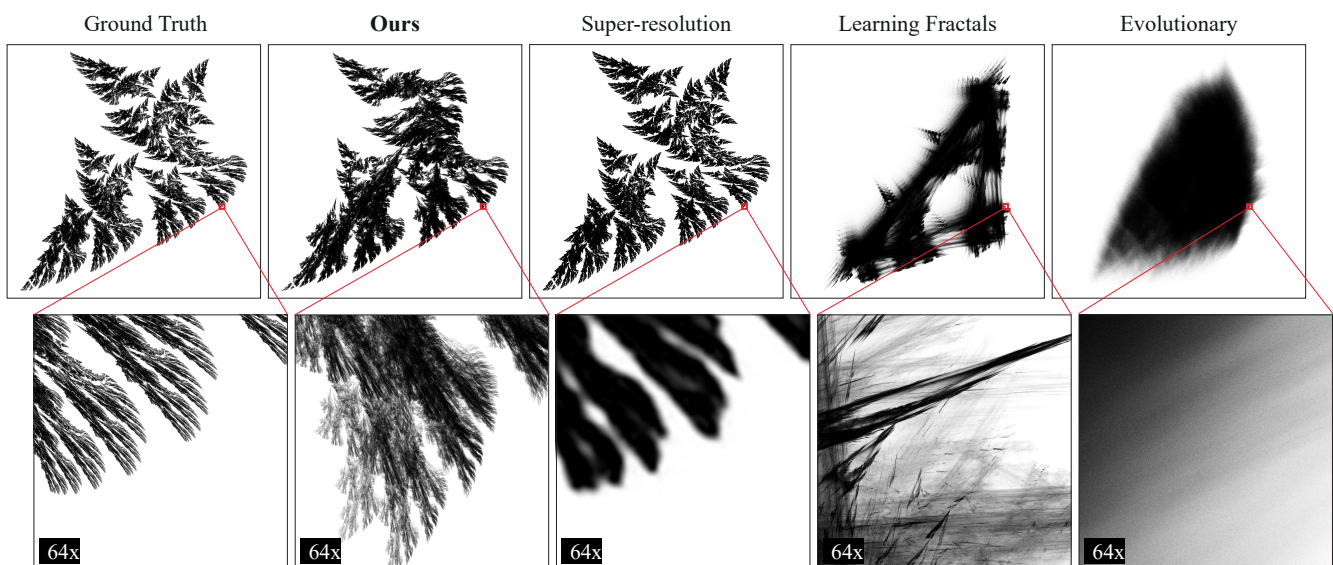


# Learning Image Fractals Using Chaotic Differentiable Point Splatting

A. Djeacoumar<sup>ID</sup>, F. Mujkanovic<sup>ID</sup>, H.-P. Seidel<sup>ID</sup>, T. Leimkühler<sup>ID</sup>

Max-Planck-Institut für Informatik, Saarbrücken, Germany



**Figure 1:** We introduce a novel method to recover a fractal description from an image containing a self-similar shape. Our hybrid optimization achieves state-of-the-art fractal inversion, enabling the synthesis of intricate details at any desired scale – illustrated here with 64x zoom-ins.

## Abstract

Fractal geometry, defined by self-similar patterns across scales, is crucial for understanding natural structures. This work addresses the fractal inverse problem, which involves extracting fractal codes from images to explain these patterns and synthesize them at arbitrary finer scales. We introduce a novel algorithm that optimizes Iterated Function System parameters using a custom fractal generator combined with differentiable point splatting. By integrating both stochastic and gradient-based optimization techniques, our approach effectively navigates the complex energy landscapes typical of fractal inversion, ensuring robust performance and the ability to escape local minima. We demonstrate the method's effectiveness through comparisons with various fractal inversion techniques, highlighting its ability to recover high-quality fractal codes and perform extensive zoom-ins to reveal intricate patterns from just a single image.

## CCS Concepts

• **Computing methodologies** → **Point-based models; Rendering; Machine learning;**

## 1. Introduction

The geometry of nature is often fractal [Man82]: Structures repeat themselves across different scales, forming self-similar patterns. A classic example is a fern, where each branch resembles both the overall fern structure and its smaller sub-branches (Fig. 2). Simi-

lar observations can be made for coastlines, terrains, trees, river networks, blood vessels, etc. Such intricate patterns can be explained by the repeated application of simple rules, referred to as *fractal codes*. Knowledge of underlying fractal codes is essential not only for describing and understanding complex sys-

tems [Man82, HBG\*95, Tur97, SHM05, Man97] but also as an effective means of data compression [Jac92, Fis94]. In this work, we develop a novel method for extracting fractal codes from an image containing a self-similar pattern using modern differentiable point splatting [KKLD23].

