PDE-SSM: A Spectral State Space Approach to Spatial Mixing in Diffusion Transformers

Eshed Gal
University of British Columbia
Vancouver, BC, Canada
eshedg@cs.ubc.ca

Moshe Eliasof
Ben-Gurion University of the Negev
Beersheba, Israel
eliasof@bgu.ac.il

Eldad Haber
University of British Columbia
Vancouver, BC, Canada
ehaber@eoas.ubc.ca

Abstract

The success of vision transformers—especially for generative modeling—is limited by the quadratic cost and weak spatial inductive bias of self-attention. We propose PDE-SSM, a spatial state-space block that replaces attention with a learnable convection–diffusion–reaction partial differential equation. This operator encodes a strong spatial prior by modeling information flow via physically grounded dynamics rather than all-to-all token interactions. Solving the PDE in the Fourier domain yields global coupling with near-linear complexity, delivering a principled and scalable alternative to attention. We integrate PDE-SSM into a flow-matching generative model to obtain the PDE-based Diffusion Transformer PDE-SSM-DiT. Empirically, PDE-SSM-DiT matches or exceeds the performance of matched Diffusion Transformers while substantially reducing compute. Our results show that, analogous to 1D settings where SSMs supplant attention, multi-dimensional PDE operators provide an efficient, inductive-bias-rich foundation for next-generation vision models.

1 Introduction

The Transformer architecture [Vaswani et al., 2017], with its self-attention mechanism, has become the de facto standard for state-of-the-art machine learning techniques, extending its dominance from Natural Language Processing [Touvron et al., 2023] to Computer Vision [Dosovitskiy et al., 2021, Peebles and Xie, 2023]. In generative modeling for images, the U-Net [Ronneberger et al., 2015] backbone with attention mechanism has become the standard for diffusion models over the years [Ho et al., 2020, Song et al., 2020, Nichol and Dhariwal, 2021, Dhariwal and Nichol, 2021, Rombach et al., 2022, Lipman et al., 2023]. However, after a long streak of success, the Diffusion Transformer (DiT) [Peebles and Xie, 2023] marked a significant milestone by replacing the U-Net backbone with a Transformer-based architecture, demonstrating that pure Transformer models can achieve or surpass the performance of convolutional U-Net backbones in large-scale diffusion models. Like other transformers, DiTs utilize deep networks. However, it is worth noting that DiTs are also constructed using attention mechanisms, which makes them fundamentally scalable. Be it modern U-Nets or DiTs, these are attention-based models which essentially operate by treating an image as a collection of patches and apply full weight within the patch and self-attention to model the relationship between patches [Dosovitskiy et al., 2021, Liu et al., 2021, Peebles and Xie, 2023].

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

Despite its empirical success, self-attention has two well-known drawbacks in the spatial domain [Tay et al., 2022], which are the focus of this paper. The limitations of self-attention in this case are as follows: *first*, its computational and memory requirements become increasingly challenging in 2D inputs, which have quadratically as many tokens as standard 1D inputs. Although there are a few recent ideas to attempt and reduce this cost (e.g. linear attention) [Katharopoulos et al., 2020, Shen et al., 2018], their performance generally falls short compared to standard attention-based methods. *Second*, self-attention is fundamentally permutation-invariant, meaning it does not inherently encode spatial locality or grid structure. Positional encodings are commonly introduced to remedy this limitation, and indeed, they allow Transformers to capture ordering and relative spatial information. Nevertheless, this mechanism is not as direct or specialized as the inductive bias present in convolutional neural networks (CNNs), which are explicitly designed around locality. Consequently, Transformer-based models often rely more heavily on data to learn robust spatial relationships.

Concurrently, a distinct lineage of sequence models has emerged to address these limitations in the 1D domain. State Space Models (SSMs), particularly structured SSMs (S4) [Gu et al., 2022] and their modern successors like Mamba [Gu and Dao, 2023], have achieved remarkable success on long-sequence tasks while offering compelling efficiency. These models are grounded in the theory of continuous-time systems, governed by a linear Ordinary Differential Equation (ODE):

$$\frac{dh(t)}{dt} = Ah(t) + Bu(t). \tag{1}$$

By imposing structure on the state matrix A, these models can be formulated as a recurrent or convolutional system, enabling them to capture extremely long-range dependencies in linear or near-linear time [Gu et al., 2022]. Their success raises a critical and compelling question: *how can the principles of SSMs be generalized from 1D time to N-D space to create a more efficient and spatially-aware foundation for vision models?*

Recent work has extended SSMs to graphs [Ceni et al., 2025, Eliasof et al., 2025, Lahoti et al., 2025], replacing A with the adjacency matrix. In the context of SSMs in 2D, initial forays into this area, such as Vision Mamba (ViM) [Zhu et al., 2024], have adapted 1D SSMs to images by "flattening" the 2D image domain into a 1D sequence, followed by applying the SSM model along forward and backward paths. While effective, this approach imposes an artificial causality on the spatial domain, effectively processing pixels/nodes in a scan order rather than respecting the true multi-dimensional structure of images. In this sense, it represents a heuristic adaptation of 1D principles, rather than a full generalization to the spatial setting. We provide an additional discussion of concurrent and complementary approaches in Appendix A.

Our approach. To bridge this gap, in this work, we propose a more principled approach. Concretely, we argue that the natural continuous-space generalization of an ODE that is the core of SSMs as shown in Equation (1), is a Partial Differential Equation (PDE). This observation has been studied for the control of PDEs in the context of chemical engineering [Morris, 2018, Foias et al., 1996].

We extend this idea and introduce the *PDE State-Space Model (PDE-SSM)*, an architectural block that models the evolution of a hidden state $h(t, \mathbf{x})$ in time t, where $\mathbf{x} \in \mathbb{R}^d$ — according to a learnable, general-form diffusion–convection–reaction equation, as follows:

$$\frac{\partial h}{\partial t} = \underbrace{\nabla \cdot (K \nabla h)}_{\text{Diffusion}} + \underbrace{\mathbf{b} \cdot \nabla h}_{\text{Convection}} + \underbrace{rh}_{\text{Reaction}}, \tag{2}$$

where the diffusion tensor K, convection vector field $\mathbf{b}$, and reaction term r are learnable parameters. This formulation provides a flexible and physically-grounded inductive bias; information propagates across the spatial domain via structured and well-studied mechanisms: (i) diffusion smooths and aggregates local information; (ii) convection directs the flow of information; and (iii) reaction models local feature transformations. Importantly, by implementing the solution to the PDE in Equation (2) in the Fourier domain, the underlying operator reduces the spatial mixing complexity to $\mathcal{O}(N \log N)$. While channel coupling introduces an $\mathcal{O}(NC_{\text{hid}}^2)$ cost, the overall near-linear scaling with respect to sequence length makes it highly scalable relative to the quadratic $\mathcal{O}(N^2)$ spatial cost of self-attention, while still capturing global interactions more directly than local convolutions. We integrate the PDE-SSM block into a standard transformer architecture, replacing the self-attention layers to create the PDE-based Diffusion Transformer, to obtain PDE-SSM$_{\text{DiT}}$.

