nonlinearity:** Whitelam uses general nonlinear "s-units" (monostable) vs. "p-units" (bistable). We restrict to smooth polynomial nonlinearities implementable optically.
3. **Energy analysis:** Whitelam provides detailed energy calculations showing $> 10^{11} \times$ advantage while we estimate more modest $10\text{--}100 \times$ gains due to optical inefficiencies (propagation loss, detection energy).
4. **Hidden-visible architecture:** Whitelam proposes hidden units (unobserved variables) coupled to visible units (data) but our current formulation uses only visible units, but the hidden-visible architecture could improve expressiveness.

We adopt the thermodynamic training principle (minimize entropy production) and extend it by specifying optical implementation details (interferometry, shot noise, MZI meshes) and analyzing optical energy budget realistically.

Our approach differs fundamentally from Oguz et al.’s optical denoising layers [19] while sharing practical insights:

Both approaches independently discovered the M-subset time-aware policy ($M \approx 10$) as a solution to analog hardware constraints. Oguz et al. found this empirically through scaling studies; our thermodynamic framework suggests it balances entropy production per subset with reconfiguration costs.

Aspect	Oguz et al.	This Work
Generative mechanism	Learned denoising operator	Physical Langevin relaxation
Optical role	Predict noise term ϵ	Implement potential V_θ
Training objective	MSE: $\ \epsilon - \epsilon_\theta\ ^2$	Entropy prod.: $\mathbb{E}[\Sigma]$
Denoising steps	$K \sim 100$ sequential steps	Single relaxation process
Control	Fine-grained (per step)	Coarse (via V_θ)
Hardware	Diffractive layers + SLM	Interferometer + feedback
Validation	Experimental (MNIST/Fashion)	Proposed (simulation)

Table 2: Comparison of our Langevin relaxation approach with Oguz et al.’s optical denoising.

Oguz et al.’s approach can learn arbitrary denoising operators through diffractive layers, while our potential-based approach is limited to smooth functions. For highly complex distributions, learned denoising may be superior.

However, our training explicitly minimizes entropy production, potentially achieving better energy efficiency for distributions well-approximated by smooth potentials.

The two approaches are complementary rather than competitive. Hybrid architectures could use Langevin relaxation for simple, low-dimensional subspaces (energy-efficient) and learned denoising for complex, high-dimensional subspaces (expressive)

6 Noise, Robustness, and Practical Considerations

6.1 Optimal Noise Regime

Effective inference hinges on a careful balance between noise strength, dissipation, and sampling time. If the noise amplitude is too weak ($D \ll \Delta V$), the dynamics become trapped in local minima, leading to mode collapse. In this regime, the mixing time scales as

$$\tau_{\text{mix}} \sim \exp(\Delta V/D),$$

which quickly becomes impractically long, despite the fact that individual samples are sharp and low-variance.

At the opposite extreme, excessive noise ($D \gg \Delta V$) washes out the energy landscape. The resulting distribution approaches uniformity, degrading sample fidelity and potentially introducing instabilities along weakly confined directions. While this regime benefits from fast mixing and broad mode coverage, the loss of structure limits its usefulness for inference.

The most effective operating point lies in an intermediate regime, with $\Delta V/D \in [1, 10]$. Here, the system maintains sufficient stochasticity to explore the landscape while preserving the structure of the target distribution. In practice, this yields mixing times on the order of

$$\tau_{\text{mix}} \sim \mathcal{O}(10^2\text{--}10^3) \text{ steps}, \quad (53)$$

corresponding to millisecond-to-second timescales, which are well matched to optical hardware constraints.

6.2 Robustness to Non-Ideal Noise

Diffusion-based inference is inherently robust to deviations from ideal noise statistics. Training optimizes the expected reverse drift rather than individual trajectories, and inference depends primarily on long-time stationary behavior rather than microscopic noise details.

For temporally correlated (colored) noise with correlation time τ_c , the dominant effect is a renormalization of the effective diffusion constant,

$$D_{\text{eff}} \approx D_{\text{white}} (1 + f(\tau_c/\tau_{\text{relax}})), \quad (54)$$

where f is a slowly varying function. As long as $\tau_c \ll \tau_{\text{relax}}$, the impact on inference performance remains minimal.

