

# Evolvable Conditional Diffusion

Zhao Wei<sup>1,2</sup>, Chin Chun Ooi<sup>1,2\*</sup>, Abhishek Gupta<sup>3</sup>, Jian Cheng Wong<sup>2</sup>, Pao-Hsiung Chiu<sup>2</sup>, Sheares Xue Wen Toh<sup>2</sup> and Yew-Soon Ong<sup>1,4</sup>

<sup>1</sup>Centre for Frontier AI Research, Agency for Science, Technology and Research, Singapore

<sup>2</sup>Institute of High Performance Computing, Agency for Science, Technology and Research, Singapore

<sup>3</sup>School of Mechanical Sciences, Indian Institute of Technology, Goa, India

<sup>4</sup>College of Computing and Data Science, Nanyang Technological University, Singapore

{wei\_zhao, ooicc}@cfar.a-star.edu.sg, abhishekgupta@iitgoa.ac.in, {wongj, chiuph, sheares\_toh}@ihpc.a-star.edu.sg, asysong@ntu.edu.sg

## Abstract

This paper presents an evolvable conditional diffusion method such that black-box, non-differentiable multi-physics models, as are common in domains like computational fluid dynamics and electromagnetics, can be effectively used for guiding the generative process to facilitate autonomous scientific discovery. We formulate the guidance as an optimization problem where one optimizes for a desired fitness function through updates to the descriptive statistic for the denoising distribution, and derive an evolution-guided approach from first principles through the lens of probabilistic evolution. Interestingly, the final derived update algorithm is analogous to the update as per common gradient-based guided diffusion models, but without ever having to compute any derivatives. We validate our proposed evolvable diffusion algorithm in two AI for Science scenarios: the automated design of fluidic topology and meta-surface. Results demonstrate that this method effectively generates designs that better satisfy specific optimization objectives without reliance on differentiable proxies, providing an effective means of guidance-based diffusion that can capitalize on the wealth of black-box, non-differentiable multi-physics numerical models common across Science.

## 1 Introduction

Diffusion models have emerged as a powerful class of deep generative models, demonstrating remarkable performance in a variety of domains such as image synthesis [Rombach *et al.*, 2022], audio generation [Kong *et al.*, 2021], biological sequence generation [Avdeyev *et al.*, 2023], and engineering design problems [Liu and Thuerey, 2024; Xu *et al.*, 2025].

While the original unconditional diffusion model generates samples randomly, conditional diffusion models are a common variant that have been proposed to enable generation

with a bias towards specific users’ requirements (e.g., aesthetics and/or specific engineering performance criteria) [Yang *et al.*, 2023]. For example, molecular linker design leverages conditioning on target protein pockets to advance drug discovery [Igashov *et al.*, 2024]. However, conditional diffusion models require paired data-sets to be available *a priori* for training, which can be difficult to obtain in large quantities.

In contrast, guided diffusion models and their variants [Dhariwal and Nichol, 2021; Bansal *et al.*, 2023; Mazé and Ahmed, 2023; Chen *et al.*, 2024] have been proposed whereby the gradient of a regressor or classifier is used to steer the denoising trajectory such that samples with specific requirements are preferentially generated. Critically, guided diffusion models offer flexibility in conditional generation as they enable a pre-trained diffusion model to generate samples which satisfy any post-hoc requirement, a notable advantage as pre-trained diffusion models continue to proliferate and scale in the era of foundation models [Li *et al.*, 2024]. This has already led to several exciting advances in AI for Science, e.g., MatterGen [Zeni *et al.*, 2025] demonstrates the ability of guided diffusion to directly generate novel, stable, synthesizable inorganic materials with specific target properties. The integration of such automated candidate generation strategies with autonomous laboratories of the future approach is key to achieving the vision of end-to-end autonomous scientific discovery [Szymanski *et al.*, 2023].

Although promising, practical applications of guided diffusion models still face certain challenges. One key issue is the need for models that can evaluate the desired performance to be differentiable in order for the gradient to be easily computed and incorporated during the denoising process. However, this precludes the use of many established multi-physics numerical models that have been developed for diverse scientific applications as the majority of such models are non-differentiable. These models are essentially black-boxes in that users typically do not have the capability to interact in an intrusive manner with the solver, e.g., molecular/protein descriptors [Gainza *et al.*, 2020; Ghiringhelli *et al.*, 2015] and physics-based simulators [Yuan *et al.*, 2023]. Hence, this disconnect between the need for gradients and the limitations of current state-of-the-art high-fidelity solvers has been a major barrier to the applicability of performance-guided diffusion

\*Corresponding author

models in enabling major scientific transitions.

The methods of evolutionary computation offer a unique pathway to overcome this barrier. In contrast to local gradient-based guidance methods, the generative processes of evolution are inherently derivative-free and exploratory, driven by random genetic variation and guided by the principle of survival of the fittest. Critically, evolutionary computation’s independence from gradient information makes it uniquely suited for exploiting the repertoire of existing black-box solvers established in many scientific domain, facilitating automated scientific discovery as described above while retaining compliance with existing known physics through evaluation and verification with state-of-the-art numerical solvers. Being a search-oriented approach, evolutionary computation is particularly effective for global exploration in complex, multi-modal landscapes, enabling it to tackle discrete, discontinuous, and high-dimensional problems as is common in many scientific domains, e.g., meta-surface design, a pressing challenge in communications and next-generation semiconductors. Moreover, its robustness in handling multiple guidance signals (fitness functions) and noisy environments reinforces its applicability to complex design challenges of today [Bali *et al.*, 2020].

It is worth highlighting that the stochastic nature of evolutionary computation also promotes diversity and can unveil unexpected and innovative solutions, with prior literature demonstrating how this can lead to the emergence of novel biological structures [Miikkulainen and Forrest, 2021]. Hence, the incorporation of in silico evolution in generative AI methods like diffusion models also engenders a similar potential to produce innovative, out-of-distribution solutions, making it an exciting engine for scientific discovery [Wong *et al.*, 2024; Lyu *et al.*, 2024; Lehman *et al.*, 2020]. Overall, evolutionary computation is a compelling alternative to traditional, gradient-based generative AI methods.

Given the gradient-free capabilities of evolutionary computation [Le *et al.*, 2013; Sung *et al.*, 2023], we seek to develop a gradient-free approach to achieve guided generation through the lens of *evolvability* [Valiant, 2009] of the diffusion model as shown in Figure 1. This method enables generation of designs that preferentially align with desired objectives even if the evaluation tool (e.g., solver) is non-differentiable. The contributions are summarized as follows:

- We propose an evolvable conditional diffusion method inspired by evolutionary algorithms. This approach treats the generation process as a black-box optimization problem, whereby the probabilistic distribution obtained from the pre-trained diffusion model is evolved to favor designs that maximize certain performance criteria. Notably, the derived update algorithm can be seen to be analogous to the update rule in conventional gradient-based guided diffusion under specific assumptions.
- The proposed method eliminates the need for differentiable models in the guided diffusion. Instead, the gradient of the fitness function is directly estimated from samples drawn from the evolved distribution, with corresponding fitness values evaluated using non-differentiable solvers.

