

DIFFUSION MODELS ARE EVOLUTIONARY ALGORITHMS

Anonymous authors

Paper under double-blind review

ABSTRACT

In a convergence of machine learning and biology, we reveal that diffusion models are evolutionary algorithms. By considering evolution as a denoising process and reversed evolution as diffusion, we mathematically demonstrate that diffusion models inherently perform evolutionary algorithms, naturally encompassing selection, mutation, and reproductive isolation. Building on this equivalence, we propose the Diffusion Evolution method: an evolutionary algorithm utilizing iterative denoising – as originally introduced in the context of diffusion models – to heuristically refine solutions in parameter spaces. Unlike traditional approaches, Diffusion Evolution efficiently identifies multiple optimal solutions and outperforms prominent mainstream evolutionary algorithms. Furthermore, leveraging advanced concepts from diffusion models, namely latent space diffusion and accelerated sampling, we introduce Latent Space Diffusion Evolution, which finds solutions for evolutionary tasks in high-dimensional complex parameter space while significantly reducing computational steps. This parallel between diffusion and evolution not only bridges two different fields but also opens new avenues for mutual enhancement, raising questions about open-ended evolution and potentially utilizing non-Gaussian or discrete diffusion models in the context of Diffusion Evolution.

1 INTRODUCTION

At least two processes in the biosphere have been recognized as capable of generalizing and driving novelty: evolution, a slow variational process adapting organisms across generations to their environment through natural selection (Darwin, 1959; Dawkins, 2016); and learning, a faster transformational process allowing individuals to acquire knowledge and generalize from subjective experience during their lifetime (Kandel, 2013; Courville et al., 2006; Holland, 2000; Dayan & Abbott, 2001). These processes are intensively studied in distinct domains within artificial intelligence. Relatively recent work has started drawing parallels between the seemingly unrelated processes of evolution and learning (Watson & Levin, 2023; Vanchurin et al., 2022; Levin, 2022; Watson et al., 2022; Kouvaris et al., 2017; Watson & Szathmary, 2016; Watson et al., 2016; Power et al., 2015; Hinton et al., 1987; Baldwin, 2018). We here argue that in particular diffusion models (Sohl-Dickstein et al., 2015; Song et al., 2020b; Ho et al., 2020; Song et al., 2020a), where generative models trained to sample data points through incremental stochastic denoising, can be understood through evolutionary processes, inherently performing natural selection, mutation, and reproductive isolation.

The evolutionary process is fundamental to biology, enabling species to adapt to changing environments through mechanisms like natural selection, genetic mutations, and hybridizations (Rosen, 1991; Wagner, 2015; Dawkins, 1996); this adaptive process introduces variations in organisms’ genetic codes over time, leading to well-adapted and diverse individuals (Mitchell & Cheney, 2024; Levin, 2023; Gould, 2002; Dennett, 1995; Smith & Szathmary, 1997; Szathmary, 2015). Evolutionary algorithms utilize such biologically inspired variational principles to iteratively refine sets of numerical parameters that encode potential solutions to often rugged objective functions (Vikhar, 2016; Golberg, 1989; Grefenstette, 1993; Holland, 1992). In another side, recent breakthroughs in deep learning have led to the development of diffusion models—generative algorithms that iteratively refine data points to sample novel yet realistic data following complex target distributions: models like Stable Diffusion (Rombach et al., 2022) and Sora (Brooks et al., 2024) demonstrate remarkable realism and diversity in generating image and video. Notably, both evolutionary processes and




Figure 1: Evolution processes can be viewed as the inverse process of diffusion, where higher fitness populations (red points) have higher final probability density. The initially unstructured parameters are iteratively refined towards high-fitness regions in parameter space.

diffusion models rely on iterative refinements that combine directed updates with undirected perturbations: in evolution, random genetic mutations introduce diversity while natural selection guides populations toward greater fitness, and in diffusion models, random noise is progressively transformed into meaningful data through learned denoising steps that steer samples toward the target distribution. This parallel raises fundamental questions: Are the mechanisms underlying evolution and diffusion models fundamentally connected? Is this similarity merely an analogy, or does it reflect a deeper mathematical duality between biological evolution and generative modeling?

To answer these questions, we first examine evolution from the perspective of generative models. By considering populations of species in the biosphere, the variational evolution process can also be viewed as a transformation of distributions: the distributions of genotypes and phenotypes. Over evolutionary time scales, mutation and selection collectively alter the shape of these distributions. Similarly, many biologically inspired evolutionary algorithms can be understood in the same way: they optimize an objective function by maintaining- and iteratively changing a large population’s distribution. In fact, this concept is central to most generative models: the transformation of distributions. Variational Autoencoders (VAEs) (Kingma, 2013), Generative Adversarial Networks (GANs) (Goodfellow et al., 2014), and diffusion models are all trained to transform simple distributions, typically standard Gaussian distributions, into complex distributions, where the samples represent meaningful images, videos, or audio, etc.

On the other hand, diffusion models can also be viewed from an evolutionary perspective. As a generative model, diffusion models transform Gaussian distributions in an iterative manner into complex, structured data-points that resemble the training data distribution. During the training phase, the data points are corrupted by adding noise, and the model is trained to predict this added noise to reverse the process. In the sampling phase, starting with Gaussian-distributed data points, the model iteratively denoises to incrementally refine the data point samples. By considering noise-free samples as the desired outcome, such a directed denoising can be interpreted as directed selection, with each step introducing slight noise, akin to mutations. Together, this resembles an evolutionary process (Fields & Levin, 2020), where evolution is formulated as a combination of deterministic dynamics and stochastic mutations within the framework of non-equilibrium thermodynamics (Ao, 2005). This aligns with recent ideas that interpret the genome as a latent space parameterization of a multi-scale generative morphogenetic process, rather than a direct blueprint of an organism (Mitchell & Cheney, 2024; Hartl et al., 2024; Levin, 2023; Gould, 2002). If one were to revert the time direction of an evolutionary process, the evolved population of potentially highly correlated high-fitness solutions will dissolve gradually, i.e., step by step and thus akin to the forward process in diffusion models, into the respectively chosen initial distribution, typically Gaussian noise, see Figure 1.

Driven by this intuition, we conduct a thorough investigation into the connections between diffusion models and evolutionary algorithms, discovering that these seemingly disparate concepts share the same mathematical foundations. This insight leads to a novel approach, the Diffusion Evolution algorithm, which directly utilizes the framework of diffusion models to perform evolutionary optimization. This can be obtained by inverting the diffusion process with the Bayesian method. Our

analytical study of Diffusion Evolution reveals promising parallels to biological evolution, naturally incorporating concepts such as mutation, hybridization, and even reproductive isolation.

This equivalence provides a new way of improving evolutionary algorithms and has the potential to unify developments in both fields. By mimicking biological evolution, evolutionary algorithms have shown promising results in numerical optimization, particularly for tasks that cannot be effectively trained using gradient-based methods (Wang et al., 2024; Goodfellow et al., 2014). These algorithms thus excel in exploring complex, rugged search spaces and finding globally optimal or near-optimal solutions (Ho & Salimans, 2022; Hansen, 2016; Hansen & Ostermeier, 2001; Sehnke et al., 2010). While the biosphere exhibits extreme diversity in lifeforms of life, many evolutionary strategies, such as CMA-ES (Hansen & Ostermeier, 2001), and PEPG (Sehnke et al., 2010), struggle to find diverse solutions (Lehman & Stanley, 2011). However, our Diffusion Evolution Algorithm offers a new approach. By naturally incorporating mutation, hybridization, and reproductive isolation, our algorithm can discover diverse solutions, mirroring the diversity of the biosphere, rather than converging on a single solution as is often the case with traditional methods. Since this parallel between diffusion and evolution exists naturally and not imposed by our design, the two fields – diffusion models and evolutionary computing – can mutually benefit from each other. For example, we demonstrate that the concept of latent diffusion (Rombach et al., 2022) and accelerated sampling (Nichol & Dhariwal, 2021) can significantly improve the performance of our Diffusion Evolution algorithm.