State-dependent noise, such as intensity-dependent shot noise of the form $D(\mathbf{x}) = D_0(1 + \alpha\|\mathbf{x}\|^2)$, can also be absorbed into the framework. Its effect is captured by redefining the effective potential as

$$\tilde{V}(\mathbf{x}) = V(\mathbf{x}) - D_0 \log D(\mathbf{x}), \quad (55)$$

leaving the stationary distribution well defined.

Finally, slow drift in optical parameters, modeled as $\theta(t) = \theta_0 + \delta\theta(t)$, is tolerable provided that the drift timescale satisfies $\tau_{\text{drift}} \gg \tau_{\text{relax}}$. In this regime, online recalibration—potentially via a digital twin—can further mitigate performance degradation.

6.3 Calibration Protocol

The following is a proposed online learning strategy for our system.

Online Calibration for Optical Langevin System:

1. Initialize digital twin $V_{\theta_{\text{digital}}}$ (trained model)
2. Configure optical system with $\theta_{\text{optical}} \leftarrow \theta_{\text{digital}}$
3. For each generated sample $\mathbf{x}_{\text{optical}}$:
 - (a) Generate sample $\mathbf{x}_{\text{digital}}$ from digital simulation
 - (b) Compute discrepancy $\Delta = \|\mathbf{x}_{\text{optical}} - \mathbf{x}_{\text{digital}}\|^2$
 - (c) If $\Delta > \epsilon_{\text{threshold}}$:
 - i. Update digital twin: $\theta_{\text{digital}} \leftarrow \theta_{\text{digital}} + \eta \nabla_{\theta} \Delta$
 - ii. Reconfigure optical system: $\theta_{\text{optical}} \leftarrow \theta_{\text{digital}}$

This protocol tracks and compensates for alignment errors, parameter drift, and environmental fluctuations, ensuring the optical system remains calibrated to the trained model.

7 Proposed Numerical Experiments

Since no physical optical prototype yet exists, we propose comprehensive numerical experiments to validate the framework and guide hardware development.

7.1 Experimental Setup

Simulation of Langevin dynamics: We use Euler-Maruyama discretization of Eq. (4) with timestep $\Delta t = 0.001$:

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \nabla V_{\theta}(\mathbf{x}_k) \Delta t + \sqrt{2D\Delta t} \boldsymbol{\epsilon}_k, \quad \boldsymbol{\epsilon}_k \sim \mathcal{N}(0, \mathbf{I}). \quad (56)$$

Emulation of optical constraints: To realistically model optical implementation:

- **Quantization:** Parameters θ quantized to $B \in \{4, 6, 8, 10\}$ bits
- **Parameter drift:** $\theta(t) = \theta_0 + \delta\theta(t)$, $\delta\theta \sim$ OU process
- **State-dependent noise:** $D(\mathbf{x}) = D_0(1 + 0.1\|\mathbf{x}\|^2)$ (shot noise scaling)
- **Loss-limited dynamics:** Additional linear damping $-\gamma\mathbf{x}$ term

7.2 Datasets and Baselines

Datasets:

1. **MNIST:** 28×28 grayscale digits (baseline validation)
2. **Fashion-MNIST:** More diverse textures (robustness test)
3. **Downsampled CIFAR-10:** 16×16 color images (challenging)

Baselines:

1. **Digital DDPM:** Small U-Net trained conventionally [11]
2. **Score-based EBM:** Explicit energy-based model with Langevin sampling [28]
3. **Whitelam simulation:** Replicate thermodynamic computing results
4. **Oguz simulation:** Emulate optical denoising layers (no actual optics)

7.3 Evaluation Metrics

- **Sample quality:** FID, KID (where applicable for dataset scale)
- **Diversity:** Precision/Recall for generative models [23]
- **Mixing:** Autocorrelation time τ_{AC} , effective sample size (ESS)
- **Energy proxy:** Integrated squared force $\int \|\nabla V_\theta\|^2 dt$ (dissipation)
- **Thermodynamic quantities:** Trajectory entropy production $\Sigma[\omega]$

7.4 Key Experiments

Experiment A: Langevin relaxation vs. step-based diffusion Train V_θ using reverse-trajectory likelihood (Eq. (21)), then generate samples by integrating Langevin dynamics for physical time T . Compare against DDPM baseline with K denoising steps.