- We successfully apply the proposed approach to two generative design problems as examples of how this method can be useful in the general area of AI for Science, i.e., the design of fluidic channel topology and frequency-selective meta-surface.

## 2 Background and Related Works

### 2.1 Diffusion Models

Diffusion models [Ho *et al.*, 2020; Song *et al.*, 2021; Song and Ermon, 2019; Song and Ermon, 2020; Song *et al.*, 2020] are a class of probabilistic generative models that learn to generate new samples by iteratively removing noise (denoising) from a random input. They have demonstrated advantages over other state-of-the-art methods like generative adversarial networks (GANs) across various tasks such as computer vision [Croitoru *et al.*, 2023], natural language processing [Li *et al.*, 2022], and topology design [Mazé and Ahmed, 2023].

Taking the famous denoising diffusion probabilistic model (DDPM) [Ho *et al.*, 2020] as an example, the training process consists of two main phases: a forward noising process and a reverse denoising process. In the forward noising process, data is progressively perturbed by adding Gaussian noise in small steps, eventually transforming the data into pure noise as per Eq. (1)

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}) \quad (1)$$

where  $\beta_t \in (0, 1)$  is a hyperparameter selected prior to model training. In contrast, the denoising process gradually removes the noise by running a learnable Markov chain to generate new data as illustrated in Eq. (2)

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t)) \quad (2)$$

where  $\boldsymbol{\mu}_\theta(\mathbf{x}_t)$  and  $\boldsymbol{\Sigma}_\theta(\mathbf{x}_t)$  are the mean and covariance of the denoising process, and are predicted through the trained neural network parameterized by  $\theta$ . In some instances, the covariance  $\boldsymbol{\Sigma}_\theta(\mathbf{x}_t)$  can be simplified as a constant [Ho *et al.*, 2020]. Once learned, new samples (e.g., images) can be generated by sampling from  $\mathcal{N}(0, 1)$  and carrying out a denoising process according to Eq. (2).

### 2.2 Guidance Methods

In practice, vanilla diffusion models produce plausible samples, but with no regard to specific user requirements or objectives. Thus, guidance in diffusion models is commonly introduced to steer the generative process towards specific desired outcomes during the denoising process [Dhariwal and Nichol, 2021]. In this approach, a separate model (e.g., neural network) which is able to predict the objective value of the condition is assumed to be available, such that its gradients can be used to steer the denoising process towards a target criteria. Some related research can be found in [Wallace *et al.*, 2023; Kim *et al.*, 2022].

For a diffusion model with an unconditional denoising process  $p_\theta(\mathbf{x}_t|\mathbf{x}_{t+1})$ , [Ho *et al.*, 2020] shows that, under reasonable assumptions, we can model the distribution as a Gaussian  $p_\theta(\mathbf{x}_t|\mathbf{x}_{t+1}) = \mathcal{N}(\mathbf{x}_t; \boldsymbol{\mu}_\theta(\mathbf{x}_{t+1}), \boldsymbol{\Sigma}_\theta(\mathbf{x}_{t+1}))$ , where the

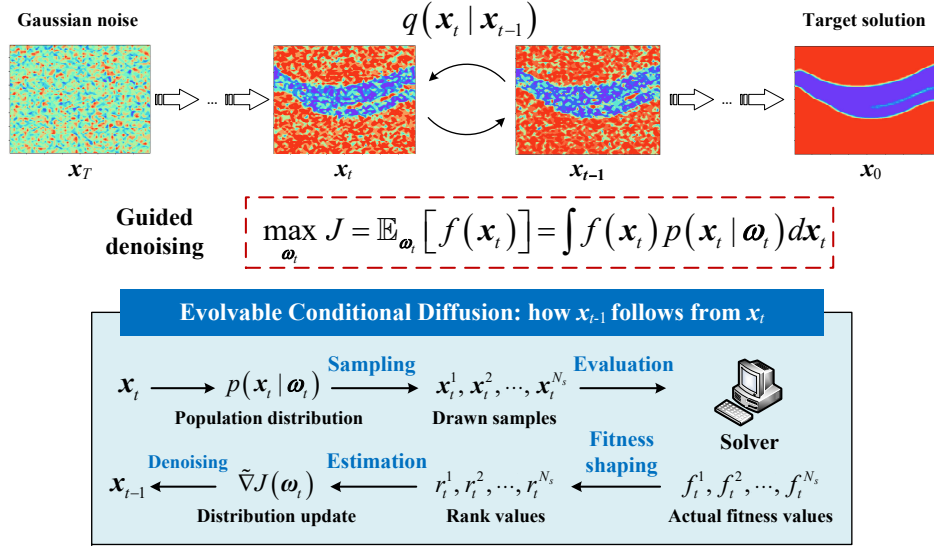


Figure 1: Framework of the proposed evolvable conditional diffusion method. The entire process is framed through the lens of probabilistic evolvability, transitioning from  $x_T$  to  $x_0$  via a sequence of mutations [Valiant, 2009]. Notably, our derived update algorithm directly utilizes evolved samples and corresponding fitness values, eliminating the need for differentiable evaluation models during guidance and extending its applicability to non-differentiable solvers.

mean  $\mu_{\theta}(x_{t+1})$  and covariance  $\Sigma_{\theta}(x_{t+1})$  are modeled by two neural networks. Crucially, a guidance process can be incorporated to post hoc bias the denoising process, and [Mazé and Ahmed, 2023] shows that the denoising guidance can in turn be formulated as a gradient-based update to the denoising Gaussian’s mean as

$$\mu_{\theta}^c(x_t) = \mu_{\theta}(x_t) + \alpha \Sigma_{\theta}(x_t) \nabla_{x_t} f(x_t) \quad (3)$$

where  $f$  is a differentiable guidance function (e.g., a regressor) that evaluates the denoising outcome; and  $\alpha$  is the gradient scaling factor.

Note that the gradient-based guidance mechanisms typically require a differentiable proxy. However, models with black-box characteristics, e.g., the finite element analysis (FEA) software Abaqus and computational fluid dynamics (CFD) software Ansys Fluent commonly used for structural and fluid analysis respectively, can evaluate objective functions but are inherently non-differentiable. This lack of differentiability makes them incompatible with gradient-based guidance mechanisms and limits their integration into classifier/regressor-guided diffusion frameworks. To address this limitation, we propose an evolvable conditional diffusion method to eliminate the need for differentiable proxies.