The *fractal inverse problem* – finding fractal codes that generate a given geometry – is a long-standing formidable challenge [Bar88] and typically framed as an optimization problem [Vrs91, CLV\*95, TCCC23]. This presents two key challenges.

First, both the large-scale recursive application of fractal codes and the rendering of high-quality images are computationally demanding, especially within an iterative optimization framework. We observe that large-scale fractal synthesis, i.e., recursively executing fractal codes hundreds of thousands of times per image, and well-designed rendering, including proper anti-aliasing, are both essential for achieving high-quality inversion results. While computer graphics research has explored the efficient forward rendering of fractals over the past few decades [Car80, HD91, Mar10, dSNLV21], there has been limited focus on improving efficiency in the inverse setting. To address this gap, we introduce a highly efficient parallel fractal generator, paired with a custom differentiable renderer. Specifically, we propose a model that optimizes the parameters of an Iterated Function System (IFS) [Hut81] to produce a point set that is differentially rasterized [ZPVBG01, KKLD23]. Our model enables scaling up fractal synthesis to a quality level that facilitates state-of-the-art fractal inversion results.

The second key challenge in fractal inversion is the low dimensionality of the parameters, which results in highly non-convex objective energy landscapes that are notoriously difficult to navigate. To address this challenge, stochastic optimization techniques, such as evolutionary algorithms [NG94, QIG17], have been employed. While these methods are effective at escaping local minima, they often follow inefficient optimization trajectories. In a different line of work, gradient-based optimization has been explored to solve the fractal inverse problem [MP98, VS99, TCCC23, BN24]. Gradients provide a strong guiding signal that enhances efficiency, but they are inherently prone to converging to local minima. In this work, we make the observation that leveraging the synergies between stochastic and gradient-based techniques significantly outperforms all previous methods in solving the fractal inverse problem. Our optimization employs an interleaved scheme that alternates between gradient descent [KB15] and simulated annealing [KGJV83], ensuring stable optimization trajectories while maintaining the ability to escape local minima.

Our approach recovers high-quality fractal codes from images. The key property of these codes is their ability to synthesize highly detailed patterns with virtually infinite resolution. These generative modeling capabilities allow for extreme zoom-ins on the original

image, revealing intricate self-similarities and complex structures at any scale (Fig. 1). We conduct extensive comparisons with both classical and modern fractal inversion techniques, demonstrating the clear superiority of our method.

In summary, our contributions are:

- A novel, highly efficient image fractal inversion technique based on differentiable point splatting.
- An effective optimization routine that relies on synergies between stochastic and gradient-based optimization.
- State-of-the-art fractal inversion results alongside an in-depth analysis of the proposed algorithm and its components.

## 2. Related Work

### 2.1. Fractals

Fractals are complex geometric patterns that repeat infinitely at different scales, revealing self-similarity through recursive or iterative processes [Man80, Man82, PJSF04]. Fractal geometry is not only crucial in many scientific disciplines [HBG\*95, Tur97, SHM05, Man97], but it has also been applied to the study of natural images [Pen84, TPRC00] and visual art [RHD08], and is an indispensable tool for synthetic scene creation [DHN85, BJM\*88, Ebe02, TGSC24].

A wide variety of algorithms for creating fractals and self-similarities exists. Grammar-based techniques use formal rules to iteratively generate complex structures [Lin68, RS80] with diverse applications [SRDT01, WWSR03, PL12]. Noise-based approaches combine random functions at different scales [MVN68, Kes82]. Escape-time algorithms generate fractals by iterating a function at each point until it escapes a boundary [MEG04, Jul18]. Furthermore, self-similarity can be achieved by tiling [Ban91, Fat01, OCNG21, CZHN22]. We build our approach on Iterated Function Systems, which generate fractals by repeatedly applying a set of geometric transformations [Hut81, BD85, Bar88, Elt87]. This method, which we review in Sec. 3, is well-suited for point-based rendering [TCCC23, BN24] and has recently been shown to combine effectively with other fractal generation paradigms [SK24].

Fractals have also been studied in 3D [HSK89, Nor82], where efficient rendering is crucial [Car80, HD91, Mar10, dSNLV21]. Similarly, our optimization scheme relies on efficient synthesis and rendering to produce high-quality fractal images – a property we found essential for achieving state-of-the-art inversions.

Our approach incorporates a deep differentiable pipeline. Fractals have been integrated into deep learning in various ways, such as through self-similar architectures [LYMF19], reframing computations as neural networks [Sta91], deep equilibrium models [BKK19], and training data augmentation [KOM\*20, AF22].