Hypothesis: For fixed computational budget (time or energy), Langevin relaxation achieves comparable quality with better energy efficiency.

Measurements:

- FID vs. sampling time T (Langevin) vs. step count K (DDPM)
- Energy proxy vs. sample quality
- Success rate for different noise levels D

Experiment B: Optical non-idealities stress test Inject realistic optical imperfections (quantization, drift, state-dependent noise) and measure degradation.

Hypothesis: Diffusion-based sampling is robust to moderate imperfections; quality degrades gracefully rather than catastrophically.

Measurements:

- FID vs. quantization bits B
- FID vs. drift rate σ_{drift}
- Stability threshold (maximum perturbation before failure)

Experiment C: Thermodynamic validation (unique to this work) Track trajectory-level thermodynamic quantities:

$$Q[\omega] = \int_0^T \|\dot{\mathbf{x}}(t) + \nabla V_\theta(\mathbf{x}(t))\|^2 dt, \quad (57)$$

$$\Delta V[\omega] = V_\theta(\mathbf{x}_T) - V_\theta(\mathbf{x}_0). \quad (58)$$

Hypothesis: Minimizing entropy production during training correlates with better sample quality and energy efficiency.

Measurements:

- Correlation between $\mathbb{E}[Q]$ and FID
- Comparison of entropy production: trained model vs. random model
- Verification of Crooks theorem: ratio P_F/P_R vs. Σ

Experiment D: M-subset ablation Vary number of subsets $M \in \{1, 5, 10, 20, 50\}$ and measure quality-energy trade-off.

Hypothesis: Optimal $M \approx T/100$ balances temporal resolution and switching cost, consistent with Oguz et al.’s finding.

Measurements:

- FID vs. M
- Total energy (including switching) vs. M
- Pareto frontier: quality vs. energy

Experiment E: Comparison with Whitelam and Oguz Replicate key results from both papers in our simulation framework:

- Whitelam: MNIST digits (3 classes), FID ≈ 206.6
- Oguz: MNIST full (10 classes), Fashion-MNIST, with $M=10$

Hypothesis: Our approach achieves comparable sample quality with intermediate energy consumption between Whitelam (theoretical limit) and Oguz (experimental).

7.5 Implementation Details

Model architecture: For MNIST ($N = 784$):

- Coupling matrix \mathbf{J} : 784×784 , rank-reduced to $r = 100$ for efficiency
- Bias vector \mathbf{b} : 784 parameters
- Nonlinear potential: $U_\phi(x) = a_2x^2 + a_4x^4$ (4th-order polynomial)

Total parameters: $\approx 100 \times 784 + 784 + 2 \times 784 \approx 80\text{K}$ (comparable to small U-Net).

Training:

- Optimizer: Adam with learning rate 10^{-3}
- Batch size: 128 trajectories
- Trajectory length: $K = 100$ steps, $T = 1.0$, $\Delta t = 0.01$
- Forward diffusion: $D_f = 1.0$
- Training time: $\sim 10\text{K}$ iterations ($\sim 1\text{M}$ trajectory samples)

Sampling:

- Initialization: $\mathbf{x}_0 \sim \mathcal{N}(0, D_f T \mathbf{I})$
- Integration time: $T_{\text{sample}} = 1.0$
- Timestep: $\Delta t_{\text{sample}} = 0.001$
- Effective steps: 1000 (but all parallel in hardware)

7.6 Expected Outcomes

Based on our theoretical analysis and prior work:

1. **Sample quality:** FID ≈ 100 –200 on MNIST (comparable to small diffusion models)
2. **Energy efficiency:** 10–100 \times reduction vs. digital (in simulation), approaching 0.01–0.1 J/sample in optical hardware
3. **Robustness:** Graceful degradation under 4–8 bit quantization and moderate drift
4. **Thermodynamic validation:** Strong correlation between entropy production minimization and sample quality
5. **Optimal M:** $M \approx 10$ for $T = 100$ (confirming Oguz et al.’s finding)