In the following sections, we will first review evolutionary strategies and diffusion models, introduce the mathematical connection between diffusion and evolution, and propose the Diffusion Evolution algorithm. Then, we will quantitatively compare our algorithm to conventional evolutionary strategies, demonstrating its capability to find multiple solutions, solve complex evolutionary tasks, and incorporate developments from diffusion model literature. Finally, the emerging connections between the derived algorithm and evolution will be discussed, along with the potentials of this finding and the limitations of our algorithm.

2 BACKGROUND

2.1 EVOLUTIONARY ALGORITHMS

The principles of evolution extend far beyond biology, offering exceptional utility in addressing complex systems across various domains. The key components of this process – imperfect replication with heredity and fitness-based selection – are sufficiently general to find applications in diverse fields. In computer and data science, for instance, evolutionary algorithms play a crucial role in optimization (Vikhar, 2016; Grefenstette, 1993; Golberg, 1989; Holland, 1992). These heuristic numerical techniques, such as CMA-ES (Hansen & Ostermeier, 2001) and PEPG (Sehnke et al., 2010), maintain and optimize a population of genotypic parameters over successive generations through operations inspired by biological evolution, such as selection of the fittest, reproduction, genetic crossover, and mutations. The goal is to gradually adapt the parameters of the entire population so individual genotypic samples, or short individuals, perform well when evaluated against an objective- or fitness function. These algorithms harness the dynamics of evolutionary biology to discover optimal or near-optimal solutions within vast, complex, and otherwise intractable parameter spaces. The evaluated numerical fitness score of an individual correlates with its probability of survival and reproduction, ensuring that individuals with superior traits have a greater chance of passing their genetic information to the next generation, thus driving the evolutionary process toward more optimal solutions. Such approaches are particularly valuable when heuristic solutions are needed to explore extensive combinatorial and permutation landscapes.

Some evolutionary algorithms operate with discrete, others with continuous sets of parameters. Here, we focus on the latter since discrete tasks can be seen as a subcategory of continuous tasks. Typically, the structure of the parameter space is a priori unknown. Thus, the initial population is often sampled from a standard normal distribution. As explained above, this initially random population is successively adapted and refined, generation by generation, to perform well on an arbitrary objective function. Thus, initially randomized parameters are successively varied by evolutionary algorithms into sets of potentially highly structured parameters that perform well on the specified task, eventually (and hopefully) solving the designated problem by optimizing the objective function. Thus, evolutionary algorithms can be understood as generative models that use

162 heuristic information about already explored regions of the parameter space (at least from the previous generation) to sample potentially better-adapted offspring individuals for the next generation (c.f., CMA-ES (Hansen et al., 2003), etc.).

166 2.2 DIFFUSION MODELS

168 Diffusion models, such as denoising diffusion probabilistic models (DDPM) (Ho et al., 2020) and denoising diffusion implicit models (DDIM) (Song et al., 2020a), have shown promising generative capabilities in image, video, and even neural network parameters (Wang et al., 2024). Similar to other generative approaches such as GANs, VAEs, and flow-based models (Dinh et al., 2016; Chen et al., 2019), diffusion models transform a simple distribution, often a Gaussian, into a more complex distribution that captures the characteristics of the training data. Diffusion models achieve this, in contrast to other techniques, via iterative denoising steps, progressively transforming noisy data into less noisy (Raya & Ambrogioni, 2024), more coherent representations (Sohl-Dickstein et al., 2015).

176 Diffusion models have two phases: diffusion and denoising. In the diffusion phase, we are blending original data points with some extent of Gaussian noise. Specifically, let \mathbf{x}_0 be the original data point and \mathbf{x}_T be the fully distorted data, then the process of diffusion can be represented as:

$$179 \mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} \boldsymbol{\epsilon}, \quad (1)$$

180 where the amount of total noise $\boldsymbol{\epsilon} \sim \mathcal{N}(0, I)$ added to the data \mathbf{x}_0 at time step $t \in [0, T]$ is controlled by α_t that is monotonously decreasing from $\alpha_0 = 1$ to $\alpha_T \sim 0$. Thus, while \mathbf{x}_0 represents the original data, \mathbf{x}_T will consist entirely of Gaussian noise. To restore such diffused data, a predictive model, typically a neural network $\boldsymbol{\epsilon}_\theta$ with parameter $\boldsymbol{\theta}$, is trained to predict the added total noise given \mathbf{x}_t and time step t . Thus, diffusion models can be trained by minimizing the loss function:

$$185 \mathcal{L} = \mathbb{E}_{\mathbf{x}_0 \sim p_{\text{data}}, \boldsymbol{\epsilon} \sim \mathcal{N}(0, I)} \|\boldsymbol{\epsilon}_\theta(\sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} \boldsymbol{\epsilon}, t) - \boldsymbol{\epsilon}\|^2, \quad (2)$$

186 where p_{data} is the distribution of training data. So, conventionally, diffusion models are understood as predicting the added noise during the diffusion process.

189 In the denoising phase, starting with a noisy pattern, the trained models are used to iteratively remove the predicted noise from current data: from $\mathbf{x}_T \sim \mathcal{N}(0, I)$, iteratively refine to $\mathbf{x}_{T-1}, \mathbf{x}_{T-2}, \dots$, until \mathbf{x}_0 . In the DDIM framework, this sampling process is given by:

$$192 \mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t-1} - \sigma_t^2} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) + \sigma_t \mathbf{w}, \quad (3)$$

195 where σ_t controls the amount of noise $\mathbf{w} \sim \mathcal{N}(0, I)$ added during the denoising phase. Notably, the schedule of α_t and σ_t will both affect the denoising process and can be chosen based on our needs under the DDIM framework.

199 3 DIFFUSION MODELS ARE EVOLUTIONARY ALGORITHMS

201 Similar to the relationship between energy and probability in statistical physics, evolutionary tasks can be connected to generative tasks by mapping fitness to probability density: higher fitness correspond to higher probability density. Thus, given a fitness function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, we can choose a mapping g to transform f into a probability density function $p(\mathbf{x}) = g[f(\mathbf{x})]$. When aligning the denoising process in a diffusion model with evolution, we want \mathbf{x}_0 to follow this density function, i.e., $p(\mathbf{x}_0 = \mathbf{x}) = g[f(\mathbf{x})]$. This requires an alternative view of diffusion models (Song et al., 2020a): diffusion models are directly predicting the original data samples from noisy versions of those samples at each time step. Given the diffusion process $\mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} \boldsymbol{\epsilon}$, we can easily express \mathbf{x}_0 in terms of the noise $\boldsymbol{\epsilon}$, and vice versa:

$$210 \mathbf{x}_0 = \frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \boldsymbol{\epsilon}}{\sqrt{\alpha_t}}, \text{ and } \boldsymbol{\epsilon} = \frac{\mathbf{x}_t - \sqrt{\alpha_t} \mathbf{x}_0}{\sqrt{1 - \alpha_t}}. \quad (4)$$

212 In diffusion models, the error $\boldsymbol{\epsilon}$ between \mathbf{x}_0 and \mathbf{x}_t is estimated by a neural network, i.e., $\hat{\boldsymbol{\epsilon}} = \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)$. Thus, Equation 4 provides an estimation $\hat{\mathbf{x}}_0$ for \mathbf{x}_0 when replacing $\boldsymbol{\epsilon}$ with $\hat{\boldsymbol{\epsilon}}$. Hence, the sampling process of DDIM (Song et al., 2020a) in Equation 3 can be written as:

$$215 \mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}} \hat{\mathbf{x}}_0 + \sqrt{1 - \alpha_{t-1} - \sigma_t^2} \hat{\boldsymbol{\epsilon}} + \sigma_t \mathbf{w}. \quad (5)$$