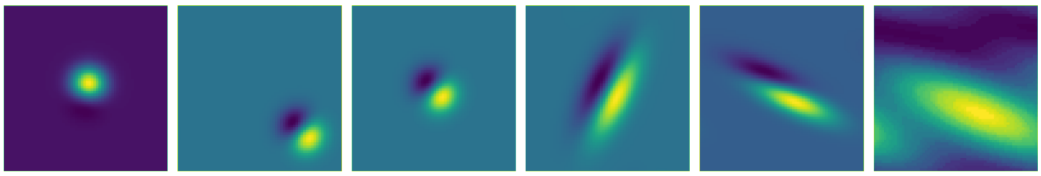

Figure 1: **Visualizing the PDE-SSM Convolutional Kernels.** By sampling the learnable parameters $\xi = (\mathcal{B}_\gamma, \zeta)$, our PDE-SSM can represent a diverse family of convolutional kernels. The examples show kernels that are (from left to right): localized, directionally blurred (anisotropic diffusion), shifted (convection), and a combination of effects. This flexibility allows our PDE-SSM model to learn a rich basis for spatial feature mixing, including non-local connections.

Our contributions are as follows:

1. We introduce PDE-SSM, a novel architectural block that provides a principled generalization of State Space Models to multi-dimensional spatial data by leveraging learnable PDEs.
2. We present PDE-SSM$_{\text{DiT}}$, a new backbone for generative vision models that utilizes the PDE-SSM for efficient and spatially-aware feature mixing, efficiently implemented via the Fourier transform.
3. Through extensive experiments on both low- and high-resolution image generation tasks, we show that PDE-SSM$_{\text{DiT}}$ is an effective generative model, achieving competitive or superior performance to the standard DiT, while being more computationally efficient, especially at higher resolutions.

2 From 1D SSM to a Spatial PDE-SSM

In this section, we develop the PDE-SSM as a principled generalization of State Space Models from 1D sequences to multi-dimensional spatial data. We begin by reformulating the 1D SSM from a differential operator perspective in Section 2.1, which provides a natural bridge to our spatial, 2D formulation in Section 2.2, followed by its efficient implementation in Section 2.3.

2.1 A Differential Operator View of SSMs

We start by utilizing the known connection between SSMs and ODEs, and their solutions [Evans, 2010]: a linear time-invariant SSM (Equation (1)), can be rewritten as the solution to an ODE:

$$\mathcal{L}h(t) = Bu(t), \quad \text{where} \quad \mathcal{L} = \frac{d}{dt} - A. \tag{3}$$

Here, $u(t)$ is the input signal, $h(t)$ is the hidden state at time t, and $A,\ B$ are learnable matrices. $\mathcal{L}$ is a first-order linear differential operator that acts on the hidden state $h(t)$. The solution can be expressed as a convolution with the Green's function (or impulse response) of the operator $\mathcal{L}$:

$$h(t) = (\mathcal{G}_A \star Bu)(t), \quad \text{where} \quad \mathcal{G}_A(t) = e^{tA}\mathbb{I}(t \geq 0). \tag{4}$$

The computation described in Equation (4) is *convolutional* — which is key to the efficiency of modern SSMs [Gu et al., 2022]. By structuring the matrix A, the convolutional kernel $\mathcal{G}_A$ can be computed efficiently, in some cases in the Fourier domain [Gu et al., 2022], allowing the model to capture long-range dependencies in near-linear time.

2.2 Generalizing to Space: The PDE-SSM Formulation

To extend the SSM framework from a 1D sequence, temporal domain to a multi-dimensional spatial domain $\mathbf{x} \in \mathbb{R}^d$, we replace the ODE with a PDE. This is a natural generalization, because PDEs are the fundamental mathematical framework for describing systems that evolve over both time and space. The resulting operator, which we call *PDE-SSM*, maps an input feature map $u(\mathbf{x})$ to an output feature map $h(\mathbf{x})$ via a spatial convolution, as follows:

$$h(\mathbf{x}) = (\mathcal{G}_\zeta \star \mathcal{B}_\gamma u)(\mathbf{x}). \tag{5}$$

Algorithm 1: **PDE-SSM Forward Pass**

Require: Input tensor $\mathbf{u} \in \mathbb{R}^{B \times C_{\text{in}} \times H \times W}$.
Require: Learnable parameters $\mathbf{R} \in \mathbb{R}^{C_{\text{hid}} \times C_{\text{in}}}, \mathbf{\Gamma}_0, \mathbf{\Gamma}_x, \mathbf{\Gamma}_y, \mathbf{K}_{xx}, \ldots, \in \mathbb{R}^{C_{\text{hid}} \times C_{\text{hid}}}, \tau$.

1. Input Embedding in Fourier Domain

$\mathbf{u}' \leftarrow \text{Conv}_{1\times1}(\mathbf{u}, \mathbf{R})$ ▷ Project to C_{hid} channels
$\hat{\mathbf{u}}' \leftarrow \text{FFT}(\mathbf{u}')$ ▷ Transform to frequency domain
Generate $\mathbf{k} = (k_x, k_y)$ as the discrete frequency grid.
$\hat{\mathcal{B}}(\mathbf{k}) \leftarrow \mathbf{\Gamma}_0 + ik_x\mathbf{\Gamma}_x + ik_y\mathbf{\Gamma}_y$ ▷ Matrix symbol for $\mathcal{B}$
$\hat{\mathbf{v}}(\mathbf{k}) \leftarrow \hat{\mathcal{B}}(\mathbf{k}) \cdot \hat{\mathbf{u}}'(\mathbf{k})$ ▷ Apply embedding (matrix–vector product)

2. Coupled PDE Evolution in Fourier Domain

$\mathbf{\Lambda}(\mathbf{k}) \leftarrow -(\mathbf{k}^\top \mathbf{K}\mathbf{k}) + \mathbf{R}_m + i(\mathbf{k} \cdot \mathbf{B})$ ▷ Assemble $C_{\text{hid}} \times C_{\text{hid}}$ matrix for each $\mathbf{k}$
$\hat{\mathcal{G}}_\zeta(\mathbf{k}) \leftarrow \exp(\tau\, \mathbf{\Lambda}(\mathbf{k}))$ ▷ Green's function symbol
$\hat{\mathbf{h}}(\mathbf{k}) \leftarrow \hat{\mathcal{G}}_\zeta(\mathbf{k}) \cdot \hat{\mathbf{v}}(\mathbf{k})$ ▷ Apply PDE operator (channelwise matrix–vector)

3. Return to Spatial Domain

$\mathbf{h} \leftarrow \text{iFFT}(\hat{\mathbf{h}})$ ▷ Back to spatial domain
return $\mathbf{h}$

Both B_γ and G_ζ are linear spatial operators and could theoretically be combined into a single operation. In practice, we apply them sequentially in the Fourier domain, since we compute the action of the matrix exponential through matrix–vector products rather than forming it explicitly.

The PDE-SSM operation, described in Equation (5), consists of two stages: (i) an input embedding $\mathcal{B}_\gamma u$ to obtain the initial state; and (ii) a spatial evolution defined by the Green's function $\mathcal{G}_\zeta$ of our learnable PDE. In what follows, we formalize these two operations.

The Embedding Operator $\mathcal{B}_\gamma$. The operator $\mathcal{B}_\gamma$ is a learnable embedding layer that transforms the input u into the hidden space. For a 2D image, which is at the focus of this paper, it reads:

$$\mathcal{B}_\gamma u = (\gamma_0 I + \gamma_x \frac{\partial}{\partial x} + \gamma_y \frac{\partial}{\partial y})(R \star u)(\mathbf{x}). \tag{6}$$

Here, R is a 1×1 (also known as pointwise) 2D convolution that mixes the input channels of u. The learnable scalars $\gamma_0, \gamma_x, \gamma_y$ then create a new feature map from a weighted sum of the channel-mixed features and their spatial gradients $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$.