### 3 Gradient-Free Conditional Diffusion Algorithm

#### 3.1 Theoretical Background

In this section, we derive a gradient-free guidance approach from first principles through the lens of probabilistic, population-based evolutionary computation. We formulate the guidance as an optimization problem where the denoising process is evolved to advance in the direction that

maximizes the fitness function  $f$ , thereby updating the parameters of the denoising distribution. In probabilistic evolution, the optimization problem can be reformulated as maximizing the following expected fitness under the underlying population distribution model [Ollivier *et al.*, 2017; Gupta *et al.*, 2022]

$$J(\omega) = \mathbb{E}_{\omega}[f(x_t)] = \int f(x_t) p(x_t | \omega) dx_t \quad (4)$$

where  $f$  is the black-box guidance function;  $\omega = (\mu_{\theta}, \Sigma_{\theta})$  represents the denoising distribution parameters, which are evolved to produce samples that are more desirable as assessed by the fitness function;  $x_t$  represents individual samples drawn from the denoising distribution. In our proposed gradient-free approach where we maximize  $J(\omega)$ ,  $f$  no longer has the same requirements for differentiability as the equivalent  $f$  in Eq. (3).

**Proposition 1.** *Given the optimization objective  $J(\omega) = \mathbb{E}_{\omega}[f(x_t)]$ ;  $\omega = (\mu_{\theta}, \Sigma_{\theta})$ . The gradient-free update rule to the mean of the denoising distribution is given by*

$$\mu_{\theta}^c = \mu_{\theta} + \alpha \tilde{\nabla}_{\mu_{\theta}} J(\omega) \quad (5)$$

where  $\tilde{\nabla}_{\mu_{\theta}} J(\omega)$  denotes the natural gradient of  $J(\omega)$  with respect to the mean  $\mu_{\theta}$ . Then, when  $\|\Sigma_{\theta}\| \rightarrow 0$ , we have  $\tilde{\nabla}_{\mu_{\theta}} J(\omega) = \Sigma_{\theta} \nabla_{x_t} f(x_t)$  which directly corresponds to the gradient-based update in Eq. (3).

*Proof.* The natural gradient to the objective  $J(\omega)$  can be formalized as the solution to the constrained optimization problem [Wierstra *et al.*, 2014]

$$\max_{\delta\omega} J(\omega + \delta\omega) \approx J(\omega) + \delta\omega^{\top} \nabla_{\omega} J(\omega) \quad (6a)$$

$$s.t. D(\omega + \delta\omega || \omega) = \epsilon \quad (6b)$$

where  $D$  is the Kullback-Leibler divergence between two probability distributions, and  $\epsilon$  is a small increment size. For  $\delta\omega \rightarrow 0$ , the solution to this can be found using a Lagrangian multiplier, yielding the necessary condition

$$\mathbf{F}\delta\omega = \phi\nabla_{\omega}J(\omega) \quad (7)$$

where  $\mathbf{F} = \mathbb{E}[\nabla_{\omega} \log p(\mathbf{x}_t|\omega) \nabla_{\omega} \log p(\mathbf{x}_t|\omega)^{\top}]$  is the Fisher information matrix, and for some constant  $\phi > 0$ . The direction of the natural gradient  $\tilde{\nabla}_{\omega}J(\omega)$  is thus given by  $\delta\omega$

$$\tilde{\nabla}_{\omega}J(\omega) = \mathbf{F}^{-1}\nabla_{\omega}J(\omega). \quad (8)$$

Given further that  $p(\mathbf{x}_t|\omega)$  is a Gaussian distribution, the Fisher information matrix for  $\mu_{\theta}$  is given by

$$\mathbf{F}_{\mu_{\theta}} = \mathbb{E}[\nabla_{\mu_{\theta}} \log p(\mathbf{x}_t|\omega) \nabla_{\mu_{\theta}} \log p(\mathbf{x}_t|\omega)^{\top}] \quad (9a)$$

$$= \mathbb{E}[(\Sigma_{\theta}^{-1}(\mathbf{x}_t - \mu_{\theta}))(\Sigma_{\theta}^{-1}(\mathbf{x}_t - \mu_{\theta}))^{\top}] \quad (9b)$$

$$= \Sigma_{\theta}^{-1} \mathbb{E}[(\mathbf{x}_t - \mu_{\theta})(\mathbf{x}_t - \mu_{\theta})^{\top}] \Sigma_{\theta}^{-1} \quad (9c)$$

$$= \Sigma_{\theta}^{-1} \Sigma_{\theta} \Sigma_{\theta}^{-1} \quad (9d)$$

$$= \Sigma_{\theta}^{-1}. \quad (9e)$$

Using this relation, the natural gradient to the objective  $J(\omega)$  can be expressed as

$$\tilde{\nabla}_{\mu_{\theta}}J(\omega) = \mathbf{F}_{\mu_{\theta}}^{-1}\nabla_{\mu_{\theta}}J(\omega) \quad (10a)$$

$$= \Sigma_{\theta} \nabla_{\mu_{\theta}} \int f(\mathbf{x}_t)p(\mathbf{x}_t|\omega)d\mathbf{x}_t. \quad (10b)$$

Note that as  $\|\Sigma_{\theta}\| \rightarrow 0$ , the density  $p(\mathbf{x}_t|\omega)$  approaches a Dirac delta function. Therefore, the translation property of the Dirac delta function allows us to simplify Eq. (10) as

$$\tilde{\nabla}_{\mu_{\theta}}J(\omega) = \Sigma_{\theta} \nabla_{\mu_{\theta}} \int f(\mathbf{x}_t)p(\mathbf{x}_t|\omega)d\mathbf{x}_t \quad (11a)$$

$$= \Sigma_{\theta} \nabla_{\mathbf{x}_t} f(\mathbf{x}_t). \quad (11b)$$

Consequently, the derived gradient-free update mathematically corresponds to the gradient-based update used in conventional guided diffusion models, i.e., Eq. (3).  $\blacksquare$

### 3.2 Implementation Details

In practice, the gradient-free update term in Eq. (5) is estimated by Monte Carlo sampling of a population of candidate solutions. By the Leibniz integral rule, we express  $\tilde{\nabla}_{\mu_{\theta}}J(\omega)$  in Eq. (10) as

$$\tilde{\nabla}_{\mu_{\theta}}J(\omega) = \Sigma_{\theta} \nabla_{\mu_{\theta}} \int f(\mathbf{x}_t)p(\mathbf{x}_t|\omega)d\mathbf{x}_t \quad (12a)$$

$$= \Sigma_{\theta} \int f(\mathbf{x}_t) \nabla_{\mu_{\theta}} p(\mathbf{x}_t|\omega) d\mathbf{x}_t. \quad (12b)$$