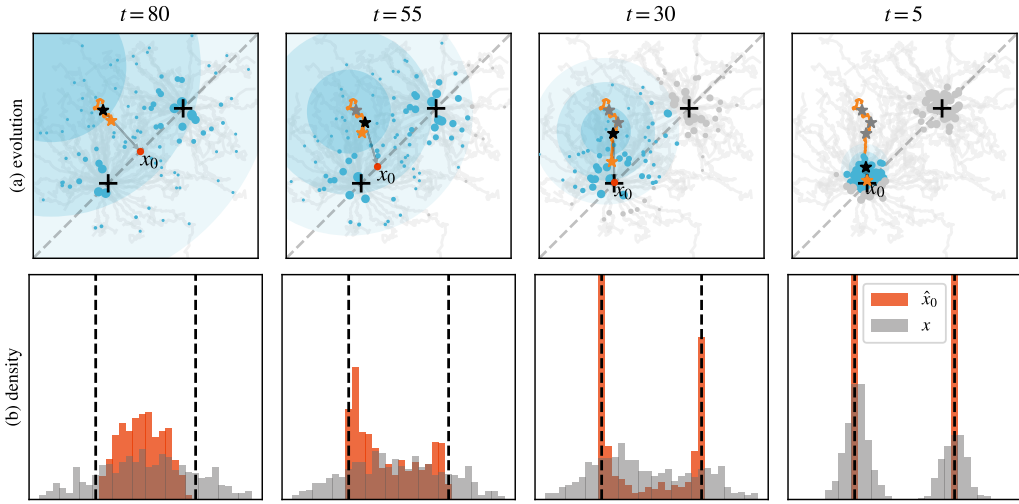


Figure 2: **(a)** Diffusion Evolution on a two-peak fitness landscape: Populations near the two black crosses have higher fitness. For each individual \mathbf{x}_t (black star), its target $\hat{\mathbf{x}}_0$ (red dots) is estimated by a weighted average of its neighbors (c.f., dots within the blue disks, respectively); larger dot-size indicates higher fitness. The individual then moves a small step forward to the next generation (orange star). As evolution proceeds, the neighbor range decreases, making the process increasingly sensitive to local neighbors, thereby enabling global competition originally, while “zooming in” eventually to balance between optimization and diversity. **(b)** By mapping the population to a 1-D space (dashed lines in (a)), we track the progress of Diffusion Evolution. As evolution progresses, both the individuals (gray) and their estimated origins (red) move closer to the targets (vertical dashed lines), with the estimated origins advancing faster.

Since the denoising step in diffusion models requires an estimation of \mathbf{x}_0 , we need to derive it from sample \mathbf{x}_t and the corresponding fitness $f(\mathbf{x}_t)$. The estimation of \mathbf{x}_0 can be expressed as a conditional probability $p(\mathbf{x}_0 = \mathbf{x} | \mathbf{x}_t)$. Using Bayes’ theorem and $p(\mathbf{x}_0 = \mathbf{x}) = g[f(\mathbf{x})]$ yields:

$$p(\mathbf{x}_0 = \mathbf{x} | \mathbf{x}_t) = \frac{p(\mathbf{x}_t | \mathbf{x}_0 = \mathbf{x})p(\mathbf{x}_0 = \mathbf{x})}{p(\mathbf{x}_t)} = \frac{p(\mathbf{x}_t | \mathbf{x})g[f(\mathbf{x})]}{p(\mathbf{x}_t)}. \quad (6)$$

Here, $p(\mathbf{x}_t | \mathbf{x}_0 = \mathbf{x})$ can be computed easily by $\mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t}\mathbf{x}, 1 - \alpha_t)$ given the design of the diffusion process, i.e., $\mathbf{x}_t = \sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{1 - \alpha_t}\epsilon$. Since deep-learning-based diffusion models are trained using mean squared error loss, the \mathbf{x}_0 estimated by \mathbf{x}_t should be the weighted average of the sample \mathbf{x} . Hence, the estimation function of \mathbf{x}_0 becomes:

$$\hat{\mathbf{x}}_0(\mathbf{x}_t, \alpha, t) = \sum_{\mathbf{x} \sim p_{\text{eval}}(\mathbf{x})} p(\mathbf{x}_0 = \mathbf{x} | \mathbf{x}_t) \mathbf{x} = \sum_{\mathbf{x} \sim p_{\text{eval}}(\mathbf{x})} g[f(\mathbf{x})] \frac{\mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t}\mathbf{x}, 1 - \alpha_t)}{p(\mathbf{x}_t)} \mathbf{x}, \quad (7)$$

where p_{eval} is the evaluation sample on which we compute the fitness score, here given by the current population $\mathbf{X}_t = (\mathbf{x}_t^{(1)}, \mathbf{x}_t^{(2)}, \dots, \mathbf{x}_t^{(N)})$ of N individuals. Equation 7 has three weight terms: The first term $g[f(\mathbf{x})]$ assigns larger weights to high fitness samples. For each individual sample \mathbf{x}_t , the second Gaussian term $\mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t}\mathbf{x}, 1 - \alpha_t)$ makes each individual only sensitive to local neighbors of evaluation samples. The third term $p(\mathbf{x}_t)$ is a normalization term. Hence, $\hat{\mathbf{x}}_0$ can be simplified to:

$$\hat{\mathbf{x}}_0(\mathbf{x}_t, \alpha, t) = \frac{1}{Z} \sum_{\mathbf{x} \in \mathbf{X}_t} g[f(\mathbf{x})] \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t}\mathbf{x}, 1 - \alpha_t) \mathbf{x}, \quad (8)$$

where Z is the normalization term:

$$Z = p(\mathbf{x}_t) = \sum_{\mathbf{x} \in \mathbf{X}_t} g[f(\mathbf{x})] \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t}\mathbf{x}, 1 - \alpha_t). \quad (9)$$

When substituting Equation 8 into Equation 4 we can express $\hat{\epsilon}$ as:

$$\hat{\epsilon}(\mathbf{x}_t, \boldsymbol{\alpha}, t) = \frac{\mathbf{x}_t - \sqrt{\alpha_t} \hat{\mathbf{x}}_0(\mathbf{x}_t, \boldsymbol{\alpha}, t)}{\sqrt{1 - \alpha_t}}, \quad (10)$$

and by substituting Equations 8 and 10 into Equation 5, we derive the Diffusion Evolution algorithm: an evolutionary optimization procedure based on iterative error correction akin to diffusion models but without relying on neural networks at all, see pseudocode in Algorithm 1. When inversely denoising, i.e., evolving from time T to 0, while increasing α_t , the Gaussian term will initially have a high variance, allowing global exploration at first. As the evolution progresses, the variance decreases giving lower weight to distant populations, leads to local optimization (exploitation). This locality avoids global competition and thus allows the algorithm to maintain multiple solutions and balance exploration and exploitation. Hence, the denoising process of diffusion models can be understood in an evolutionary manner: $\hat{\mathbf{x}}_0$ represents an estimated high fitness parameter target. In contrast, \mathbf{x}_t can be considered as diffused from high-fitness points. The first two parts in the Equation 5, i.e., $\sqrt{\alpha_{t-1}}\hat{\mathbf{x}}_0 + \sqrt{1 - \alpha_{t-1} - \sigma_t^2}\hat{\epsilon}$, guide the individuals towards high fitness targets in small steps. The last part of Equation 5, $\sigma_t\mathbf{w}$, is an integral part of diffusion models, perturbing the parameters in our approach similarly to random mutations.

Algorithm 1 Diffusion Evolution

Require: Population size N , parameter dimension D , fitness function f , density mapping function g , total evolution steps T , diffusion schedule $\boldsymbol{\alpha}$ and noise schedule $\boldsymbol{\sigma}$.