The PDE Evolution and its Green's Function $\mathcal{G}_\zeta$. The core of our model is the diffusion-convection-reaction PDE from Equation (2), which governs the evolution of the state $h(\mathbf{x}, t)$ from an initial condition $h(\mathbf{x}, t = 0) = \mathcal{B}_\gamma u(\mathbf{x})$. The PDE-SSM state is obtained by solving the PDE (2) in the Fourier domain to derive an analytic expression for its Green's function. To this end, let $\hat{h}(\mathbf{k}, t)$ be the spatial Fourier transform of $h(\mathbf{x}, t)$, where $\mathbf{k} = (k_x, k_y)$ is the corresponding frequency vector. In Fourier space, the spatial derivatives $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$ transform into multiplication by $i\mathbf{k}$: $\nabla \to i\mathbf{k}$. Overall, applying the Fourier transform to the PDE yields:

$$\frac{\partial \hat{h}(\mathbf{k}, t)}{\partial t} = (i\mathbf{k})^T K (i\mathbf{k}) \hat{h}(\mathbf{k}, t) + \mathbf{b} \cdot (i\mathbf{k}) \hat{h}(\mathbf{k}, t) + r\hat{h}(\mathbf{k}, t) \tag{7}$$

$$= \underbrace{(-\mathbf{k}^T K \mathbf{k} + r + i(\mathbf{b} \cdot \mathbf{k}))}_{\lambda(\mathbf{k})} \hat{h}(\mathbf{k}, t). \tag{8}$$

This transforms the PDE into a simple, linear ODE for each frequency $\mathbf{k}$, with the solution:

$$\hat{h}(\mathbf{k}, \tau) = e^{\tau \lambda(\mathbf{k})} \hat{h}(\mathbf{k}, 0) \tag{9}$$

The term $e^{\tau\lambda(\mathbf{k})}$ is the Fourier transform of the Green's function, $\hat{\mathcal{G}}_\zeta(\mathbf{k})$, also known as the symbol of the operator, and the kernel is defined in the frequency domain by the parameters $\zeta = (K, \mathbf{b}, r, \tau)$:

$$\hat{\mathcal{G}}_\zeta(\mathbf{k}) = \exp\left(\tau(-\mathbf{k}^T K\mathbf{k} + r + i(\mathbf{b}\cdot\mathbf{k}))\right). \tag{10}$$

Theoretical Properties of PDE-SSM. Beyond computational efficiency, the PDE-SSM operator inherits structural guarantees from its PDE formulation. The diffusion term can provide structural damping. However, without explicit constraints on K, b, r, or the coupled operators, the reaction and multi-channel coupling terms may introduce eigenvalues with positive real parts. Therefore, unconditional spectral stability would require additional constraints. Expressivity is preserved through the combination of diffusion, convection, and reaction terms, which together span a rich class of kernels that includes localized filters, global translations, and nonlocal smoothers, as illustrated in Figure 1. Nonlocality is inherent as well: the integration time τ directly controls the receptive field, and as τ increases, the operator aggregates information across arbitrarily large spatial scales, paralleling the long-range memory of 1D SSMs but in the spatial domain. These guarantees make PDE-SSM a principled and theoretically grounded alternative to attention for spatial mixing.

The closed-form expression in Eq. (10) enables efficient computation while allowing the learnable parameters to encode physically meaningful inductive biases:

- **Diffusion (K):** The term $-\tau\mathbf{k}^\top K\mathbf{k}$ is a non-positive quadratic form, acting as a low-pass filter that damps high frequencies. Anisotropic K induces direction-dependent smoothing.
- **Convection ($\mathbf{b}$):** The term $i\tau(\mathbf{b}\cdot\mathbf{k})$ induces a phase shift $e^{i\tau(\mathbf{b}\cdot\mathbf{k})}$, corresponding to a spatial translation via the Fourier shift theorem, enabling directed nonlocal relationships.
- **Reaction (r):** The scalar τr uniformly scales responses across all frequencies, modeling global amplification ($r > 0$) or suppression ($r < 0$).
- **Time (τ):** The integration time modulates the strength of all effects: as $\tau \to 0$, the operator approaches the identity, while larger τ values yield increasingly global kernels.

The reaction term r is learned rather than fixed. Diffusion damps high-frequency components, while reaction provides uniform amplification or suppression across frequencies. Since K, b, and r are optimized jointly, their balance is determined by the flow-matching objective.

We illustrate different kernels obtained by our PDE-SSM convolution kernel in Equation (10) in Figure 1. We note that, in the above discussion, we have presented the computation for a single channel. In practice, one needs to consider the multi-channel case. Let C be the number of channels. In this case, in 2D, the matrix K becomes a 4D tensor of size $C \times C \times 2 \times 2$, the vector $\mathbf{b}$ becomes a $C \times C \times 2 \times 1$ tensor and r becomes a $C \times C \times 1 \times 1$ tensor. The convolution is performed for each channel pair, and the resulting outputs are aggregated across channels, similar to the way that standard multi-channel 2D convolutions are implemented in modern deep learning frameworks.

2.3 Efficient Implementation with Multi-Channel Coupling

The analytic form of the kernel in the Fourier domain in Equation (10) allows for an efficient implementation. To model rich interactions where channels interact, i.e., mix, we generalize the PDE-SSM parameters $\zeta = (K, \mathbf{b}, r, \tau)$ from scalars to matrices that operate on the channel dimension. This extension transforms the decoupled C PDE-SSM systems into a system of coupled PDEs, which can be solved simultaneously in the Fourier domain. The procedure is detailed in Algorithm 1. In this formulation, the parameters for embedding ($\mathbf{\Gamma}$'s) and evolution ($\mathbf{K}, \mathbf{B}, \mathbf{R}_m$) are now $C_{\text{hid}} \times C_{\text{hid}}$ matrices. Consequently, the operations in lines 6, 8, 9, and 10 in Algorithm 1 become matrix-matrix operations on the channel direction.

Computational Complexity. Channel-coupling impacts the computational cost. While the FFTs remain $\mathcal{O}(C_{\text{hid}} \cdot N \log N)$, the Fourier-domain operations have the following complexity: *Symbol Computation (Lines 5, 8, Algorithm 1)*—assembling the $\mathbf{\Lambda}(\mathbf{k})$ matrix for each of the N frequencies involves matrix additions and scalar–matrix products, costing $\mathcal{O}(N \cdot C_{\text{hid}}^2)$; *Matrix–Vector Products (Lines 6, 10, Algorithm 1)*—these cost $\mathcal{O}(N \cdot C_{\text{hid}}^2)$. Computing the matrix exponential explicitly at each frequency would incur an additional $\mathcal{O}(NC_{\text{hid}}^3)$ cost. For small channel dimensions, we compute the Green's function explicitly. For larger C_{hid}, we instead compute the action of the matrix

exponential using an Arnoldi approximation with a fixed Krylov dimension of 16, reducing this cost to $\mathcal{O}(NC_{\text{hid}}^2)$. The total complexity is therefore dominated by the Fourier-domain channel operations, leading to an overall cost of $\mathcal{O}(C_{\text{hid}} \cdot N \log N + N \cdot C_{\text{hid}}^2)$. The $N \cdot C_{\text{hid}}^2$ term dominates when the number of hidden channels is large. While both PDE-SSM and standard self-attention share this $\mathcal{O}(N \cdot C_{\text{hid}}^2)$ channel-mixing cost, our approach achieves substantial practical speedups by replacing attention's prohibitive $\mathcal{O}(N^2)$ spatial bottleneck with efficient $\mathcal{O}(N \log N)$ spectral mixing. This asymptotic analysis is unchanged for deeper or wider models, as model scale affects the number of layers and C_{hid} but not the structure of the PDE-SSM operator.