8 Discussion, Limitations, and Future Directions

This work contributes a unified perspective at the intersection of generative modeling, thermodynamic computing, and optical hardware by developing a rigorous theoretical framework that connects Onsager–Machlup path probabilities, entropy production, score matching, and Fokker–Planck dynamics in the context of optical generative models. It explicitly synthesizes Whitelam’s thermodynamic computing principles with the optical implementation showing how diffusion-like generative processes can be realized through physical relaxation in interferometric optical networks with realistic noise and dissipation. Building on this architecture, we provide a concrete optical Langevin design and a quantitative energy analysis, demonstrating potential 10–100 \times efficiency gains over digital implementations grounded in thermodynamic limits, while drawing clear connections to classical concepts such as Kramers escape, fluctuation–dissipation relations, and non-equilibrium statistical mechanics.

Translating the framework to physical hardware raises significant challenges, including the implementation of smooth and controllable optical nonlinearities, long-term calibration and phase stability in interferometric networks, and scalability to high-dimensional data beyond small benchmarks such as MNIST. In addition, the smooth potential-based parameterization may be less expressive than neural network denoisers for highly complex data distributions, and energy-efficient Langevin relaxation can incur long mixing times that limit sample throughput compared to fast sequential denoising approaches. Finally, in purely passive optical systems, the fluctuation–dissipation theorem constrains noise and dissipation, preventing their independent optimization without active noise injection, which would increase system complexity and energy cost.

Future work should prioritize experimental realization through the construction of an optical prototype, even at a modest scale. Initial efforts could focus on low-dimensional problems or heavily downsampled data, using commercially available fiber-optic Mach–Zehnder interferometer meshes for linear couplings and optoelectronic feedback loops to implement nonlinearities via photodiodes and phase modulators. Such a prototype would enable direct validation of energy consumption and noise characteristics, allowing meaningful comparison between theoretical thermodynamic limits and measured performance.

Beyond proof-of-concept demonstrations, several architectural extensions could substantially expand the framework’s capabilities. Introducing hidden–visible structures, analogous to restricted Boltzmann machines, would increase expressiveness by allowing additional latent degrees of freedom that evolve under the same Langevin dynamics but are not directly observed. Hybrid optical–digital systems offer another promising direction, where optical Langevin relaxation captures coarse structure in early diffusion stages while digital or learned optical denoisers refine fine details, or where optics operate in a low-dimensional latent space paired with a digital high-dimensional decoder. More speculative avenues include quantum optical implementations, leveraging squeezed states or entanglement to surpass classical noise limits, and connections to neuromorphic computing, where stochastic dynamics and local learning rules suggest natural synergies between optical linear operations and electronic spiking nonlinearities, with energy efficiency compared across digital, analog, optical, and quantum platforms.

Several theoretical extensions warrant deeper investigation. One direction concerns optimal control, specifically identifying time-dependent potentials $V_\theta(x, t)$ that minimize entropy production under fixed time or throughput constraints. Closely related is the question of whether fluctuation theorems such as the Crooks and Jarzynski relations can be directly tested in optical hardware, providing experimental access to non-equilibrium thermodynamics at the level of generative computation. Further work is also needed to clarify fundamental thermodynamic limits, including Landauer-style lower bounds on the energy cost of generative modeling, and to formalize the connection between the proposed Langevin computer and continuous-spin Boltzmann machines within a unified statistical-mechanical framework.

Beyond image generation, the proposed approach has potential applications across a range of domains where efficient sampling and stochastic dynamics are central. These include scientific computing tasks such as posterior sampling for inverse problems, materials discovery through the generation of molecular or crystalline configurations, and global optimization using parallel Langevin dynamics. More broadly, the framework points toward physical simulation and analog computing paradigms in which inference or optimization processes are embedded directly in hardware rather than emulated digitally.

9 Conclusion

We have presented a generative modeling framework that synthesizes thermodynamic computing principles with optical continuous computing, implementing diffusion models as physical Langevin relaxation in optical hardware. By training model parameters to minimize expected entropy production (maximizing reverse-trajectory likelihood), we embed the entire generative process in a passive optical system whose natural stochastic dynamics produce structured samples from noise.