Ensure: $\alpha_0 \sim 1, \alpha_T \sim 0, \alpha_i > \alpha_{i+1}, 0 < \sigma_i < \sqrt{1 - \alpha_{i-1}}$

```

1:  $[\mathbf{x}_T^{(1)}, \mathbf{x}_T^{(2)}, \dots, \mathbf{x}_T^{(N)}] \leftarrow \mathcal{N}(0, I^{N \times D})$  ▷ Initialize population
2: for  $t \in [T, T - 1, \dots, 2]$  do
3:    $\forall i \in [1, N]: Q_i \leftarrow g[f(\mathbf{x}_t^{(i)})]$  ▷ Fitness are cached to avoid repeated evaluations
4:   for  $i \in [1, 2, \dots, N]$  do
5:      $Z \leftarrow \sum_{j=1}^N Q_j \mathcal{N}(\mathbf{x}_t^{(i)}; \sqrt{\alpha_t} \mathbf{x}_t^{(j)}, 1 - \alpha_t)$ 
6:      $\hat{\mathbf{x}}_0 \leftarrow \frac{1}{Z} \sum_{j=1}^N Q_j \mathcal{N}(\mathbf{x}_t^{(i)}; \sqrt{\alpha_t} \mathbf{x}_t^{(j)}, 1 - \alpha_t) \mathbf{x}_t^{(j)}$ 
7:      $\mathbf{w} \leftarrow \mathcal{N}(0, I^D)$ 
8:      $\mathbf{x}_{t-1}^{(i)} \leftarrow \sqrt{\alpha_{t-1}} \hat{\mathbf{x}}_0 + \sqrt{1 - \alpha_{t-1} - \sigma_t^2} \frac{\mathbf{x}_t^{(i)} - \sqrt{\alpha_t} \hat{\mathbf{x}}_0}{\sqrt{1 - \alpha_t}} + \sigma_t \mathbf{w}$ 
9:   end for
10: end for

```

Figure 2(a) demonstrates the detailed evolution process of a multi-target fitness landscape with two optimal points (see exact fitness function in Appendix A.1). Each individual estimates high fitness parameter targets and moves toward the target along with random mutations. The high fitness parameter targets $\hat{\mathbf{x}}_0$ are estimated based on their neighbors' fitness scores (neighbors are shown in blue disks, with radius proportional to $\sqrt{1 - \alpha_t} / \sqrt{\alpha_t}$). The estimated targets $\hat{\mathbf{x}}_0$ typically move faster than the individuals while the individuals are successively refined in small denoising steps in the direction of the estimated target, see Figure 2(b). Although $\hat{\mathbf{x}}_0$ often have higher fitness, they exhibit lower diversity, hence they are used as a goal of individuals instead of the final solutions. This difference also provides flexibility in balancing between more greedy and more diverse strategies.

4 EXPERIMENTS

We conduct two sets of experiments to study Diffusion Evolution in terms of diversity and solving complex reinforcement learning tasks. Moreover, we utilize techniques from the diffusion models literature to improve Diffusion Evolution. In the first experiment, we adopt an accelerated sampling method (Nichol & Dhariwal, 2021) to significantly reduce the number of iterations. In the second experiment, we propose Latent Space Diffusion Evolution, inspired by latent space diffusion models (Rombach et al., 2022), allowing us to deploy our approach to complex problems with high-dimensional parameter spaces through exploring a lower-dimensional latent space.

4.1 MULTI-TARGET EVOLUTION

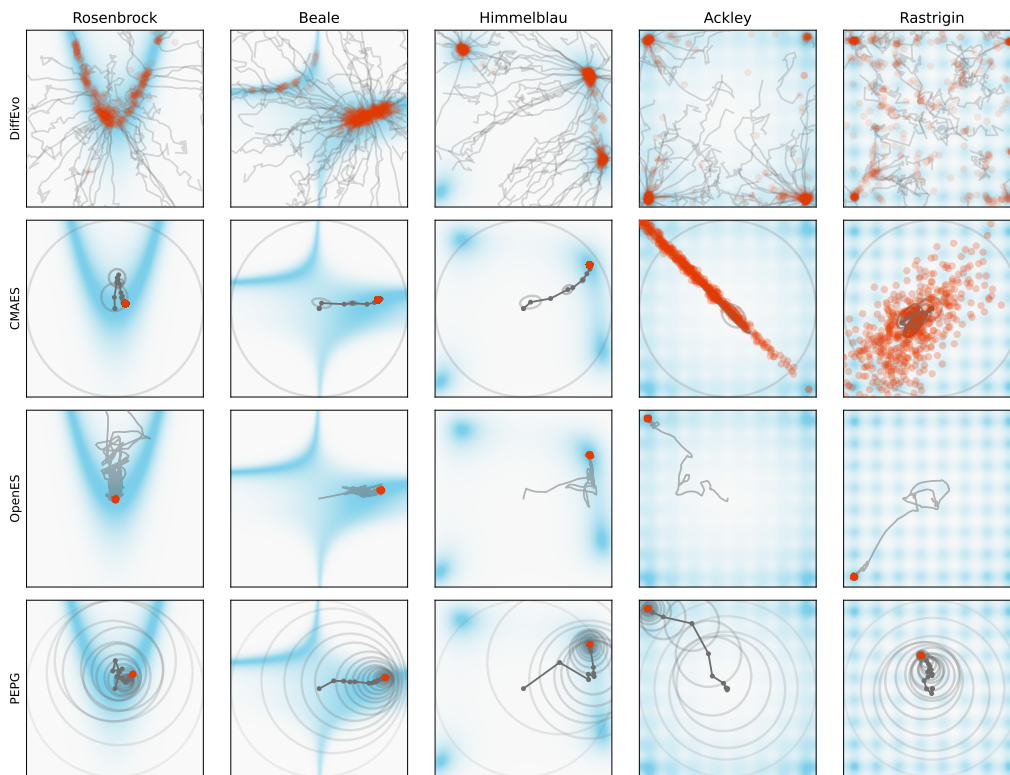


Figure 3: One of the benchmark experiments (columns) on different fitness functions with selected evolutionary algorithms (rows). The blue regions represent high fitness, while white indicates low fitness in a two-dimensional parameter space. All fitness functions are mapped to a 0 to 1 range to make them comparable, as detailed in the Appendix. The Diffusion Evolution algorithm finds multiple optimal solutions in the 2D benchmarks while maintaining genetic diversity. Red dots indicating the final population and gray lines show the trajectories of populations; for simplicity, only the trajectories of 64 individuals are plotted here. In the CMAES and PEPG experiments, the gray ellipsoids represent the estimated covariances at each step, and the gray lines represent the history of estimated averages. The red dots indicate the final population. In the OpenES experiments, the red dots indicate the final population, and the gray line represents the parameter trajectory.

To compare our method to selected mainstream evolutionary algorithms, we choose five different two-dimensional fitness landscapes as benchmarks: The Rosenbrock and Beale fitness functions have a single optimal point, while the Himmelblau, Ackley, and Rastrigin functions have multiple optimal solutions, see Appendix A.4 for more details; we compare our method to other evolutionary strategies, including CMA-ES (Hansen et al., 2003), OpenES (Salimans et al., 2017), and PEPG (Sehnke et al., 2010). The experiments show that our Diffusion Evolution algorithm can find diverse solutions on the Himmelblau, Ackley, and Rastrigin functions, while other methods struggle (see Figure 3 and Table 1). CMA-ES, OpenES, and PEPG either focus on finding a single solution or get distracted by multiple high-fitness peaks, leading to sub-optimal results. Our experiments demonstrate that the Diffusion Evolution algorithm can identify diverse solutions and adapt to various fitness landscapes.

The most time-consuming part of evolutionary algorithms is often the fitness evaluation. In this experiment, we adopt an accelerated sampling method from the diffusion models literature to reduce the number of iteration. As proposed by Nichol & Dhariwal (2021), instead of the default α_t scheduling in DDPM, a cosine scheduling $\alpha_t = \cos(\pi t/T)/2 + 1/2$ leads to better performance

Table 1: Average entropy of the top-64 elite populations after evolution, with average fitness (0 to 1) in brackets. Higher entropy indicates greater diversity, and higher fitness reflects better solutions. Bold numbers are the highest value per line, and underlined numbers are the second-highest.