3 Using PDE-SSM within a Diffusion Transformer

We now discuss the viability of the PDE-SSM block in the context of Diffusion Transformer (DiT), to address the task of image generation [Lipman et al., 2023, Song et al., 2021].

Background on Flow-Matching. We focus our attention on flow-matching techniques as discussed in [Lipman et al., 2023], where the goal is to train a transformer-based network that solves the following optimization problem:

$$\min_{\theta} \mathbb{E}_{u,z} \int_0^1 \|v_\theta(tu + (1-t)z, t) - (u - z)\|^2 \, dt, \tag{11}$$

where u is the original ("clean") image, z is Gaussian noise and t is time. During training, we are given a corrupt, noisy image

$$u_t = tu + (1-t)z, \tag{12}$$

and the goal is to train a network v_θ that recovers the velocity $u - z$. This task is closely related to image denoising, because given v_θ, one can estimate the clean image u by setting:

$$u \approx u_t + (1-t)v_\theta. \tag{13}$$

During training, one uses the clean images by generating randomly corrupted images with different levels of noise, governed by the parameter t, and uses the network v_θ to recover the velocity $u - z$. The loss for a given triplet, (u, z, t), is computed by the integrand of (11). The network weights θ are trained using the AdamW [Kingma and Ba, 2014, Loshchilov and Hutter, 2019] optimizer.

Background on DiTs. Diffusion Transformer networks use attention-based blocks in the velocity network v_θ. To this end, the image is "patchified", i.e., it is separated into patches, where each patch is comprised of $k \times k$ pixels. Given an image of shape $C \times N \times N$ where C is the number of channels and N is the number of pixels in each axis, each patch is flattened to a $C \cdot k \cdot k$ vector, which is considered as the feature vector that is associated with the patch. DiTs uses a fully-connected layer on the patches, yielding a coarse image of N_{patch}^2 patches with C_{hid} channels, where $N_{patch} = N/k$ (assuming that N_{patch} is an integer). The image is then flattened to generate a 2D tensor with dimensions $C_{hid} \times N_{patch}^2$. An attention layer is then used between the patches. Overall, a standard DiT block takes the computational form of:

$$h_{j+\frac{1}{2}} = h_j + \text{ATT}(h_j)\,, \tag{14}$$

$$h_{j+1} = h_{j+\frac{1}{2}} + \text{MLP}\left(h_{j+\frac{1}{2}}\right), \tag{15}$$

where ATT is an attention layer [Vaswani et al., 2017], and MLP is a standard multilayer-perceptron. Note that the MLP operates on all patches individually (similar to a dilated convolution), while the attention models generate interactions between patches. The PDE-SSM operator itself is linear. Nonlinearities in the full architecture arise from the pointwise MLPs interleaved with PDE-SSM inside each DiT block.

Patch size and complexity. The patch size k in DiTs critically controls the computation. Large k inflates the MLP—whose dense layers scale as $C_{\text{hid}} \times C \times k \times k$—but reduces the number of patches to $N_{\text{patch}}^2 = N^2/k^2$, so attention couples fewer tokens; since attention cost grows quadratically in the number of patches, this yields a scaling proportional to $(N_{\text{patch}}^2 \times N_{\text{patch}}^2) = N_{\text{patch}}^4$.

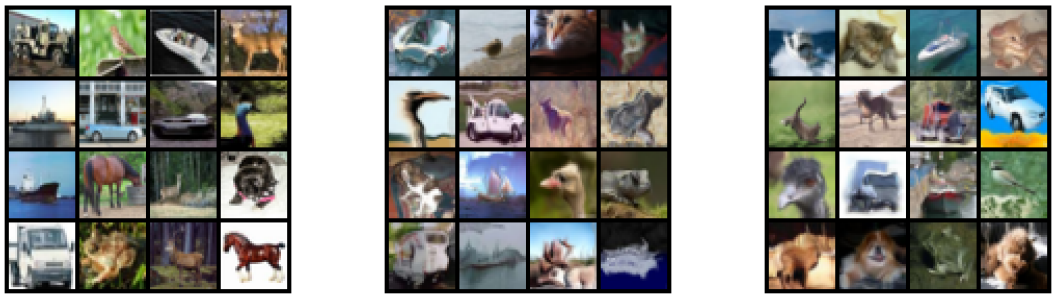

(a) Real images (b) DiT (c) PDE-SSM-DiT

Figure 2: CIFAR-10 Images: **(a)** real images; **(b)** DiT; **(c)** PDE-SSM-DiT. Visual quality is comparable, in congruence with Table 1.



(a) Real images (b) DiT (c) PDE-SSM-DiT

Figure 3: LSUN-Churches generations: **(a)** real images; **(b)** DiT; **(c)** PDE-SSM-DiT.

Conversely, a small k shrinks the MLP yet explodes the patch count, making attention dominant. Empirically, a smaller k improves generation quality [Peebles and Xie, 2023], further emphasizing attention as the bottleneck. Our PDE-SSM variant keeps the DiT block of (14) unchanged except for swapping ATT with a PDE-SSM block. By the analysis in Section 2.3, this yields a complexity of $\mathcal{O}\left(N_{\text{patch}}^2 \log N_{\text{patch}}\right)$, allowing us to choose small k without incurring prohibitive attention cost.

4 Experiments

We evaluate *PDE-SSM-DiT* by replacing DiT attention with a spatiotemporal state-space module, holding all other components, training schedules, and sampling fixed for a clean ablation. We state the guiding research questions and evaluation protocol, then report results on four benchmarks. We analyze training dynamics via a lightweight FID proxy and study scalability under matched compute as model size and resolution increase. All experiments ran on an NVIDIA RTX6000 Ada GPU. Datasets and hyperparameters are in Appendix B.

Research Questions. Our experiments seek to address the following:

- **RQ1** (*PDE-SSM vs. Attention*): At fixed training objective, schedule, and sampler, does PDE-SSM-DiT match or improve the generative quality of attention-based DiT?
- **RQ2** (*Plug-and-play replacement*): When swapping only the block type (attention $\rightarrow$ PDE-SSM) in an off-the-shelf DiT, do we retain performance without retuning hyperparameters?
- **RQ3** (*Scaling and efficiency*): How do wall-clock and asymptotic complexity evolve with image and patch resolution compared to attention?
- **RQ4** (*Domain robustness*): Are the gains consistent across class-diverse (CIFAR-10, ImageNet) and structure- or texture-centric datasets (Oxford-Flowers, LSUN-Churches)?

Evaluation Protocol. We conduct all experiments in the official DiT codebase (`https://github.com/facebookresearch/DiT`). We only replace the attention block from (14) with our PDE-SSM block defined in (2). All remaining components are held fixed, including normalization layers, positional embeddings, loss definitions, training schedules, time discretization, and the sampling procedure. We report results for three systems: DiT (unmodified DiT with attention), UNet (a standard diffusion baseline), and PDE-SSM-DiT (DiT where attention is replaced by PDE-SSM). Because end-to-end diffusion and flow-matching pipelines vary across prior work in the training domain (pixel vs. latent) [Rombach et al., 2022], time grids, learning objectives [Tong et al., 2023, Song et al., 2023], guidance mechanisms [Ho and Salimans, 2022, Dhariwal and Nichol, 2021], and samplers [Lu et al., 2022], our protocol yields a direct, plug-and-play comparison that isolates the architectural swap under matched settings. Together with the results reported below, this controlled evaluation highlights the effectiveness of PDE-SSM while minimizing confounding changes in training or inference.