By applying the log-likelihood trick, i.e.,  $\nabla_{\mu_{\theta}} p(\mathbf{x}_t|\omega) = \nabla_{\mu_{\theta}} \log p(\mathbf{x}_t|\omega) p(\mathbf{x}_t|\omega)$ , Eq. (12) is reformulated as

$$\tilde{\nabla}_{\mu_{\theta}}J(\omega) = \Sigma_{\theta} \int f(\mathbf{x}_t) \nabla_{\mu_{\theta}} \log p(\mathbf{x}_t|\omega) p(\mathbf{x}_t|\omega) d\mathbf{x}_t \quad (13a)$$

$$\approx \Sigma_{\theta} \frac{1}{N_s} \sum_{i=1}^{N_s} f(\mathbf{x}_t^i) \nabla_{\mu_{\theta}} \log p(\mathbf{x}_t^i|\omega) \quad (13b)$$

$$= \Sigma_{\theta} \frac{1}{N_s} \sum_{i=1}^{N_s} f(\mathbf{x}_t^i) \Sigma_{\theta}^{-1} (\mathbf{x}_t^i - \mu_{\theta}) \quad (13c)$$

$$= \frac{1}{N_s} \sum_{i=1}^{N_s} f(\mathbf{x}_t^i) (\mathbf{x}_t^i - \mu_{\theta}). \quad (13d)$$

Note that Eq. (13b) is approximated by the Monte Carlo estimate of the search gradient, which is obtained without ever having to compute the derivatives of the guidance signal  $f$ .

Since Eq. (13d) is sensitive to the magnitude and extreme values of the fitness function, a fitness shaping [Wierstra *et al.*, 2014] approach is applied to transform the actual fitness values into rank values. This replaces  $f$  with a rank-based function  $r$  as

$$r(\mathbf{x}_t) = a + b \cdot \text{rank}(f(\mathbf{x}_t)) \quad (14)$$

where  $\text{rank}(\mathbf{x}_t)$  denotes the rank values of all samples in  $\mathbf{x}_t$ , determined based on the actual fitness value  $f(\mathbf{x}_t)$ , and a lower rank value indicates a smaller actual fitness value;  $a$  and  $b$  are constants used for rescaling. It is evident that Eq. (14) controls the range of fitness values and thereby ensures the invariance of our method to order-preserving fitness transformations. As the fitness function  $f$  is now replaced by a rank-based function  $r$ , the natural gradient in Eq. (13d) can be further expressed as

$$\tilde{\nabla}_{\mu_{\theta}}J(\omega) \approx \frac{1}{N_s} \sum_{i=1}^{N_s} r(\mathbf{x}_t^i) (\mathbf{x}_t^i - \mu_{\theta}) \quad (15a)$$

$$= \frac{1}{N_s} \mathbf{M}_t \mathbf{R}_t \quad (15b)$$

where  $\mathbf{M}_t = [m_{ni}]_{N \times N_s}$  with  $m_{ni} := \mathbf{x}_{tn}^i - \mu_{\theta}^n$ , where  $N$  is the dimensionality of the samples, and  $N_s$  is the number of samples drawn from the current distribution  $(\mu_{\theta}(\mathbf{x}_t), \Sigma_{\theta}(\mathbf{x}_t))$ ;  $\mathbf{R}_t = [r_t^1, \dots, r_t^{N_s}]^{\top}$ .

### 3.3 Algorithm

The steps of the evolvable conditional diffusion method are outlined in Algorithm 1, assuming that a pre-trained diffusion model parameterized by  $\theta$  that predicts  $(\mu_{\theta}, \Sigma_{\theta})$  is available.

The initial distribution for the evolutionary optimization problem is given by  $\omega = \omega_t = (\mu_{\theta}(\mathbf{x}_t), \Sigma_{\theta}(\mathbf{x}_t))$ , and  $N_s$  samples are drawn from the distribution  $\mathcal{N}(\mu_{\theta}(\mathbf{x}_t), \Sigma_{\theta}(\mathbf{x}_t))$  in Line 4. The actual fitness values for these samples are computed accordingly in Line 5. To decrease the influence of extreme values, rank-based fitness shaping as per Eq. (14) is applied to transform the actual

---

**Algorithm 1** Evolvable conditional diffusion method
 

---

**Input:** number of denoising step  $T$ , number of samples for gradient estimation  $N_s$ , dimensionality of the noised sample  $N$ , gradient scaling factor  $\alpha$ , unconditional diffusion model  $(\mu_\theta, \Sigma_\theta)$ , and model for evaluating the fitness  $f(\cdot)$

**Output:** final guided sample  $x_0$

```

1:  $x_T \leftarrow \text{sampling from } \mathcal{N}(0, I)$ 
2: for  $t \in [T, T-1, \dots, 1]$  do
3:   for  $i \in [1, 2, \dots, N_s]$  do
4:      $x_t^i \leftarrow \text{sampling from } \mathcal{N}(\mu_\theta(x_t), \Sigma_\theta(x_t))$ 
5:      $f_t^i \leftarrow f(x_t^i)$ 
6:   end for
7:    $r_t^1, r_t^2, \dots, r_t^{N_s} \leftarrow \text{FitnessShaping}(f_t^1, f_t^2, \dots, f_t^{N_s})$ 
8:    $M_t \leftarrow [m_{ni}]_{N \times N_s}$  with  $m_{ni} := x_{t_n}^i - \mu_\theta^n$ 
9:    $R_t \leftarrow [r_t^1, \dots, r_t^{N_s}]^\top$ 
10:   $\mu_\theta^c(x_t) \leftarrow \mu_\theta(x_t) + \alpha M_t R_t$ 
11:   $x_{t-1} \leftarrow \text{sampling from } \mathcal{N}(\mu_\theta^c(x_t), \Sigma_\theta(x_t))$ 
12: end for
13: return  $x_0$ 
    
```

---

fitness values to rank values. A single gradient ascent step at  $\omega_t$  based on Eq. (15b) is then performed with a gradient scaling factor  $\alpha$  to update the denoising distribution in the direction of higher fitness values in Line 10. The conditional denoising mean  $\mu_\theta^c(x_t)$  is now obtained. With conditional sampling, a new sample for the next denoising step is generated in Line 11. Subsequently, the distribution parameters for the next denoising step  $\omega_{t-1} = (\mu_\theta(x_{t-1}), \Sigma_\theta(x_{t-1}))$  are determined, and the initial  $\omega$  for this step is then set using  $\omega_{t-1}$ . This process is repeated iteratively until the final denoising state  $x_0$  is reached.

As the updates are approximated directly through the use of drawn samples from the evolved distribution and corresponding rank values, this method is applicable even when the gradients of the fitness function evaluator are difficult to obtain (e.g., in the use of non-differentiable, black-box numerical solvers).

## 4 Experiments

### 4.1 Design of Fluidic Channel Topology

**Problem Statement.** In this section, we validate the effectiveness of our proposed evolvable conditional diffusion on the design of fluidic channel topology, a fundamental problem with implications on the performance of heat exchangers across multiple domains, including semiconductor chip cooling, battery pack liquid cooling, and chemical reactor cooling. Specifically, we seek to optimize the geometry of fluid channels to minimize pressure drop across the inlet and outlet ( $\Delta p$ ) as pressure drop is a proxy for the amount of energy required to drive flow through the system, and designs with reduced pressure drop typically have superior performance in terms of energy efficiency.