Task	Diffusion Evolution	CMA-ES	OpenES	PEPG
Rosenbrock	4.93 (1.00)	0.00 (1.00)	<u>1.02 (1.00)</u>	0.86 (1.00)
Beale	4.21 (1.00)	0.00 (1.00)	<u>1.06 (1.00)</u>	0.32 (1.00)
Himmelblau	2.58 (1.00)	<u>0.07 (1.00)</u>	0.05 (1.00)	0.00 (1.00)
Ackley	<u>2.49 (1.00)</u>	3.96 (0.52)	0.00 (1.00)	0.02 (<u>0.96</u>)
Rastrigin	<u>3.29 (0.82)</u>	5.77 (0.18)	0.00 (1.00)	0.00 (<u>0.11</u>)

when T is small. With this, we can significantly reduce the number of fitness evaluations while maintaining sampling diversity and quality.

To systematically compare different methods, we repeated the evolution 100 times for each method. In all experiments, the fitness functions were rescaled from 0 to 1, with 1 representing the highest fitness (see Appendix A.4). Each experiment was conducted with a population of 512 and 25 iterations, except for the OpenES method, which requires 1000 steps to converge. To quantify diversity, we then calculated Shannon entropy of the final population by gridding the space and counting the individuals in different grid cells (we select the top-64 fitness individuals, focusing solely on elite individuals). The results in Table 1 show that our method consistently finds more diverse solutions without sacrificing fitness performance. While CMA-ES shows higher entropy on the Ackley and Rastrigin functions, it finds significantly lower fitness solutions compared to Diffusion Evolution, suggesting it is distracted by multiple solutions rather than finding diverse ones (see examples in Figure 3).

4.2 LATENT SPACE DIFFUSION EVOLUTION

Here, we apply the Diffusion Evolution method to reinforcement learning tasks (Sutton & Barto, 1998) to train neural networks for controlling the cart-pole system (Barto et al., 1983). This system has a cart with a hinged pole, and the objective is to keep the pole vertical as long as possible by moving the cart sideways while not exceeding a certain range, see Figure 4(d). The game is terminated if the pole-angle exceeds $\pm 12^\circ$ or the cart position exceeds ± 2.4 . Thus, longer duration yield higher fitness. We use a two-layer neural network of 58 parameters to control the cart, with inputs being the current position, velocity, pole angle, and pole angular velocity. The output of the neural network determines whether to move left or right. See more details about the neural network in Appendix A.5.1. The task is considered solved when a fitness score (cumulative reward) of 500 is reached consistently over several episodes.

Deploying our original Diffusion Evolution method to this problem results in poor performance and lack of diversity, see Figure 4(b-c). To address this issue, we propose *Latent Space Diffusion Evolution*: inspired by the latent space diffusion model (Rombach et al., 2022), we map individual parameters into a lower-dimensional latent space in which we perform the Diffusion Evolution Algorithm. This approach significantly improves performance and restores diversity. The key insight comes from the Gaussian term in Equation 8 for estimating the original points \hat{x}_0 : the distance between parameters increases with higher dimensions, making the evolution more local and slower. Moreover, the parameter or genotype space may have dimensions that don’t effectively impact fitness, known as sloppiness (Gutenkunst et al., 2007). Assigning random values to these dimensions often doesn’t affect fitness, similar to genetic drift or neutral genes, suggesting the true high-fitness genotype distribution is lower-dimensional. The straightforward approach is directly denoising in a lower-dimensional latent space z and estimating high-quality targets z_0 via:

$$\hat{z}_0(z_t, \alpha, t) = \sum_z p(z|z_t)z = \frac{1}{Z} \sum_{x \sim p_{\text{eval}}(x)} g[f'(z)]\mathcal{N}(z_t; \sqrt{\alpha_t}z, 1 - \alpha_t)z. \quad (11)$$

However, this approach requires a decoder and a new fitness function f' for z , which can be challenging to obtain. To circumvent this, we approximate the latent diffusion by using the latent space only to calculate the distance between individuals. While we don’t know the exact distribution of x a priori, a random projection can often preserve the distance relationships between populations, as

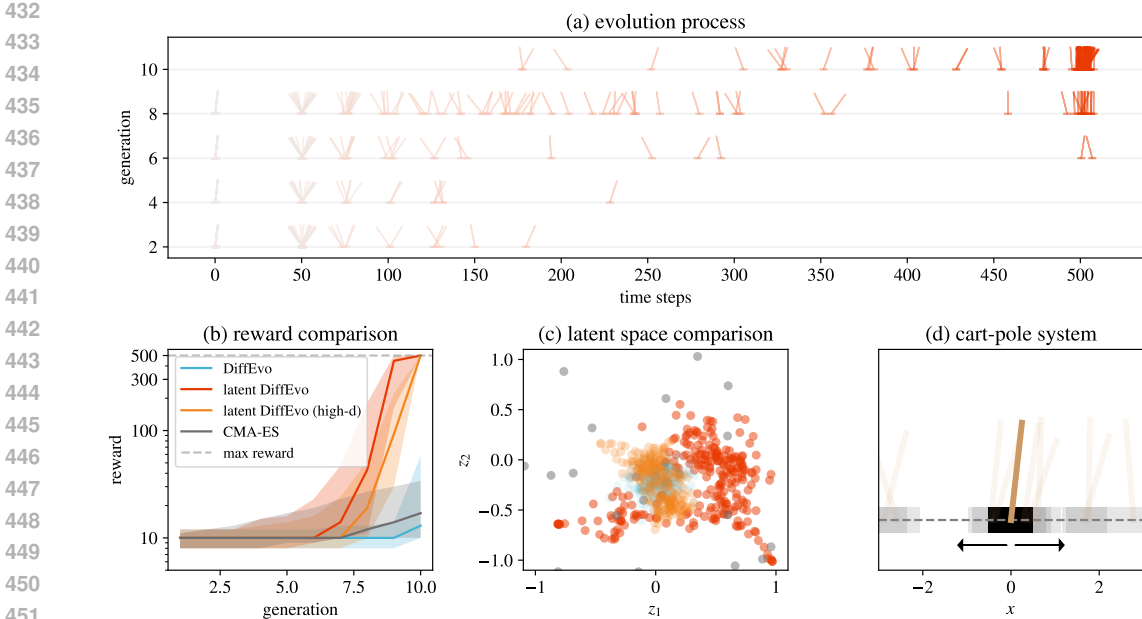


Figure 4: **(a)** Evolution process of cart-pole tasks: The horizontal axis shows survival time, and the vertical axis represents generations. Each point indicates an individual’s state (pole angle, cart shift) at their final survival. As the evolution progresses, more systems survive longer and achieve higher rewards. **(b)** Compared to the original Diffusion Evolution (blue), the Latent Space Diffusion Evolution method (red) significantly improves performance, while the CMA-ES method (gray) fails to find any solutions in given generations. This latent method can even be applied to high-dimensional spaces (orange), with dimensions as high as 17,410. Each experiment is repeated 100 times, with medians (solid lines) and ranges (25% to 75% quantile) shown as shaded areas. **(c)** Projecting the parameters of individuals into a latent space visualize their diversity. The same projection is used for all results (except for the high-dimensional experiment, which has a different original dimension). This indicates enhanced diversity with the latent method. **(d)** The cart-pole system consists of a pole hinged to the cart. And the controller balances the pole by moving the cart left or right.

suggested by the Johnson-Lindenstrauss lemma (Johnson, 1984). To do this, we change Equation 11 to:

$$\hat{x}_0(\mathbf{x}_t, \alpha, t) = \sum_{\mathbf{x} \sim p_{\text{eval}}(\mathbf{x})} p(\mathbf{x}|\mathbf{x}_t)\mathbf{x} = \frac{1}{Z} \sum_{\mathbf{x} \sim p_{\text{eval}}(\mathbf{x})} g[f(\mathbf{x})]\mathcal{N}(z_t; \sqrt{\alpha_t}\mathbf{z}, 1 - \alpha_t)\mathbf{x}, \quad (12)$$

where $\mathbf{z} = \mathbf{E}\mathbf{x}$, $\mathbf{E}_{ij} \sim \mathcal{N}^{(d,D)}(0, 1/D)$, D is the dimension of \mathbf{x} , and d is the dimension of latent space (Johnson, 1984), see Algorithm 2 in Appendix. Here we choose $d = 2$ in our experiments for better visualization. The results show a significant improvement in both fitness and diversity, see Figure 4(b-c). We also found that this latent evolution can still operate in a much larger dimensional parameter space, utilizing a three-layer neural network with 17,410 parameters, while still achieving strong performance. Combined with accelerated sampling method, we can solve the cart pole task in only 10 generations, with 512 population size, one fitness evaluation per individual. These highlights the potential of using tools and theories from the diffusion model domain in evolutionary optimization tasks and *vice versa*, opening up broad opportunities to improve and understand evolution from a new perspective.