Training and sampling. Unless stated otherwise, we use the default DiT hyperparameters and the standard sampling integrators from the public release cited above. Although these settings were tuned for the original DiT, we keep optimization schedules and data augmentations unchanged so that differences isolate the effect of the block replacement, directly addressing **RQ2**. We report mean metrics over three random seeds when applicable.

Metrics. We report Fréchet Inception Distance (FID) and, where relevant, mean-squared-error (MSE).

Datasets. We utilize the CIFAR-10 [Krizhevsky, 2009], CelebA-HQ [Karras et al., 2017], ImageNet64 [Deng et al., 2009], Imagenet256 [Krizhevsky et al., 2012], LSUN-Churches [Yu et al., 2015], and Oxford-Flowers [Nilsback and Zisserman, 2008]. CIFAR-10 and ImageNet stress class diversity; CelebA-HQ, Oxford-Flowers, and LSUN-Churches emphasize fine structure and global coherence, addressing **RQ4**. We provide additional details regarding the datasets and experimental settings in Appendix B.

4.1 Image Generation with PDE-SSM-DiT

CIFAR-10. CIFAR-10 is a compact and diverse benchmark. We compare a small DiT (∼30–34M parameters), a U-Net of comparable size, and PDE-SSM-DiT with a matched parameter budget. Results in Table 1 shows parity in fidelity, while qualitative samples in Figure 2 further confirm that the block swap preserves generative quality while unlocking the scaling benefits analyzed in Section 4.2 (**RQ1, RQ2**).

Table 1: PDE-SSM-DiT matches attention-based DiT and U-Net performance on CIFAR-10, indicating *quality parity at equal settings*, isolating the effect of the block swap (**RQ1, RQ2**).

Method	#Params (M)	MSE↓	FID (50K)↓
DiT	29.6	3.82×10^{-2}	4.25
U-Net	32.6	3.78×10^{-2}	4.19
PDE-SSM-DiT (Ours)	34.2	$\mathbf{3.76 \times 10^{-2}}$	**4.18**

ImageNet. On ImageNet64, PDE-SSM-DiT achieves a modest yet consistent improvement over DiT under identical training and sampling conditions, as reported in Table 2. The model also remains competitive with U-Net at comparable parameter counts, while being strictly plug-and-play (**RQ1, RQ2, RQ4**). Literature-reported results are provided for reference, though they may differ in experimental setup and sample counts. We present additional results for **ImageNet** in Appendix E.

Table 2: Unconditional Generation on ImageNet-64×64.

Method	Params (M)	FID↓
DiT	33	22.9
U-Net	32	22.2
PDE-SSM-DiT (Ours)	**31**	**22.1**

CelebA-HQ64, CelebA-HQ256, LSUN-Churches and **Oxford-Flowers:** On structure- and texture-centric datasets, PDE-SSM-DiT improves FID relative to DiT, indicating benefits for long-range coherence with small patches (**RQ1, RQ4**). Samples are shown in Figure 3, Figure 4 and Appendix F, while results are reported in Tables 3 and 4.

Table 3: FID (2048 features) on CelebA-HQ64, LSUN-Churches and Oxford-Flowers.

Method	CelebA HQ64	LSUN Churches	Oxford Flowers
DiT	9.17	55.04	61.50
U-Net	11.06	51.56	57.20
PDE-SSM-DiT (Ours)	**7.01**	**40.71**	**49.56**



(a) Real images (b) DiT (c) PDE-SSM-DiT

Figure 4: Oxford-Flowers generations: **(a)** real images; **(b)** DiT; **(c)** PDE-SSM-DiT.

Table 4: Training results on CelebA HQ 256

Architecture	Parameters	FID	MSE
DiT	32.99 M	21.08	8.02e-2
PDE-SSM-DiT(Ours)	26.48 M	18.36	5.37e-2

4.2 Computational Efficiency and Scaling

A key benefit of PDE-SSM-DiT is favorable scaling with the number of patches. While attention scales as $O(N_{\text{patch}}^4)$ when accounting for 2D tokenization, our implementation scales as $O(N_{\text{patch}}^2 \log N_{\text{patch}})$. In practice, attention encourages large patch sizes $k \times k$ to curb quadratic attention cost, shrinking the receptive field for a fixed token budget. PDE-SSM-DiT maintains efficiency at small patches, preserving long-range dependencies without the attention bottleneck (**RQ3**). Table 5 quantifies parameter counts and per-step wall-clock across resolutions, and we show additional timing results in Appendix D. In Appendix C, we show the spatial effect of the SSM model.

Main Takeaways. (i) At small patches and higher resolutions, PDE-SSM-DiT yields substantial wall-clock savings vs. attention, enabling finer tokenization without exploding cost. (ii) At very large patches, attention narrows the gap, but this comes at the expense of the global receptive field.

5 Conclusion and Discussion

Limitations. PDE-SSM is a structured spatial mixing operator and differs from self-attention in an important way. Self-attention computes content-dependent pairwise interactions between tokens, which can be beneficial when precise instance-specific correspondence between distant regions is required. In contrast, PDE-SSM imposes a spectral, translation-equivariant mixing structure whose adaptivity comes through learned channels, nonlinear layers, and depth. This makes the method attractive when global spatial mixing and favorable scaling are the main bottlenecks, but it may be less suitable for tasks requiring explicit token-to-token matching or highly instance-specific routing. We discuss broader impact of out method in Appendix G.

Conclusion. We extend state-space models from 1D to n^d and use the resulting PDE-SSM as a drop-in replacement for attention inside DiT for image generation. The key idea is to replace the 1D ODE with a space-time PDE that combines diffusion, advection, and reaction with a trainable integration horizon, producing learned nonlocal kernels whose field of view ranges from a few pixels to the entire image. Empirically, PDE-SSM-DiT matches or improves upon attention-based DiT while scaling more favorably: attention has a spatial complexity of $O(N^2)$ in tokens N, whereas our spectral solver reduces this spatial mixing to $O(N \log N)$, enabling faster training and more favorable scaling at high resolutions.

A broader takeaway is that strong generators share nonlocal mixing. UNets achieve it through multiscale downsampling, attention through pairwise correlations, and SSM or PDE-SSM through dynamical evolution that couples distant states. The emerging principle is simple: nonlocality is all you need.

Table 5: DiT vs. PDE-SSM-DiT training step runtimes (in seconds). Parameters in millions (M).

Image	Patch	Params (M)		Time (s)	
	Size	Attn	PDE-SSM	Attn	PDE-SSM
32	2	9.01	12.76	1.09×10^{-1}	2.14×10^{-2}
32	4	8.98	12.76	3.34×10^{-2}	2.40×10^{-2}
32	8	9.13	12.76	1.60×10^{-2}	2.40×10^{-2}
64	2	9.31	12.78	4.66×10^{-1}	5.58×10^{-2}
64	4	9.06	12.78	9.92×10^{-2}	5.77×10^{-2}
64	8	9.15	12.78	3.24×10^{-2}	5.84×10^{-2}
128	2	10.49	12.88	3.00×10^{0}	2.24×10^{-1}
128	4	9.35	12.88	4.70×10^{-1}	2.05×10^{-1}
128	8	9.22	12.88	1.44×10^{-1}	1.96×10^{-1}
256	2	15.21	13.27	3.42×10^{1}	8.48×10^{-1}
256	4	10.53	13.27	3.32×10^{0}	8.47×10^{-1}
256	8	9.52	13.27	5.07×10^{-1}	8.96×10^{-1}

Acknowledgments and Disclosures

ME acknowledges support from the the Israel Science Foundation through a personal grant (ISF 3062/26) and an equipment grant (ISF 3066/26), as well as support from the Israeli Ministry of Innovation, Science & Technology. The authors declare no competing interests.