This pressure drop can be obtained conventionally by solving the steady-state incompressible Navier-Stokes (N-S) equations as per Eq. (16) [Wei *et al.*, 2023]

$$\nabla \cdot \vec{u} = 0 \quad (\vec{u} \cdot \nabla) \vec{u} = \frac{1}{\text{Re}} \nabla^2 \vec{u} - \nabla p \quad (16)$$

where  $\vec{u}$  represents the velocity vector;  $p$  denotes the pressure;  $\text{Re} = 500$  is the Reynolds number under investigation.

**Implementation Details.** For demonstration of the evolvable conditional diffusion model, we first assume availability of a black-box solver that can provide evaluations of the pressure drop when provided a specific channel topology and a pre-trained diffusion model that can generate random fluidic channel topologies. Examples of such black-box solvers include CFD numerical solvers such as Ansys Fluent. During the denoising process, evaluations of the objective (i.e.,  $\Delta p$ ) corresponding to the final denoising state  $x_0$  can be obtained by calling this black-box solver. Figure 2 presents sample CFD results for a representative fluidic topology. As a proof-of-concept, a regressor was trained to provide the black-box fitness evaluation using a paired dataset comprising topology designs and their corresponding  $\Delta p$ . Using the pre-trained diffusion model, guidance is applied for the second half of the denoising process with 30 samples evaluated per denoising step for gradient estimation. The entire denoising process consists of 100 steps, and the input design representation has a spatial resolution of  $64 \times 64$ .

**Results.** We generate 1000 samples for assessment of the effectiveness of our proposed method. The histogram of  $\Delta p$  obtained with and without gradient-free guidance during denoising across these 1000 samples is presented in Figure 3 (a). The results indicate that the distribution of  $\Delta p$  of the samples generated with guidance is significantly lower compared to the baseline generated samples. Figure 3 (b) further illustrates that all 1000 samples obtained via conditional diffusion have reduced  $\Delta p$  relative to the corresponding sample generated without guidance in the denoising process. Furthermore, increasing the gradient scaling factor  $\alpha$  clearly biases the denoising process towards designs with even lower  $\Delta p$ , demonstrating the potential for generating designs which better satisfy one's criteria. The results highlight the effectiveness of the proposed evolvable conditional diffusion in scenarios where the solver is non-differentiable.

Two representative samples from the 1000 generated samples are presented in Figure 4 to illustrate the final images (i.e.,  $x_0$ ) generated with and without guidance. Notably, when  $\alpha = 5$ , the outlet area is observed to be widened, causing a reduction in  $\Delta p$  compared to the baseline design.  $\Delta p$  curves for the final 50 denoising steps of conditional and unconditional diffusion are presented in Figure 5. When the guidance is applied during denoising,  $\Delta p$  decreases as denoising progresses, with a more significant drop observed for larger  $\alpha$ . In contrast,  $\Delta p$  in the baseline denoising process remains fairly consistent, further illustrating the effectiveness of the proposed evolvable conditional diffusion.

To further demonstrate the effectiveness of the proposed approach, we applied guidance for a smaller fraction of the denoising steps in the generation process. Specifically, the gradient-free guidance was applied only for the last 10 steps (as opposed to the previously presented 50 steps). The histogram of  $\Delta p$  and difference in  $\Delta p$  between samples generated with and without guidance are presented for 1000 gen-

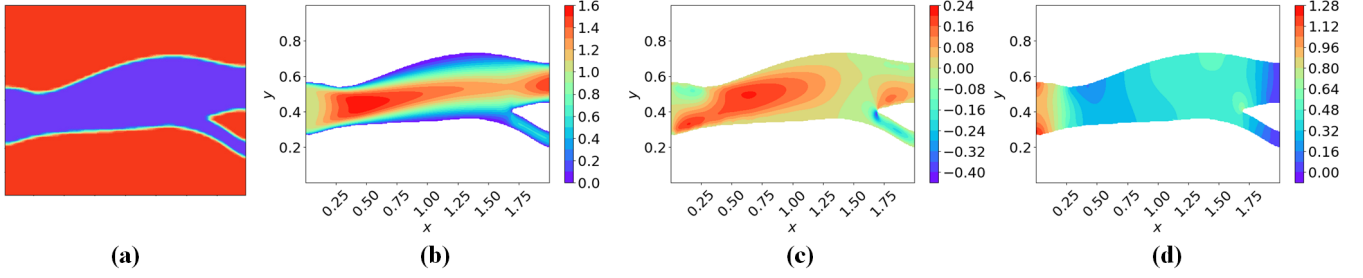


Figure 2: CFD results for a representative topology. (a) Channel topology (i.e., the final denoising state  $x_0$  generated by the diffusion model). (b)  $u$ -velocity contour obtained by CFD. (c)  $v$ -velocity contour obtained by CFD. (d)  $p$  contour obtained by CFD.

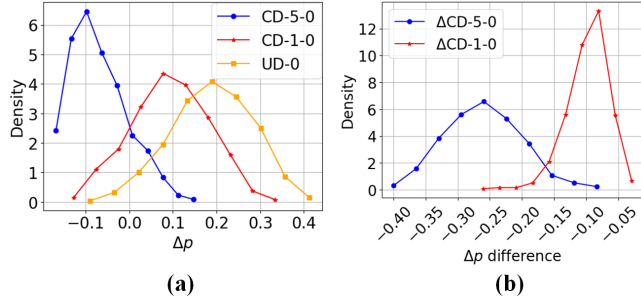


Figure 3: Histograms of  $\Delta p$  and the difference in  $\Delta p$  between samples generated with 50-step guidance and without guidance for 1000 test samples in fluidic channel topology design. During the denoising process,  $\Delta p$  is normalized by  $\ln \Delta p/5$ . (a) Histogram of  $\Delta p$ . “CD-5-0” represents  $\Delta p$  for samples generated with guidance and  $\alpha = 5$ . “CD-1-0” represents  $\Delta p$  for samples generated with guidance and  $\alpha = 1$ . “UD-0” represents  $\Delta p$  for samples generated without guidance. (b) Histogram of the difference in  $\Delta p$  across paired samples generated with and without guidance. “ $\Delta CD-5-0$ ” represents the difference in  $\Delta p$  between “CD-5-0” and “UD-0”. “ $\Delta CD-1-0$ ” represents the difference in  $\Delta p$  between “CD-1-0” and “UD-0”.

erated samples in Figure 6. The results show that  $\Delta p$  for the 1000 samples generated with guidance applied during denoising for just 10 steps remain significantly lower than the baseline, further demonstrating the method’s effectiveness.