Our approach bridges recent advances on thermodynamic computing and optical diffusion, offering a distinct path via Langevin relaxation rather than learned denoising. We provide mathematical formulations connecting Onsager-Machlup path probabilities, entropy production, score matching, and Fokker-Planck equations, along with detailed optical implementation designs and energy analyses.

While primarily theoretical at present, this work establishes a foundation for future experimental optical generative systems. With estimated 10–100× energy advantages over digital implementations and natural parallelism from optical propagation, thermodynamically-trained optical diffusion represents a promising direction for energy-efficient, high-throughput generative modeling at the intersection of physics, machine learning, and photonics.

This is particularly relevant due to the escalating energy demands of generative AI models that already accounts for up to 22% of U.S. household electricity consumption annually. This approach not only minimizes entropy production and heat dissipation but also paves the way for sustainable AI deployment in resource-constrained environments, potentially reducing the sector’s carbon footprint and enabling broader accessibility to high-fidelity generative technologies

References

- [1] Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [2] Dennis V Christensen, Regina Dittmann, Bernabé Linares-Barranco, Abu Sebastian, Manuel Le Gallo, Andrea Redaelli, Manan Suri, Stephan Petrović, Ulrich Ramacher, et al. 2022 roadmap on neuromorphic computing and engineering. *Neuromorphic Computing and Engineering*, 2(2):022501, 2022.
- [3] William R Clements, Peter C Humphreys, Benjamin J Metcalf, W Steven Kolthammer, and Ian A Walmsley. Optimal design for universal multiport interferometers. *Optica*, 3(12):1460–1465, 2016.
- [4] Thomas M Conte, Erik P DeBenedictis, Natesh Ganesh, Todd Hylton, John Paul Strachan, R Stanley Williams, Alexander Alemi, Lee Altenberg, Gavin E Crooks, Joseph A Faria, et al. Thermodynamic computing. In *2021 IEEE International Conference on Rebooting Computing (ICRC)*, pages 1–10. IEEE, 2021.
- [5] Gavin E Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60(3):2721, 1999.
- [6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.
- [7] Johannes Feldmann, Nathan Youngblood, Maxim Karpov, Helge Gehring, Xuan Li, Maxime Stappers, Manuel Le Gallo, Xiao Fu, Anton Lukashchuk, Arslan S Raja, et al. Parallel convolutional processing using an integrated photonic tensor core. *Nature*, 589(7840):52–58, 2021.
- [8] Johannes Feldmann, Nathan Youngblood, C David Wright, Harish Bhaskaran, and Wolfram HP Pernice. All-optical spiking neurosynaptic networks with atomic quantization. *Nature*, 589(7843):518–523, 2021.
- [9] Joseph W Goodman. *Introduction to Fourier optics*. Roberts and Company Publishers, 2005.
- [10] Peter Hänggi, Peter Talkner, and Michal Borkovec. Reaction-rate theory: fifty years after kramers. *Reviews of modern physics*, 62(2):251, 1990.
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [12] Jordan M Horowitz and Todd R Gingrich. Thermodynamic uncertainty relations constrain non-equilibrium fluctuations. *Nature Physics*, 16(1):15–20, 2020.
- [13] Hendrik Anthony Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4):284–304, 1940.
- [14] Ryogo Kubo. The fluctuation-dissipation theorem. *Reports on progress in physics*, 29(1):255, 1966.
- [15] Sabrina Maniscalco. Thermodynamic computing: A new paradigm for energy-efficient computation. In *2022 26th International Conference on Nanotechnology (IEEE-NANO)*, pages 1–2. IEEE, 2022.