References

Naman Agarwal, Daniel Suo, Xinyi Chen, and Elad Hazan. Spectral state space models. 2023. URL https://arxiv.org/abs/2312.06837. arXiv:2312.06837 [cs.LG].

Andrea Ceni, Alessio Gravina, Claudio Gallicchio, Davide Bacciu, Carola-Bibiane Schonlieb, and Moshe Eliasof. Message-passing state-space models: Improving graph learning with modern sequence modeling. *arXiv preprint arXiv:2505.18728*, 2025.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. IEEE, 2009.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:8780–8794, 2021. URL https://arxiv.org/abs/2105.05233.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021. URL https://openreview.net/forum?id=YicbFdNTTy.

Moshe Eliasof, Alessio Gravina, Andrea Ceni, Claudio Gallicchio, Davide Bacciu, and Carola-Bibiane Schönlieb. Graph adaptive autoregressive moving average models. In *Forty-second International Conference on Machine Learning*, 2025.

Lawrence C Evans. *Partial differential equations*, volume 19. American mathematical society, 2010.

Ciprian Foias, Hitay Özbay, and Allen Tannenbaum. *Robust control of infinite dimensional systems: frequency domain methods*. Springer, 1996.

Daniel Y. Fu, Elliot L. Epstein, Eric Nguyen, Armin W. Thomas, Michael Zhang, Tri Dao, Atri Rudra, and Christopher Ré. Simple hardware-efficient long convolutions for sequence modeling. In *Proceedings of the 40th International Conference on Machine Learning (ICML 2023)*, 2023. URL https://arxiv.org/abs/2302.06646.

Karan Goel, Albert Gu, Chris Donahue, and Christopher Ré. It's raw! audio generation with state-space models. In *Proceedings of the 39th International Conference on Machine Learning (ICML 2022)*, 2022. URL https://arxiv.org/abs/2202.09729.

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752, 2023.

Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations (ICLR)*, 2022.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2022. URL https://arxiv.org/abs/2207.12598.

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

Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.

Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, pages 5156–5165. PMLR, 2020.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2014.

Alex Krizhevsky. Learning multiple layers of features from tiny images., 2009.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012.

Aakash Lahoti, Tanya Marwah, Ratish Puduppully, and Albert Gu. Chimera: State space models beyond sequences, 2025. URL https://openreview.net/forum?id=Sfmk5amxFu.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.

Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems (NeurIPS)*, 35:5775–5787, 2022. URL https://arxiv.org/abs/2206.00927.

Kirsten Morris. Control of systems governed by partial differential equations. *The Control Systems Handbook*, 2018.

Alexander Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. *International Conference on Machine Learning (ICML)*, pages 8162–8171, 2021. URL https://arxiv.org/abs/2102.09672.

Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. *Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing*, pages 722–729, 2008.

William Peebles and Saining Xie. Scalable diffusion models with transformers. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. URL https://arxiv.org/abs/2212.09748.

Xuebin Qin, Hang Dai, Xiaobin Hu, Deng-Ping Fan, Ling Shao, and Luc Van Gool. Highly accurate dichotomous image segmentation. *European Conference on Computer Vision*, 2022.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

Zhuoran Shen, Mingyuan Zhang, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Efficient attention: Attention with linear complexities. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 3531–3539. IEEE, 2018.

Jiaxin Shi, Ke Alexander Wang, and Emily B. Fox. Sequence modeling with multiresolution convolutional memory. In *Proceedings of the 40th International Conference on Machine Learning (ICML 2023)*, 2023.

Jimmy T. H. Smith, Andrew Warrington, and Scott W. Linderman. Simplified state space layers for sequence modeling. In *Proceedings of the International Conference on Learning Representations (ICLR 2023)*, 2023. URL https://arxiv.org/abs/2208.04933.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

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 (ICLR)*, 2021.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023. URL https://arxiv.org/abs/2303.01469.

Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena: A benchmark for efficient transformers. In *Proceedings of the 9th International Conference on Learning Representations (ICLR 2021)*, 2021. URL https://arxiv.org/abs/2011.04006.

Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey, 2022. URL https://arxiv.org/abs/2009.06732.

Yao Teng, Yue Wu, Han Shi, Xuefei Ning, Guohao Dai, Yu Wang, Zhenguo Li, and Xihui Liu. Dim: Diffusion mamba for efficient high-resolution image synthesis. *arXiv.org*, 2024.

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2023.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Jing Nathan Yan, Jiatao Gu, and Alexander M. Rush. Diffusion models without attention. *Computer Vision and Pattern Recognition*, 2023.

Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 486–494, 2015.

Michael Zhang, Khaled K. Saab, Michael Poli, Tri Dao, Karan Goel, and Christopher Re. Effectively modeling time series with simple discrete state spaces. In *Proceedings of the 11th International Conference on Learning Representations (ICLR 2023)*, 2023.

Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision mamba: Efficient visual representation learning with bidirectional state space model. International Conference on Machine Learning, 2024.

A Additional Related Work

We discuss additional related work that is relevant to our paper.

[Agarwal et al., 2023] introduced spectral state space models that employ spectral filtering for long-range sequence dependencies, while [Fu et al., 2023] proposed hardware-efficient long convolutions for sequence modeling. Both of these works remain in the sequence modeling setting, whereas our approach generalizes state space methods to multi-dimensional spatial domains via PDE operators.

In the generative modeling domain, [Goel et al., 2022] explored the use of state-space models for audio generation, and [Shi et al., 2023] developed a multiresolution convolutional memory approach for capturing long-range dependencies. These models highlight the applicability of SSM-like architectures to generative tasks, but they are constrained to one-dimensional or heuristic extensions of sequential structures. By contrast, PDE-SSM provides a principled extension to higher-dimensional spatial data and integrates directly into diffusion transformers.

[Smith et al., 2023] introduced simplified state space layers that offer efficient implementations of SSMs and serve as alternative baselines, and [Zhang et al., 2023] proposed discrete state space formulations for time series. Both approaches focus on efficiency and simplicity in the temporal domain, whereas our work emphasizes spatial generalization and non-locality through PDE formulations. Finally, [Tay et al., 2021] proposed the Long Range Arena benchmark for evaluating the scalability of models on long-range dependencies. While prior work is primarily evaluated on sequence tasks, PDE-SSM demonstrates scalability on image generation benchmarks, showcasing the benefits of our PDE-based perspective for vision applications.

Recent works investigate alternatives to attention within diffusion models. [Teng et al., 2024] replace attention with Mamba-based selective state space layers to accelerate high-resolution diffusion, but their formulation operates on flattened 1D token sequences. [Qin et al., 2022] introduce a dichotomous refinement mechanism for segmentation, emphasizing the importance of strong spatial priors, an idea aligned with our PDE-based inductive bias. Finally, [Yan et al., 2023] propose DiffuSSM, which replaces attention with state-space layers inside diffusion models; however, DiffuSSM adapts 1D SSMs to images through sequential scanning rather than modeling true multidimensional structure. In contrast, PDE-SSM provides a learnable, physically grounded PDE operator that enables global non-local mixing with an efficient closed-form spectral implementation. Fourier Neural Operators (FNOs) and spectral convolutions also operate in the frequency domain, but typically learn unconstrained complex-valued Fourier weights, often on a truncated set of low-frequency modes. In contrast, PDE-SSM uses a spectral symbol derived from the closed-form Green's function of a learnable convection–diffusion–reaction PDE. This parameterization naturally captures directed transport through convection and global amplification or suppression through reaction.