## 4.2 Meta-Surface Design

**Problem Statement.** In this section, we demonstrate the effectiveness of the proposed guidance for design of frequency-selective meta-surfaces. Meta-surface design plays a crucial role in modern electronic devices, including 5G telecommunications. Electrical engineers typically rely on high-fidelity electromagnetic simulation tools (e.g., Ansys HFSS) in order to obtain and analyze electromagnetic behavior of these complex surfaces [Yang *et al.*, 2022]. These simulators solve Maxwell’s equations to provide the transmission and reflectance responses of various meta-surface designs. As an example, a meta-surface can be designed to match a given profile of the real ( $T_r$ ) and imaginary ( $T_i$ ) components of the transmission. A sample band-pass-type profile, where the magnitude of  $T_r$  and  $T_i$  must resemble the parabolic profile

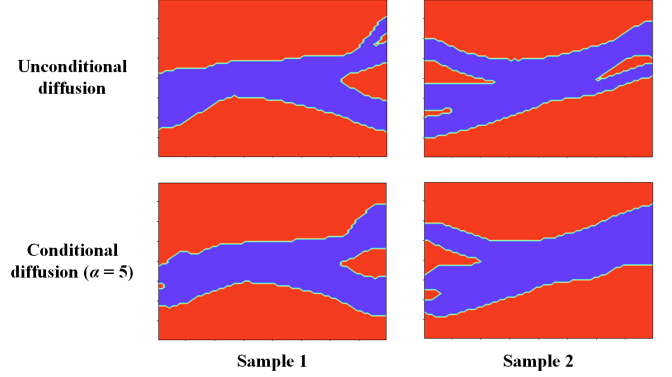


Figure 4: Examples of fluidic channel topologies generated with and without guidance.

$y = 1 - 2 \times (x - 0.5)^2$  across a specific frequency range, is one of many possible common design specifications engineers encounter. In this instance, the mean absolute error (MAE) between the predicted and target magnitudes of  $T_r$  and  $T_i$  is a key design goal, and minimizing MAE can be a key metric for guiding the design process. We present two representative meta-surface designs along with their corresponding  $T_r$ ,  $T_i$ , and magnitude of  $T_r$  and  $T_i$  curves in Figure 7.

**Implementation Details.** Following the previous example, we assume the availability of a black-box solver capable of evaluating the transmission profile for a given meta-surface design, as well as a pre-trained diffusion model that can generate random meta-surface designs. Examples of such black-box solvers include typical computational electromagnetics numerical solvers such as Ansys HFSS. During the denoising process, evaluations of the objective (i.e., MAE relative to the desired target profile) corresponding to the final design  $x_0$  can be obtained by calling this black-box solver. As a proof-of-concept, a regressor is trained to provide the black-box fitness evaluation using a paired dataset comprising meta-surface designs and their corresponding transmission profiles. Using the pre-trained diffusion model, guidance is applied for the second half of the denoising process with 30 samples evaluated per denoising step for gradient estimation. The entire denoising process consists of 100 steps, and the input design representation has a spatial resolution of  $64 \times 64$ .






Figure 5: Curves illustrating the change in  $\Delta p$  for the final 50 denoising steps of generation with and without guidance. “CD-5”, “CD-1” and “UD” indicate  $\Delta p$  is plotted for the scenarios where guidance with a scaling factor of  $\alpha = 5$ ,  $\alpha = 1$ , and  $\alpha = 0$  respectively, are applied.




Figure 6: Histograms of  $\Delta p$  and the difference in  $\Delta p$  between samples generated with 10-step guidance and without guidance for 1000 test samples in fluidic channel topology design. During the denoising process,  $\Delta p$  is normalized by  $\ln \Delta p/5$ . (a) Histogram of  $\Delta p$ . “CD-50-0” represents  $\Delta p$  for samples generated with  $\alpha = 50$ . “UD-0” represents  $\Delta p$  for samples generated without guidance. (b) Histogram of the difference in  $\Delta p$  across the paired samples generated with and without guidance. “ $\Delta$ CD-50-0” represents the difference in  $\Delta p$  between “CD-50-0” and “UD-0”.

**Results.** We generate 1000 samples to evaluate the effectiveness of the proposed evolvable conditional diffusion method. The histogram of the MAE and the difference in MAE between samples generated with and without guidance is presented in Figure 8. Figure 8 (a) clearly shows that the samples generated with guidance during denoising better match the target profile (i.e., show lower MAE values) than samples generated without guidance. The median MAE of samples generated with guidance is  $\approx 0.1$ , indicating a good match to the target profile. Figure 8 (b) further shows that the MAE of meta-surface designs generated with guidance is consistently lower than the baseline samples across all 1000 generated samples, emphasizing the effectiveness of the proposed gradient-free guidance.

Two samples from the 1000 test samples are selected for detailed analysis. Figure 9 illustrates that the MAE between the transmission profile for samples generated with guidance during denoising and the target profile steadily decreases to a small value, whereas the MAE remains large in the baseline case. Similarly, in Figure 10, the transmission profile of a generated meta-surface design aligns closely with the target

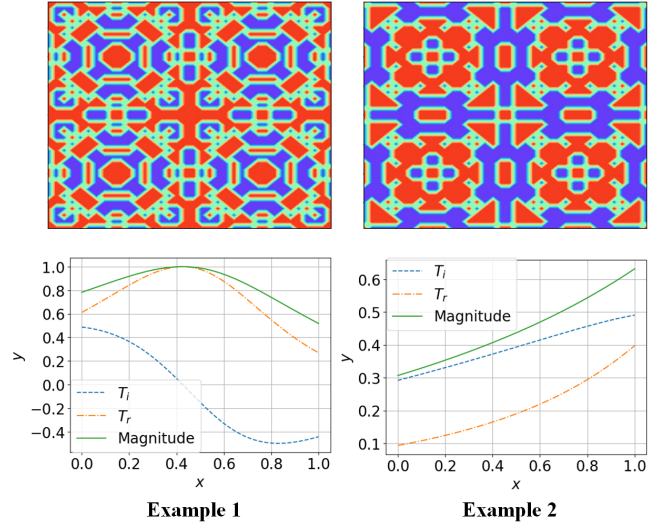


Figure 7: Two representative meta-surface designs along with curves depicting the real and imaginary components of their transmission profiles (i.e.,  $T_r$  and  $T_i$  respectively), and the corresponding transmission magnitudes.

profile. In contrast, the transmission profile of the baseline design generated without guidance exhibits significant discrepancy from the target. These results highlight the effectiveness of the proposed evolvable conditional diffusion.