5 DISCUSSION

By aligning diffusion models with evolutionary processes, we demonstrate that *diffusion models are evolutionary algorithms, and evolution can be viewed as a generative process*. The Diffusion Evolution process inherently includes mutation, selection, hybridization, and reproductive isolation,

486 indicating that diffusion and evolution are *two sides of the same coin*. Our Diffusion Evolution algo-
 487 rithm leverages this theoretical connection to improve solution diversity without compromise quality
 488 too much compared to standard approaches. By integrating latent space diffusion and accelerated
 489 sampling, our method scales to high-dimensional spaces, enabling the training of neural network
 490 agents in reinforcement learning environments with exceptionally short training time.

491 This equivalence between the two fields offers valuable insights from both deep learning and evolu-
 492 tionary computation. Through the lens of machine learning, the evolutionary process can be viewed
 493 as nature’s way of learning and optimizing strategies for survival of species over generations. Sim-
 494 ilarly, our Diffusion Evolution algorithm iteratively refines estimation of high-fitness parameters,
 495 continuously learning and adapting to the fitness landscape. This positions evolutionary algorithms
 496 and diffusion models not merely as optimization tools, but also as learning frameworks that enhance
 497 understanding and functionality through iterative refinement. Conversely, framing evolution as a dif-
 498 fusion process offers a concrete mathematical formulation. In contrast to previous approaches (Ao,
 499 2005), we provide an explicit and implementable evolutionary framework.

500 The connection between diffusion and evolution enables mutual contributions between the two
 501 fields. Diffusion models are extensively studied in the contexts of controlling, optimization, and
 502 probability theory, offering robust tools to analyze and enhance evolutionary algorithms. In our ex-
 503 periments, leveraging concepts from diffusion models enabled flexible strategies while maintaining
 504 the effectiveness of evolutionary processes. For instance, accelerated sampling methods (Nichol
 505 & Dhariwal, 2021) can be applied seamlessly to Diffusion Evolution to accelerate the optimiza-
 506 tion process. Latent diffusion models (Rombach et al., 2022) inspired our Latent Space Diffusion
 507 Evolution, enabling evolution in high-dimensional spaces and substantially improving performance.
 508 Other advancements in the diffusion model field hold the potential to enhance our understanding
 509 of evolutionary processes. For instance, non-Gaussian noise diffusion models (Bansal et al., 2024),
 510 discrete denoising diffusion models, and theoretical studies of diffusion models through the lens of
 511 non-equilibrium thermodynamics, such as Langevin dynamics (Song et al., 2020b) or spontaneous
 512 symmetry breaking (Raya & Ambrogioni, 2024), unveil entirely new possibilities and perspectives
 513 for understanding and advancing evolutionary methods.

514 However, this parallel we draw here between evolution and diffusion models also gives rise to several
 515 challenges and open questions. While diffusion models, by design, have a finite number of sampling
 516 steps, evolution is inherently open-ended. How can Diffusion Evolution be adapted to support open-
 517 ended evolution? Could other diffusion model implementations yield different evolutionary methods
 518 with diverse and unique features? Can advancements in diffusion models help introduce inductive
 519 biases into evolutionary algorithms? How do latent diffusion models correlate with neutral genes?
 520 Additionally, can insights from the field of evolution enhance diffusion models? These questions
 521 highlight the potential of this duality and synergy between diffusion and evolution.

522 REFERENCES

- 523 Ping Ao. Laws in darwinian evolutionary theory. *Physics of life Reviews*, 2(2):117–156, 2005.
- 524 James Mark Baldwin. A new factor in evolution. *Diacronia*, pp. 1–13, 2018.
- 525 Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie Li, Hamid Kazemi, Furong Huang, Micah Gold-
 526 blum, Jonas Geiping, and Tom Goldstein. Cold diffusion: Inverting arbitrary image transforms
 527 without noise. *Advances in Neural Information Processing Systems*, 36, 2024.
- 528 Andrew G Barto, Richard S Sutton, and Charles W Anderson. Neuronlike adaptive elements that can
 529 solve difficult learning control problems. *IEEE transactions on systems, man, and cybernetics*,
 530 pp. 834–846, 1983.
- 531 Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe
 532 Taylor, Troy Luhman, Eric Luhman, et al. Video generation models as world simulators, 2024.
- 533 Ricky TQ Chen, Jens Behrmann, David K Duvenaud, and Jörn-Henrik Jacobsen. Residual flows for
 534 invertible generative modeling. *Advances in Neural Information Processing Systems*, 32, 2019.
- 535 Aaron C Courville, Nathaniel D Daw, and David S Touretzky. Bayesian theories of conditioning in
 536 a changing world. *Trends in cognitive sciences*, 10(7):294–300, 2006.

- 540 Charles Darwin. The origin of species, 1859-1959. *Bios*, 30(2):67–72, 1959.
- 541
- 542 Richard Dawkins. *The blind watchmaker: Why the evidence of evolution reveals a universe without*
543 *design*. WW Norton & Company, 1996.
- 544 Richard Dawkins. *The selfish gene*. Oxford university press, 2016.
- 545
- 546 Peter Dayan and L. F. Abbott. *Theoretical neuroscience: computational and mathematical modeling*
547 *of neural systems*. Computational neuroscience. Massachusetts Institute of Technology Press,
548 2001. ISBN 978-0-262-04199-7.
- 549 Daniel C Dennett. Darwin’s dangerous idea. *The Sciences*, 35(3):34–40, 1995.
- 550
- 551 Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv*
552 *preprint arXiv:1605.08803*, 2016.
- 553
- 554 Chris Fields and Michael Levin. Does evolution have a target morphology? *Organisms. Journal of*
555 *Biological Sciences*, 4(1):57–76, 2020.
- 556 David E Golberg. Genetic algorithms in search, optimization, and machine learning. *Addion wesley*,
557 1989(102):36, 1989.
- 558
- 559 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
560 Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information*
561 *processing systems*, 27, 2014.
- 562 Stephen Jay Gould. *The structure of evolutionary theory*. Harvard university press, 2002.
- 563
- 564 John J Grefenstette. Genetic algorithms and machine learning. In *Proceedings of the sixth annual*
565 *conference on Computational learning theory*, pp. 3–4, 1993.
- 566
- 567 Ryan N Gutenkunst, Joshua J Waterfall, Fergal P Casey, Kevin S Brown, Christopher R Myers, and
568 James P Sethna. Universally sloppy parameter sensitivities in systems biology models. *PLoS*
569 *computational biology*, 3(10):e189, 2007.
- 570 Nikolaus Hansen. The cma evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772*, 2016.
- 571
- 572 Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution
573 strategies. *Evolutionary computation*, 9(2):159–195, 2001.
- 574 Nikolaus Hansen, Sibylle D Müller, and Petros Koumoutsakos. Reducing the time complexity of
575 the derandomized evolution strategy with covariance matrix adaptation (cma-es). *Evolutionary*
576 *computation*, 11(1):1–18, 2003.
- 577
- 578 Benedikt Hartl, Sebastian Risi, and Michael Levin. Evolutionary implications of self-assembling
579 cybernetic materials with collective problem-solving intelligence at multiple scales. *Entropy*, 26
580 (7):532, 2024.
- 581 Geoffrey E Hinton, Steven J Nowlan, et al. How learning can guide evolution. *Complex systems*, 1
582 (3):495–502, 1987.
- 583
- 584 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint*
585 *arXiv:2207.12598*, 2022.
- 586
- 587 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*
588 *neural information processing systems*, 33:6840–6851, 2020.
- 589
- 590 John H Holland. *Adaptation in natural and artificial systems: an introductory analysis with appli-*
591 *cations to biology, control, and artificial intelligence*. MIT press, 1992.
- 592
- 593 John H Holland. *Emergence: From chaos to order*. OUP Oxford, 2000.
- William B Johnson. Extensions of lipshitz mapping into hilbert space. In *Conference modern*
analysis and probability, 1984, pp. 189–206, 1984.