B Datasets and Experimental Settings

We elaborate on our datasets and experimental settings.

CIFAR-10 CIFAR-10 [Krizhevsky, 2009] is a canonical benchmark in computer vision, composed of 60,000 color images at a resolution of 32×32 pixels. The dataset is evenly divided into 10 object categories, including animals (e.g., birds, cats, dogs) and vehicles (e.g., airplanes, automobiles, ships). With 50,000 training images and 10,000 test images.

ImageNet64 ImageNet64 [Deng et al., 2009] is a downsampled variant of the large-scale ImageNet dataset, where each image is resized to 64×64 pixels. It retains the rich diversity of the original ImageNet with over 1,000 categories spanning natural objects, animals, and everyday items. Despite the reduced resolution.

ImageNet256 ImageNet [Krizhevsky et al., 2012] is a large-scale visual recognition benchmark containing over 1.2 million training images across 1,000 object categories. The dataset spans a wide range of natural objects, animals, and everyday scenes.

CelebA-HQ CelebA-HQ [Karras et al., 2017] is a high-quality refinement of the CelebA dataset. The CelebA-HQ64 subset contains 30,000 face images at 64×64 resolution, emphasizing attributes such as hair style, facial expression, and background context.

LSUN-Churches The LSUN dataset [Yu et al., 2015] contains millions of large-scale images categorized into different scene types. We focus on the LSUN-Churches subset, which consists of over 125,000 images of church images captured under diverse lighting conditions, architectural styles, and viewpoints. For our experiments, we use downsampled 64×64 versions, making the dataset suitable for structure-centric generative tasks that demand global spatial coherence.

Oxford-Flowers The Oxford 102 Flowers dataset [Nilsback and Zisserman, 2008] consists of 8,189 images across 102 flower categories, with substantial intra-class variation in scale, pose, and background. For this dataset, we report the FID score on the same number of images as the dataset size.

Our experimental setting and details are summarized in Table 6.

Table 6: Experimental settings and hyperparameters for different experiments.

	CIFAR-10	ImageNet	CelebA-HQ	LSUN-Churches	Oxford-Flowers
Max iter	400,000	400,000	400,000	400,000	400,000
Learning rate	5×10^{-4}	5×10^{-4}	5×10^{-4}	5×10^{-4}	5×10^{-4}
Hidden size	384	384	384	512	512
Batch size	64	64	64	32	32
Image size	32	64	64	64	64
Layers	12	12	12	12	12
DiT Attention heads	6	6	6	6	6

C SSM Maps

To better understand the behavior of our state-space layers, we visualize the SSM maps produced by our model. These maps highlight the global structures that the SSM captures across different images, demonstrating how information propagates over large spatial extents. Figure 5 shows the SSM maps alongside their corresponding input images. In both the simple synthetic example and the real LSUN-Church [Yu et al., 2015] image, the SSM map reflects the underlying object. This illustrates the strong global bias of the SSM mechanism, which aggregates information beyond local neighborhoods and responds to the broader shape and composition of the input.



(a) Original Simple (b) SSM Map - Simple (c) Original Church (d) SSM Map - Church

Figure 5: SSM maps visualization: **(a)** real simple image; **(b)** simple image SSM map; **(c)** real church image; **(d)** church SSM map.

D Additional Timing Results

To further demonstrate the efficiency of our SSM approach, we provide timing results for different model sizes and patch resolutions. In Table 7, we report the forward-pass time for the XL models across patch sizes, for an input size of 256 and batch size 4. As the patch size decreases, our method shows increasingly significant speedups, with especially large gains at the smallest patch resolution. In Table 8, we show that even at a fixed small patch size of 2, the SSM-based variants consistently outperform their DiT counterparts across all model scales.

We additionally report end-to-end inference and training performance across multiple resolutions and channel widths in Tables 9 and 10, including peak memory and throughput. We also isolate the cost of the spatial mixing operators in Table 11. As the resolution increases, PDE-SSM-DiT exhibits increasingly favorable throughput relative to DiT.

Table 7: Iteration time vs. patch size for DiT-XL and PDE-SSM-DiT-XL models.

Model	**Patch Size**	**Time (s)**
DiT-XL	8	0.1288
DiT-XL	4	0.9633
DiT-XL	2	10.4166
PDE-SSM-DiT-XL(Ours)	8	0.1509
PDE-SSM-DiT-XL(Ours)	4	0.6315
PDE-SSM-DiT-XL(Ours)	2	2.3376

Table 8: Iteration time for DiT and PDE-SSM-DiT models.

Model	**Variant**	**Time (s)**
DiT	XL	10.4162
DiT	L	5.9323
DiT	B	2.1223
DiT	S	0.9863
PDE-SSM-DiT(Ours)	XL	2.3435
PDE-SSM-DiT(Ours)	L	1.6020
PDE-SSM-DiT(Ours)	B	0.4773
PDE-SSM-DiT(Ours)	S	0.1519

E Additional ImageNet Results

For ImageNet64, we additionally report the individual results across three random seeds in Table 12. PDE-SSM-DiT achieves the lowest mean FID and the smallest standard deviation, indicating that the improvement over the DiT and U-Net baselines is consistent across runs.

Table 9: Inference profiling across resolutions and channel widths.

Resolution	Width	Architecture	Time (s)	Peak Mem. (GB)	Throughput (img/s)
64	384	DiT	1.179×10^{-2}	0.22	339.4
64	384	PDE-SSM-DiT	1.112×10^{-2}	0.21	359.6
64	768	DiT	2.774×10^{-2}	0.69	144.2
64	768	PDE-SSM-DiT	2.745×10^{-2}	0.68	145.7
128	384	DiT	8.228×10^{-2}	0.44	48.6
128	384	PDE-SSM-DiT	3.186×10^{-2}	0.46	125.6
128	768	DiT	1.939×10^{-1}	1.11	20.6
128	768	PDE-SSM-DiT	1.197×10^{-1}	1.36	33.4
256	384	DiT	9.907×10^{-1}	1.29	4.0
256	384	PDE-SSM-DiT	1.517×10^{-1}	1.46	26.4
256	768	DiT	2.129×10^{0}	2.81	1.9
256	768	PDE-SSM-DiT	4.790×10^{-1}	4.07	8.3

Table 10: Training profiling, including the forward pass, backward pass, and optimizer step.

Resolution	Width	Architecture	Time (s)	Peak Mem. (GB)	Throughput (img/s)
64	384	DiT	3.905×10^{-2}	1.98	102.5
64	384	PDE-SSM-DiT	3.206×10^{-2}	2.33	124.8
64	768	DiT	9.794×10^{-2}	4.74	40.8
64	768	PDE-SSM-DiT	9.315×10^{-2}	6.54	42.9
128	384	DiT	3.143×10^{-1}	6.66	12.7
128	384	PDE-SSM-DiT	9.647×10^{-2}	8.10	41.5
128	768	DiT	7.032×10^{-1}	14.08	5.7
128	768	PDE-SSM-DiT	3.797×10^{-1}	21.84	10.5

We report results for our PDE-SSM method on ImageNet256, summarized in Table 13. We use the DiT-S/2 architecture and its equivalent PDE-SSM counterpart, trained for 300k iterations, and evaluate all models using classifier-free guidance. As shown in the table, our method achieves performance comparable to the DiT baseline on the challenging ImageNet dataset, while offering improved modeling efficiency through structured state-space parameterization.