## 5 Discussion

Guidance-based generation with diffusion models offers the advantage of seamless integration of desired performance or criteria into the diffusion denoising process using pre-trained diffusion models. However, a significant drawback is the reliance on differentiable proxies, a rule that precludes the use of many real-world black-box, non-differentiable simulators. This limitation is particularly critical in scientific domains, where accurate evaluation is best achieved by the use of established, physics-based simulators. While a common paradigm is to train a surrogate model of the guiding objective to provide a differentiable approximation of the objective, this often requires thousands of labeled samples which can be prohibitively expensive, particularly for high-dimensional design spaces (e.g., pixel-level geometry parameterizations or meta-surface unit cell designs). While it is possible to utilize finite difference, such approximations are known to scale poorly for high-dimensional inputs common in physical systems. Alternative scalable gradient-free approaches such as the simultaneous perturbation stochastic approximation (SPSA) proposed in [Shen *et al.*, 2025] are also known to suffer from high variance (noise) in high-dimensional spaces.

To address this challenge, we propose an evolvable conditional diffusion method which is fundamentally different from the current paradigm. By integrating evolution strategies directly into the denoising process, we eliminate the need for any *a priori* surrogate model. Further, this approach is more robust to noise as it derives stochastic gradient estimates from a population of sample evaluations.




Figure 8: Histograms of MAE and the difference in MAE between samples generated with 50-step guidance and without guidance for 1000 test samples in meta-surface design. (a) Histogram of MAE. “CD-5-0” represents MAE of meta-surface designs generated using  $\alpha = 5$ . “UD-0” represents MAE of meta-surface designs generated without any guidance during the denoising process. (b) Histogram of the difference in MAE across the paired samples generated with and without guidance. “ $\Delta$ CD-5-0” represents the difference in MAE between “CD-5-0” and “UD-0”.




Figure 9: Curves illustrating the change in MAE for the final 50 denoising steps of generation with and without guidance. “CD-5” and “UD” indicate MAE is plotted for the scenarios where guidance with a scaling factor of  $\alpha = 5$  and  $\alpha = 0$  respectively are applied.

More critically, the scalability of our evolvable conditional diffusion algorithm depends on the “intrinsic” dimensionality of the problem, which is often low even for high-resolution images where only a small number of pixel values affect the target objectives. Empirically, this conforms with our proposed method’s strong performance in high-dimensional settings, as the evolution strategies appear to effectively exploit key informative dimensions of the optimization problem, allowing their performance to scale with the intrinsic complexity or effective dimensionality of the task, rather than with the full dimensionality of the input search space [Salimans *et al.*, 2017]. While the scalability of such evolution strategies was prominently demonstrated by OpenAI in their seminal work [Salimans *et al.*, 2017], our method further extends this observation by explicitly incorporating the covariance matrix into the diffusion denoising process. Note that the ability to scale efficiently to high dimensions is yet another explanation of how our method fundamentally departs from recent surrogate-guided diffusion methods [Tan *et al.*, 2025], where the amount of data needed to train a differentiable surrogate model becomes intractable with increasing problem dimensionality.




Figure 10: Predicted transmission profiles of designs generated with and without guidance in the denoising process. “Target” indicates target profile which was used as guidance during denoising. “CD-5” indicates the predicted profile is obtained through denoising with a guidance of  $\alpha = 5$ . “UD” indicates the predicted profile is obtained with no guidance during denoising.

Lastly, while [Shen *et al.*, 2025] derive their gradient-free guidance algorithm by adopting and building upon the update rule in [Mazé and Ahmed, 2023], Proposition 1 shows that this update rule actually emerges from first principles as a mathematical consequence of the probabilistic evolution method, without requiring any prior assumptions. This positions our method as a more principled and scalable alternative for high-dimensional inverse design and optimization tasks, including but not limited to the two examples of fluidic channel topology and meta-surface design as demonstrated in this work.

## 6 Conclusion

In this work, we propose an evolvable conditional diffusion method by harnessing the generative capacity of probabilistic evolutionary computation. This approach evolves the diffusion process towards denoising distributions with higher fitness, thereby improving the conformance of generated samples to specified, desirable criteria. Notably, the derived update algorithm is provably analogous to the update in standard gradient-based guided diffusion models, but without the need for computing derivatives of the guidance signal.

We validate the effectiveness of the proposed evolvable conditional diffusion method through two design applications common in engineering: fluidic channel topology design and frequency-selective meta-surface design. Results demonstrate that the proposed method can effectively generate high-quality designs better aligned with desired performance criteria. Furthermore, the proposed method is highly adaptable and can be easily extended to other domains that require optimization through non-differentiable evaluations, as evidenced by our demonstrations in the two distinct fields of fluid dynamics and electromagnetic. This significantly broadens the applicability of diffusion models for design across a multitude of scientific and engineering domains.

## Acknowledgments

This research is partially supported by the National Research Foundation, Singapore, and Civil Aviation Authority under



the Aviation Transformation Programme - “ATM-Met Integration of Convective Weather Forecast and Impact Forecast Solutions Supporting Singapore Air Traffic Operations” (Award No. ATP2.0-ATM-MET.I2R) and A\*STAR under BMRC CSF (Award No. C240314053). This research is partially supported by the National Research Foundation, Singapore and DSO National Laboratories under the AI Singapore Programme - “Design Beyond What You Know: Material-Informed Differential Generative AI (MIDGAI) for Light-Weight High-Entropy Alloys and Multi-functional Composites (Stage 1a)” (AISG Award No. AISG2-GC-2023-010). This research is also partially supported by the Ramanujan Fellowship from the Science and Engineering Research Board, Government of India (Grant No. RJF/2022/000115).