- [16] Danijela Marković, Alice Mizrahi, Damien Querlioz, and Julie Grollier. Physics for neuromorphic computing. *Nature Reviews Physics*, 2(9):499–510, 2020.
- [17] David AB Miller. Are optical transistors the logical next step? *Nature Photonics*, 4(1):3–5, 2010.
- [18] Mario Miscuglio and Volker J Sorger. Photonic neural networks: a comprehensive review. *APL Photonics*, 7(3):031101, 2022.
- [19] Ilker Oguz, Niyazi Ulas Dinc, Mustafa Yildirim, Junjie Ke, Innfarn Yoo, Qifei Wang, Feng Yang, Christophe Moser, and Demetri Psaltis. Optical diffusion models for image generation. *arXiv preprint arXiv:2407.10897*, 2024.
- [20] Lars Onsager and Stefan Machlup. Fluctuations and irreversible processes. *Physical Review*, 91(6):1505, 1953.
- [21] Juan MR Parrondo, Jordan M Horowitz, and Takahiro Sagawa. Thermodynamics of information. *Nature Physics*, 11(2):131–139, 2015.
- [22] Michael Reck, Anton Zeilinger, Herbert J Bernstein, and Philip Bertani. Experimental realization of any discrete unitary operator. *Physical Review Letters*, 73(1):58, 1994.
- [23] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. *Advances in Neural Information Processing Systems*, 31, 2018.
- [24] Catherine D Schuman, Shruti R Kulkarni, Maryam Parsa, J Parker Mitchell, Prasanna Date, and Bill Kay. Opportunities for neuromorphic computing algorithms and applications. *Nature Computational Science*, 2(1):10–19, 2022.
- [25] Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75(12):126001, 2012.
- [26] Bhavin J Shastri, Alexander N Tait, Thomas Ferreira de Lima, Wolfram HP Pernice, Harish Bhaskaran, C David Wright, and Paul R Prucnal. Photonics for artificial intelligence and neuromorphic computing. *Nature Photonics*, 15(2):102–114, 2021.
- [27] Yichen Shen, Nicholas C Harris, Scott Skirlo, Mihika Prabhu, Tom Baehr-Jones, Michael Hochberg, Xin Sun, Shijie Zhao, Hugo Larochelle, Dirk Englund, et al. Deep learning with coherent nanophotonic circuits. *Nature Photonics*, 11(7):441–446, 2017.
- [28] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.
- [29] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations*, 2021.
- [30] Gordon Wetzstein, Aydogan Ozcan, Sylvain Gigan, Shanhui Fan, Dirk Englund, Marin Soljacic, Cornelia Denz, David AB Miller, and Demetri Psaltis. Inference in artificial intelligence with deep optics and photonics. *Nature*, 588(7836):39–47, 2020.
- [31] Stephen Whitelam. Generative thermodynamic computing. *arXiv preprint arXiv:2506.15121*, 2025.
- [32] Yong Zhong et al. On-chip photonic neural networks: a review. *Nanophotonics*, 12(1):1–38, 2023.

A Noise Calibration Protocol

Protocol A: Ornstein-Uhlenbeck Calibration To estimate the effective diffusion coefficient D from hardware measurements:

1. **Linearize drift:** Configure system near a stable point where drift is approximately linear: $\dot{\mathbf{x}} \approx -\mathbf{A}(\mathbf{x} - \boldsymbol{\mu})$.
2. **Collect time series:** Record $\{\mathbf{x}(t_n)\}_{n=0}^N$ at uniform intervals Δt after transient decay.
3. **Estimate drift matrix:** Compute increments $\Delta \mathbf{x}_n = \mathbf{x}_{n+1} - \mathbf{x}_n$ and fit:

$$\hat{\mathbf{A}} = - \left(\sum_n \frac{\Delta \mathbf{x}_n \tilde{\mathbf{x}}_n^\top}{\Delta t} \right) \left(\sum_n \tilde{\mathbf{x}}_n \tilde{\mathbf{x}}_n^\top \right)^{-1}, \quad (59)$$

where $\tilde{\mathbf{x}}_n = \mathbf{x}_n - \bar{\mathbf{x}}$.

4. **Estimate diffusion from residuals:** Compute residuals:

$$\mathbf{r}_n = \Delta \mathbf{x}_n + \hat{\mathbf{A}} \tilde{\mathbf{x}}_n \Delta t. \quad (60)$$

The diffusion coefficient is:

$$\hat{D} = \frac{1}{2N\Delta t \cdot N_{\text{dim}}} \sum_n \|\mathbf{r}_n\|^2, \quad (61)$$

assuming isotropic noise.

5. **Validation:** Check that residuals are uncorrelated and autocorrelation of $\mathbf{x}(t)$ matches predicted decay from $\hat{\mathbf{A}}$.

This protocol allows online calibration during operation, adapting to environmental changes - temperature or optical power.