Figure 6 illustrates that DiT and PDE-SSM-DiT exhibit similar convergence behavior (panel b) and reach FID scores comparable to the dataset's own internal FID (panel a). Here, internal FID refers to the baseline obtained by comparing the dataset against itself, serving as a reference for the best achievable score.

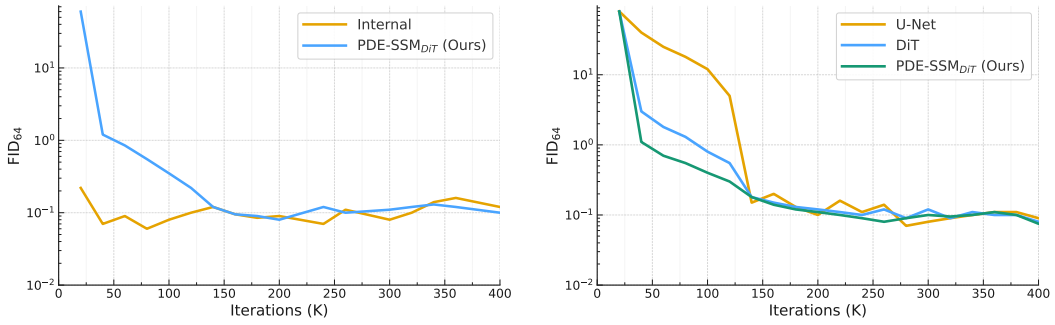

(a) PDE-SSM-DiT `FID64` vs. internal baseline.

(b) `FID64` across DiT, PDE-SSM-DiT, and U-Net.

Figure 6: ImageNet64 training. (a) The achieved FID score is consistent with the internal FID score. (b) All methods converge at a similar rate and to an FID score that is similar.

Table 11: Profiling of the isolated spatial mixing operator.

Resolution	Width	Operator	Time (s)	Peak Mem. (GB)	Throughput (img/s)
64	384	Attention	4.681×10^{-4}	0.06	8546.1
64	384	PDE-SSM	4.426×10^{-4}	0.08	9037.1
64	768	Attention	1.063×10^{-3}	0.10	3761.9
64	768	PDE-SSM	8.945×10^{-4}	0.24	4471.9
128	384	Attention	5.084×10^{-3}	0.17	786.8
128	384	PDE-SSM	9.018×10^{-4}	0.25	4435.5
128	768	Attention	1.082×10^{-2}	0.33	369.6
128	768	PDE-SSM	4.482×10^{-3}	0.87	892.5
256	384	Attention	7.423×10^{-2}	0.62	53.9
256	384	PDE-SSM	4.508×10^{-3}	0.94	887.4
256	768	Attention	1.531×10^{-1}	1.24	26.1
256	768	PDE-SSM	1.752×10^{-2}	3.39	228.3

Table 12: ImageNet64 FID across three random seeds.

Method	Run 1	Run 2	Run 3	Mean $\pm$ Std
DiT	22.6	23.2	22.9	22.9 ± 0.30
U-Net	21.7	21.8	23.1	22.2 ± 0.78
PDE-SSM-DiT	22.4	22.0	21.9	$\mathbf{22.1 \pm 0.26}$

F Additional Samples

We present qualitative results for PDE-SSM-DiT and DiT on CelebA-HQ 256 in Figure 7, where the two models exhibit similar visual quality.

G Broader Impact

While our work on PDE-SSM primarily presents a foundational architectural contribution that improves the efficiency and scaling of vision models, we recognize its broader societal implications. By reducing the computational bottleneck of spatial mixing from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$ via spectral methods, our approach lowers the barrier to entry for training high-resolution generative models and helps mitigate the carbon footprint associated with large-scale AI research. Conversely, increasing the accessibility and efficiency of powerful generative architectures like Diffusion Transformers can inadvertently facilitate the scalable creation of malicious synthetic media, such as deepfakes used for disinformation or fraud. Mitigating these dual-use risks requires concerted downstream efforts, and we strongly encourage practitioners deploying PDE-SSM-based models to implement robust safeguards, including rigorous dataset curation, post-generation safety classifiers, and cryptographic watermarking to provenance AI-generated content.

Table 13: Imagenet256 results.

Model	FID↓	sFID↓	IS↑	Precision ↑	Recall ↑
DiT	48.44	9.04	30.78	0.45	0.53
PDE-SSM-DiT-S-2 (Ours)	49.37	11.21	34.09	0.47	0.48

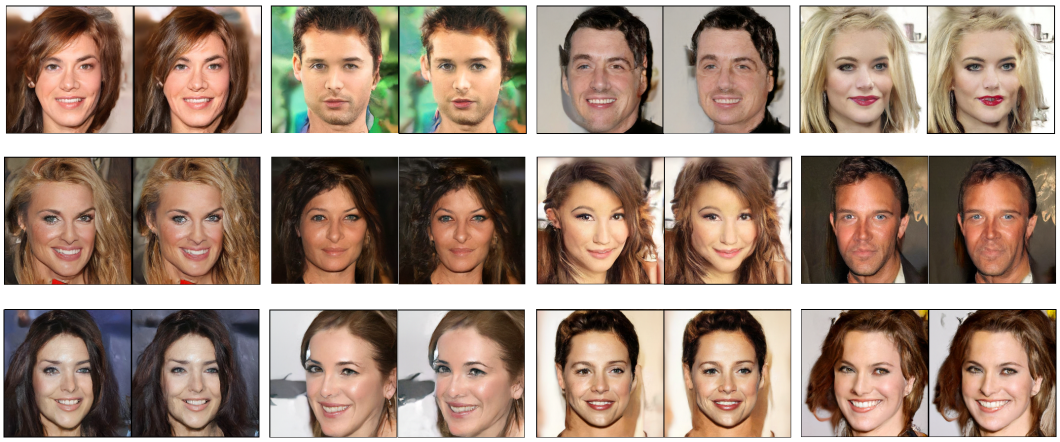

Figure 7: Sample celebrity faces generated by PDE-SSM-DiT (left) and DiT (right), trained on CelebA-HQ 256.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction state the proposed PDE-SSM block, its integration into DiT, and the controlled empirical comparisons used to support the main claims.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper discusses the main limitations of PDE-SSM relative to self-attention and the scope of the empirical evaluation in Section 5.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [N/A]

Justification: The paper contains mathematical derivations of the proposed operator, but it does not present formal theoretical claims in the form of theorems or lemmas.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper describes the PDE-SSM architecture, datasets, evaluation metrics, and training hyperparameters in section 4, with detailed experimental settings provided in Appendix B.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: Code is not included with the submission to preserve anonymity, but the implementation will be released upon acceptance; the paper provides architectural and experimental details needed to reproduce the main results.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: The paper specifies the datasets, optimizer, evaluation protocol, and main hyperparameters, with additional training details summarized in section 4 and Appendix B.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: The main comparisons in section 4 use a fixed and statistically stable evaluation protocol inherited from the DiT setup, with matched sampling and evaluation conditions across methods.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Section 4 states that all experiments were run on an NVIDIA RTX6000 Ada GPU.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research uses standard public image-generation benchmarks and follows the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: We discuss brpader impact in Appendix G.

    Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: Safeguards for responsible release are not applicable for this paper.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The paper credits the existing datasets and codebase used in the experiments through the corresponding citations and descriptions.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not release new datasets, code, or pretrained model assets as part of the submission.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing experiments or research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing experiments or research with human subjects, so IRB approval is not applicable.

    Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: The core method and experiments do not use LLMs as an important, original, or non-standard component.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.