- 594 Eric R. Kandel. *Principles of Neural Science, Fifth Edition*. McGraw Hill Professional, 2013. ISBN
595 978-0-07-139011-8.
- 596
- 597 Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- 598
- 599 Kostas Kouvaris, Jeff Clune, Loizos Kounios, Markus Brede, and Richard A Watson. How evolu-
600 tion learns to generalise: Using the principles of learning theory to understand the evolution of
601 developmental organisation. *PLoS computational biology*, 13(4):e1005358, 2017.
- 602 Joel Lehman and Kenneth O Stanley. Abandoning objectives: Evolution through the search for
603 novelty alone. *Evolutionary computation*, 19(2):189–223, 2011.
- 604
- 605 Michael Levin. Technological approach to mind everywhere: an experimentally-grounded frame-
606 work for understanding diverse bodies and minds. *Frontiers in systems neuroscience*, 16:768201,
607 2022.
- 608 Michael Levin. Darwin’s agential materials: evolutionary implications of multiscale competency in
609 developmental biology. *Cellular and Molecular Life Sciences*, 80(6):142, 2023.
- 610
- 611 Kevin J Mitchell and Nick Cheney. The genomic code: The genome instantiates a generative model
612 of the organism. *arXiv preprint arXiv:2407.15908*, 2024.
- 613
- 614 Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models.
615 In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.
- 616
- 617 Daniel A Power, Richard A Watson, Eörs Szathmáry, Rob Mills, Simon T Powers, C Patrick Don-
618 caster, and Blažej Czapp. What can ecosystems learn? expanding evolutionary ecology with
learning theory. *Biology direct*, 10:1–24, 2015.
- 619
- 620 Gabriel Raya and Luca Ambrogioni. Spontaneous symmetry breaking in generative diffusion mod-
621 els. *Advances in Neural Information Processing Systems*, 36, 2024.
- 622
- 623 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
624 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-
ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- 625
- 626 Robert Rosen. *Life itself: a comprehensive inquiry into the nature, origin, and fabrication of life*.
627 Columbia University Press, 1991.
- 628
- 629 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.
- 630
- 631 Frank Sehnke, Christian Osendorfer, Thomas Rückstieß, Alex Graves, Jan Peters, and Jürgen
632 Schmidhuber. Parameter-exploring policy gradients. *Neural Networks*, 23(4):551–559, 2010.
- 633
- 634 John Maynard Smith and Eors Szathmary. *The major transitions in evolution*. OUP Oxford, 1997.
- 635
- 636 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
637 learning using nonequilibrium thermodynamics. In *International conference on machine learn-
ing*, pp. 2256–2265. PMLR, 2015.
- 638
- 639 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv
preprint arXiv:2010.02502*, 2020a.
- 640
- 641 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
642 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint
arXiv:2011.13456*, 2020b.
- 643
- 644 Richard S Sutton and Andrew G Barto. Reinforcement learning: an introduction mit press. *Cam-
645 bridge, MA*, 22447:10, 1998.
- 646
- 647 Eörs Szathmáry. Toward major evolutionary transitions theory 2.0. *Proceedings of the National
Academy of Sciences*, 112(33):10104–10111, 2015.

648 Vitaly Vanchurin, Yuri I Wolf, Mikhail I Katsnelson, and Eugene V Koonin. Toward a theory
649 of evolution as multilevel learning. *Proceedings of the National Academy of Sciences*, 119(6):
650 e2120037119, 2022.

651 Pradnya A Vikhar. Evolutionary algorithms: A critical review and its future prospects. In *2016*
652 *International conference on global trends in signal processing, information computing and com-*
653 *munication (ICGTSPICC)*, pp. 261–265. IEEE, 2016.

654 Andreas Wagner. *Arrival of the fittest: How nature innovates*. Current, 2015.

655 Kai Wang, Zhaopan Xu, Yukun Zhou, Zelin Zang, Trevor Darrell, Zhuang Liu, and Yang You.
656 Neural network diffusion. *arXiv preprint arXiv:2402.13144*, 2024.

657 Richard Watson and Michael Levin. The collective intelligence of evolution and development. *Col-*
658 *lective Intelligence*, 2(2):26339137231168355, 2023.

659 Richard A Watson and Eörs Szathmáry. How can evolution learn? *Trends in ecology & evolution*,
660 31(2):147–157, 2016.

661 Richard A Watson, Rob Mills, CL Buckley, Kostas Kouvaris, Adam Jackson, Simon T Powers,
662 Chris Cox, Simon Tudge, Adam Davies, Loizos Kounios, et al. Evolutionary connectionism:
663 algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco
664 and evolutionary transitions. *Evolutionary biology*, 43:553–581, 2016.

665 Richard A Watson, Michael Levin, and Christopher L Buckley. Design for an individual: con-
666 nectionist approaches to the evolutionary transitions in individuality. *Frontiers in Ecology and*
667 *Evolution*, 10:823588, 2022.

668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 A APPENDIX

703 A.1 TWO-PEAKS MODELS

704 We first apply our method to a simple two-dimensional fitness function, with two optimal points,
705 to demonstrate its behavior and capability of finding multiple solutions. We choose a continuous
706 mixed Gaussian density function. The fitness function is a mixed Gaussian density function with
707 means located at $(1, 1)$ and $(-1, -1)$. The fitness function is:
708

$$709 f(x, y) = [\mathcal{N}((x, y); \boldsymbol{\mu}_1, \sigma^2) + \mathcal{N}((x, y); \boldsymbol{\mu}_2, \sigma^2)] / 2, \quad (13)$$

710 where $\boldsymbol{\mu}_1 = (1, 1)$ and $\boldsymbol{\mu}_2 = (-1, -1)$. And the $\sigma = 0.1$.
711

712 A.2 ALPHA AND NOISE SCHEDULE

713 In our experiments, we tested three different noise schedules for α_t . The first is a simple linear
714 schedule, used in Figure 2:

$$715 \alpha_t = 1 - \frac{1}{T}. \quad (14)$$

716 The second is the schedule used in DDPM, which can be approximated by:

$$717 \alpha_t = \exp\left(-\beta_0 t - \frac{\gamma t^2}{T}\right), \quad (15)$$

718 where β_0 and γ are hyperparameters. These are calculated by constraining $\alpha_0 = 1 - \varepsilon$ and $\alpha_T = \varepsilon$,
719 with $\varepsilon = 10^{-4}$ as the default.
720

721 The third schedule is the cosine schedule proposed by Nichol & Dhariwal (2021) and is used in both
722 Figures 3 and 4:

$$723 \alpha_t = \frac{1}{2} \cos\left(\frac{\pi t}{T}\right) + \frac{1}{2}. \quad (16)$$

724 For σ_t , we follow the DDIM setting with a slight modification for better control:

$$725 \sigma_t = \sigma_m \sqrt{\frac{1 - \alpha_{t-1}}{1 - \alpha_t}} \sqrt{1 - \frac{\alpha_t}{\alpha_{t-1}}}, \quad (17)$$