## References

- [Avdeyev *et al.*, 2023] Pavel Avdeyev, Chenlai Shi, Yuhao Tan, Kseniia Dudnyk, and Jian Zhou. Dirichlet diffusion score model for biological sequence generation. In *Proceedings of the International Conference on Machine Learning*, pages 1276–1301. PMLR, 2023.
- [Bali *et al.*, 2020] Kavitesh Kumar Bali, Abhishek Gupta, Yew-Soon Ong, and Puay Siew Tan. Cognizant multitasking in multiobjective multifactorial evolution: MO-MFEA-II. *IEEE Transactions on Cybernetics*, 51(4):1784–1796, 2020.
- [Bansal *et al.*, 2023] Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 843–852, 2023.
- [Chen *et al.*, 2024] Minshuo Chen, Song Mei, Jianqing Fan, and Mengdi Wang. An overview of diffusion models: Applications, guided generation, statistical rates and optimization. *arXiv preprint arXiv:2404.07771*, 2024.
- [Croitoru *et al.*, 2023] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):10850–10869, 2023.
- [Dhariwal and Nichol, 2021] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.
- [Gainza *et al.*, 2020] Pablo Gainza, Freyr Sverrisson, Federico Monti, Emanuele Rodola, Davide Boscaini, Michael M Bronstein, and Bruno E Correia. Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. *Nature Methods*, 17(2):184–192, 2020.
- [Ghiringhelli *et al.*, 2015] Luca M Ghiringhelli, Jan Vybiral, Sergey V Levchenko, Claudia Draxl, and Matthias Scheffler. Big data of materials science: Critical role of the descriptor. *Physical Review Letters*, 114(10):105503, 2015.
- [Gupta *et al.*, 2022] Abhishek Gupta, Lei Zhou, Yew-Soon Ong, Zefeng Chen, and Yaqing Hou. Half a dozen real-world applications of evolutionary multitasking, and more. *IEEE Computational Intelligence Magazine*, 17(2):49–66, 2022.
- [Ho *et al.*, 2020] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [Igashov *et al.*, 2024] Ilia Igashov, Hannes Stärk, Clément Vignac, Arne Schneuing, Victor Garcia Satorras, Pascal Frossard, Max Welling, Michael Bronstein, and Bruno Correia. Equivariant 3d-conditional diffusion model for molecular linker design. *Nature Machine Intelligence*, 6:417–427, 2024.
- [Kim *et al.*, 2022] Heeseung Kim, Sungwon Kim, and Sungroh Yoon. Guided-TTS: A diffusion model for text-to-speech via classifier guidance. In *International Conference on Machine Learning*, pages 11119–11133. PMLR, 2022.
- [Kong *et al.*, 2021] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*, 2021.
- [Le *et al.*, 2013] Minh Nghia Le, Yew Soon Ong, Stefan Menzel, Yaochu Jin, and Bernhard Sendhoff. Evolution by adapting surrogates. *Evolutionary Computation*, 21(2):313–340, 2013.
- [Lehman *et al.*, 2020] Joel Lehman, Jeff Clune, Dusan Misevic, Christoph Adami, Lee Altenberg, Julie Beaulieu, Peter J Bentley, Samuel Bernard, Guillaume Beslon, David M Bryson, et al. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artificial Life*, 26(2):274–306, 2020.
- [Li *et al.*, 2022] Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-LM improves controllable text generation. *Advances in Neural Information Processing Systems*, 35:4328–4343, 2022.
- [Li *et al.*, 2024] Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, et al. Derivative-free guidance in continuous and discrete diffusion models with soft value-based decoding. *arXiv preprint arXiv:2408.08252*, 2024.
- [Liu and Thuerey, 2024] Qiang Liu and Nils Thuerey. Uncertainty-aware surrogate models for airfoil flow simulations with denoising diffusion probabilistic models. *AIAA Journal*, 62(8):2912–2933, 2024.
- [Lyu *et al.*, 2024] Yueming Lyu, Kim Yong Tan, Yew Soon Ong, and Ivor Tsang. Covariance-adaptive sequential black-box optimization for diffusion targeted generation. *arXiv preprint arXiv:2406.00812*, 2024.
- [Mazé and Ahmed, 2023] François Mazé and Faez Ahmed. Diffusion models beat GANs on topology optimization. In

- Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 9108–9116, 2023.
- [Miikkulainen and Forrest, 2021] Risto Miikkulainen and Stephanie Forrest. A biological perspective on evolutionary computation. *Nature Machine Intelligence*, 3(1):9–15, 2021.
- [Ollivier et al., 2017] Yann Ollivier, Ludovic Arnold, Anne Auger, and Nikolaus Hansen. Information-geometric optimization algorithms: A unifying picture via invariance principles. *Journal of Machine Learning Research*, 18(18):1–65, 2017.
- [Rombach et al., 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.
- [Salimans et al., 2017] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.
- [Shen et al., 2025] Yuchen Shen, Chenhao Zhang, Sijie Fu, Chenghui Zhou, Newell Washburn, and Barnabas Poczos. Chemistry-inspired diffusion with non-differentiable guidance. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [Song and Ermon, 2019] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.
- [Song and Ermon, 2020] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in Neural Information Processing Systems*, 33:12438–12448, 2020.
- [Song et al., 2020] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [Song et al., 2021] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34:1415–1428, 2021.
- [Sung et al., 2023] Nicholas Sung, Jian Cheng Wong, Chin Chun Ooi, Abhishek Gupta, Pao-Hsiung Chiu, and Yew-Soon Ong. Neuroevolution of physics-informed neural nets: Benchmark problems and comparative results. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*, pages 2144–2151, 2023.
- [Szymanski et al., 2023] Nathan J Szymanski, Bernardus Rendy, Yuxing Fei, Rishi E Kumar, Tanjin He, David Milsted, Matthew J McDermott, Max Gallant, Ekin Dogus Cubuk, Amil Merchant, et al. An autonomous laboratory for the accelerated synthesis of novel materials. *Nature*, 624(7990):86–91, 2023.
- [Tan et al., 2025] Kim Yong Tan, Yueming Lyu, Ivor Tsang, and Yew-Soon Ong. Fast direct: Query-efficient online black-box guidance for diffusion-model target generation. In *International Conference on Learning Representations*, 2025.
- [Valiant, 2009] Leslie G Valiant. Evolvability. *Journal of the ACM*, 56(1):1–21, 2009.
- [Wallace et al., 2023] Bram Wallace, Akash Gokul, Stefano Ermon, and Nikhil Naik. End-to-end diffusion latent optimization improves classifier guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7280–7290, 2023.
- [Wei et al., 2023] Zhao Wei, Jian Cheng Wong, Nicholas Sung, Abhishek Gupta, Chin Chun Ooi, Pao-Hsiung Chiu, My Ha Dao, and Yew-Soon Ong. How to select physics-informed neural networks in the absence of ground truth: A pareto front-based strategy. In *1st Workshop on the Synergy of Scientific and Machine Learning Modeling@ICML*, 2023.
- [Wierstra et al., 2014] Daan Wierstra, Tom Schaul, Tobias Glasmachers, Yi Sun, Jan Peters, and Jürgen Schmidhuber. Natural evolution strategies. *The Journal of Machine Learning Research*, 15(1):949–980, 2014.
- [Wong et al., 2024] Melvin Wong, Jiao Liu, Thiago Rios, Stefan Menzel, and Yew Soon Ong. LLM2FEA: Discover novel designs with generative evolutionary multitasking. *arXiv preprint arXiv:2406.14917*, 2024.
- [Xu et al., 2025] Qingshan Xu, Jiao Liu, Melvin Wong, Caishun Chen, and Yew-Soon Ong. Looks great, functions better: Physics compliance text-to-3D shape generation. In *International Joint Conference on Neural Networks*, 2025.
- [Yang et al., 2022] Zhong Liang Ou Yang, Yang Jiang, Jian Cheng Wong, Pao-Hsiung Chiu, Weijiang Zhao, My Ha Dao, and Chin Chun Ooi. Physics compliance as a metric for neural network uncertainty. In *IEEE Symposium Series on Computational Intelligence*, pages 1–7, 2022.
- [Yang et al., 2023] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.
- [Yuan et al., 2023] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. PhysDiff: Physics-guided human motion diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16010–16021, 2023.
- [Zeni et al., 2025] Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbé, Shoko Ueda, et al. A generative model for inorganic materials design. *Nature*, pages 1–3, 2025.