726 where $0 \leq \sigma_m \leq 1$ is a hyperparameter to control the magnitude of noise. We use $\sigma_m = 1$ for
727 most experiments, except for the experiment demonstrating the process in Figure 2, which requires
728 a lower noise magnitude $\sigma_m = 0.1$ for better visualization.
729

730 A.3 NEIGHBOR OF INDIVIDUALS

731 In Figure 2, we use blue discs to represent the neighbors of each individual. Its mean and standard
732 deviation can be derived from Equation 8. By transforming it into an equivalent form with \mathbf{x} as the
733 variable and \mathbf{x}_t as the parameter, we have:

$$734 \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t} \mathbf{x}, 1 - \alpha_t) = \frac{1}{\sqrt{\alpha_t}} \mathcal{N}\left(\mathbf{x}; \frac{\mathbf{x}_t}{\sqrt{\alpha_t}}, \sqrt{\frac{1 - \alpha_t}{\alpha_t}}\right). \quad (18)$$

735 Hence, this Gaussian term can be transformed into a form where \mathbf{x} is the variable, which is more
736 intuitive for describing \mathbf{x} as the neighbors of $\mathbf{x}_t / \sqrt{\alpha_t}$. Thus, the discs have $\boldsymbol{\mu} = \mathbf{x}_t / \sqrt{\alpha_t}$ and
737 $r^2 = (1 - \alpha_t) / \alpha_t$.
738

739 A.4 TWO-DIMENSIONAL FITNESS FUNCTION

740 To benchmark the solution diversity and performance, we choose five different fitness functions to
741 compare our method with other evolutionary strategies. All the functions depend on variables x and
742 y , with the objective being to minimize or maximize the function value. Specifically, we constrain
743 the range of x and y to $(-4, 4)$ and set the objective of the Rastrigin function to be maximization
744

Table 2: Fitness functions used in our experiments.

Name	Formula	features
Rosenbrock	$f(x, y) = 100(y - x^2)^2 + (1 - x)^2$	The minimal value position is $(x, y) = (1, 1)$, where $f(1, 1) = 0$.
Beale	$f(x, y) = (1.5 - x + xy)^2 + (2.25 - x + xy^2)^2 + (2.625 - x + xy^3)^2$	The minimal value position is $(x, y) = (3, 0.5)$, where $f(3, 0.5) = 0$.
Himmelblau	$f(x, y) = (x^2 + y - 11)^2 + (x + y^2 - 7)^2$	This function has four minimal value points, they are: $f(3.0, 2.0) = 0.0$, $f(-2.81, 3.13) = 0.0$, $f(-3.78, -3.28) = 0.0$, $f(3.58, -1.85) = 0.0$
Ackley	$f(x, y) = -20 \exp\left(-0.2\sqrt{\frac{x^2+y^2}{2}}\right) - \exp\left(\frac{\cos 2\pi x + \cos 2\pi y}{2}\right) + e + 20$	When restricting the range of x, y between -4 to 4, the maximal points are located at the four corners.
Rastrigin	$f(x) = An + \sum_{i=1}^2 [x_i^2 - A \cos(2\pi x_i)]$	Here $A = 10$. Similar as the Ackley function above, when restricting the range of x, y , the maximal points are located at the four corners.

instead of minimization to benchmark the capability of finding multiple solutions. To standardize the comparison, we apply the following transformation to convert target values to the highest fitness:

$$F(x, y) = \frac{\epsilon}{\epsilon + \|f(x, y) - f^*\|^2 / s^2} \quad (19)$$

Here, $\epsilon = 10^{-3}$ is used to avoid singular values, and f^* is the target fitness value, with s as the scale factor to make different functions comparable. For each fitness function, we determine f^* as the minimal or maximal value, depending on the optimization objective. The scale factor is determined by the standard deviation of fitness around their optimal points, with ranges adjusted for different functions. After this transformation, the lowest fitness is near zero, and the highest fitness is one.

A.4.1 ESTIMATING ENTROPY TO QUANTIFY DIVERSITY

To quantify the diversity of the solutions, we divided the 2-D space into 80×80 grids and counted the frequencies of elite solutions within this space. We intentionally used this simple and coarse method to quantify entropy in order to eliminate the contribution of local diversity, focusing solely on the diversity of solutions across different basins. The entropy is calculated by:

$$H = \sum_{i=1}^N P_i \log_2 P_i, \quad (20)$$

where P_i is the probability of having a sample in grid i .

810 A.5 CART-POLE EXPERIMENT

811 A.5.1 NEURAL NETWORK

812 The controller of the cart-pole system has four observational inputs: the current position, velocity,
813 pole angle, and pole angular velocity. The system accepts two actions: move left or right. To model
814 the controller, we use artificial neural networks with an input layer of 4 neurons corresponding to
815 the four observations and an output layer of 2 neurons corresponding to the two actions. The action
816 is determined by which output neuron has the higher value.
817

818 Our standard experiment uses a one-hidden-layer neural network with the hidden layer of 8 neurons,
819 resulting in $(4 \times 8 + 8) + (8 \times 2 + 2) = 58$ parameters. We also use a deeper neural network with two
820 hidden layers (each has 128 neurons), totaling $(4 \times 128 + 128) + (128 \times 128 + 128) + (128 \times 2 + 2) =$
821 17410 parameters. Both neural networks use the ReLU activation function.
822

823 A.6 LATENT SPACE DIFFUSION EVOLUTION

824 Following is the pseudocode for Latent Space Diffusion Evolution algorithm. The difference from
825 the original Diffusion Evolution (Algorithm 1) is indicated in [blue](#).
826
827

828 **Algorithm 2** Latent Space Diffusion Evolution

829 **Require:** Population size N , parameter dimension D , latent space dimension d , fitness function f ,
830 density mapping function g , total evolution steps T , diffusion schedule α and noise schedule σ .

831 **Ensure:** $\alpha_0 \sim 1, \alpha_T \sim 0, \alpha_i > \alpha_{i+1}, 0 < \sigma_i < \sqrt{1 - \alpha_{i-1}}$

832 1: $\mathbf{E}^{(d,D)} \leftarrow \mathcal{N}(0, 1/D)$ ▷ Initialize the random mapping
833 2: $[\mathbf{x}_T^{(1)}, \mathbf{x}_T^{(2)}, \dots, \mathbf{x}_T^{(N)}] \leftarrow \mathcal{N}(0, I^{N \times D})$ ▷ Initialize population
834 3: **for** $t \in [T, T - 1, \dots, 2]$ **do**
835 4: $\forall i \in [1, N] : Q_i \leftarrow g[f(\mathbf{x}_t^{(i)})]$ ▷ Fitness are cached to avoid repeated evaluations
836 5: $\forall i \in [1, N] : \mathbf{z}_t^{(i)} \leftarrow \mathbf{E}\mathbf{x}_t^{(i)}$ ▷ Encode individual parameters into latent space
837 6: **for** $i \in [1, 2, \dots, N]$ **do**
838 7: $Z \leftarrow \sum_{j=1}^N Q_j \mathcal{N}(\mathbf{z}_t^{(i)}; \sqrt{\alpha_t} \mathbf{z}_t^{(j)}, 1 - \alpha_t)$
839 8: $\hat{\mathbf{x}}_0 \leftarrow \frac{1}{Z} \sum_{j=1}^N Q_j \mathcal{N}(\mathbf{z}_t^{(i)}; \sqrt{\alpha_t} \mathbf{z}_t^{(j)}, 1 - \alpha_t) \mathbf{x}_t^{(j)}$
840 9: $\mathbf{w} \leftarrow \mathcal{N}(0, I^D)$
841 10: $\mathbf{x}_{t-1}^{(i)} \leftarrow \sqrt{\alpha_{t-1}} \hat{\mathbf{x}}_0 + \sqrt{1 - \alpha_{t-1} - \sigma_t^2} \frac{\mathbf{x}_t^{(i)} - \sqrt{\alpha_t} \hat{\mathbf{x}}_0}{\sqrt{1 - \alpha_t}} + \sigma_t \mathbf{w}$
842 11: **end for**
843 12: **end for**