### 2.2. Fractal Inversion and Procedural Modeling

Obtaining fractal codes from an image is a notably intricate process [BEHL86, Bar88, Vrs91]. Since direct methods are only feasible for simple synthetic cases [Ber97, SDG\*93], the task is generally formulated as an optimization problem. Stochastic techniques have been explored to address the highly non-convex energy landscapes involved. These include evolutionary algorithms optimizing the



**Figure 2:** A fern exhibiting self-similarities.

fractal code directly [LVL93, NG94, Lan95, GCI00], a partitioning of the fractal [SB06], or a polar reparameterization [CLRS00]. Similarly, swarm intelligence algorithms solve the problem by modeling stochastic natural processes [QIG17, GI18, GID\*21]. Additionally, quadratic programming [FV94] and expectation maximization [BdR17] have been applied to solve this task.

In a separate line of research, gradient-based approaches have been explored [VS99]. Differentiability has been achieved by using point set distances [MP98], moments [VR89, RZ94], signed distance functions [Kim15], and, recently, differentiable point splatting [TCCC23, BN24, Sco24]. Additionally, experiments have been conducted on directly regressing fractal codes using neural networks [GD21]. The approach most similar to ours is that of Tu et al. [TCCC23]. However, unlike their focus on representing low-resolution images with points arising from an Iterated Function System, we aim to recover infinite-resolution fractal structures. To achieve this, we utilize high-quality renderings produced by an efficient forward model, along with a hybrid optimization framework to prevent getting stuck in local minima.

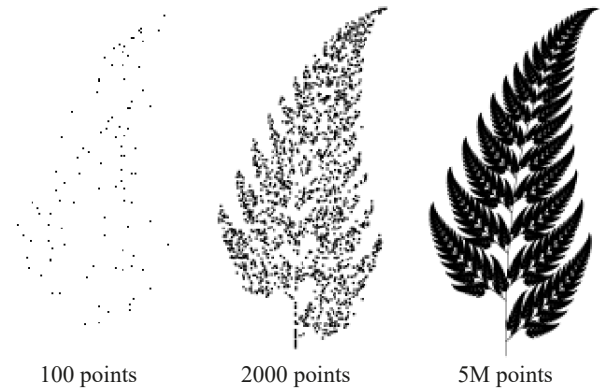
In addition to applications in data compression [Jac92, BH93, Fis94] and symmetry detection [LHOK\*10, LSS\*17], discovering self-similarities has also been studied as a powerful tool for generative modeling [PXM\*22, KRWM22, ZWSW23, Mer23]. Inverse procedural generative models have been studied in various forms, including tiling [VSLD13], noise models [GSV\*14, HN18], handcrafted spatial patterns [LP00], handcrafted material-specific procedures [GHYZ20], and both fixed [HDR19, SLH\*20] and learnable [HHD\*22] procedural node graphs. However, to our knowledge, little to no inverse procedural models generate details at infinite scales, thus most of them lose detail when magnifying way past the scale of the reference image. In contrast, in our work we learn a generative model that automatically extrapolates the target structure from a single scale to infinitely finer scales, constituting a significant advantage for infinite-resolution image synthesis.

### 2.3. Multiscale Image Representation and Reconstruction

Representing, synthesizing, and reconstructing images at multiple scales is a fundamental task in computer graphics and computer vision.

Coarser representations of an image can be formalized using linear scale-space theory [Iij59, Wit87] and efficiently implemented using pyramids [Bur81, Wil83]. In recent years, neural multiscale representations have gained significant attention [FSWK20, CLW21, PNS\*22, LVVPW22, STB\*22, MESK22, BGF\*23, MNT\*24]. In contrast to our approach, all these approaches assume that a maximum resolution exists and a corresponding image is available. However, we leverage image pyramids as part of our multiscale supervision structure.

The reverse task, inferring finer-scale images from coarser ones, is an active area of research in the super-resolution community. For an overview, we refer to a recent survey [MRF\*23]. Similar to our setting, the internal statistics of an image in the form of recurring patches has been studied [SI07, ZI11] to increase resolution [GBI09, SCI18, BKSI19], remove blur [MI14], or synthesize new layouts [SDM19, SBII19, ZZB\*18]. Unlike these methods, we



**Figure 3:** An IFS fractal generated using the chaos game (Eq. 2) with varying numbers of points.

model self-similarity using analytic functions. This results in a narrower application domain, but enables image synthesis without scale limits.

Another approach to obtaining a multiscale representation is to fuse images from different scales. This is often done with an understanding of how the source images are related to each other [KHE\*10, HMK\*19, MAAB\*17, TLW\*19, LFS21], but unstructured image collections have also been explored [WDJ\*24]. In contrast, our approach takes a single image as input and infers generative rules to synthesize infinite-resolution details.

### 2.4. Point Splatting

Point-based rendering has a long and rich history [GP11]. Early work focused on efficient hardware-accelerated point samples [GD98], while advancements in “splatting” reconstruction kernels, such as EWA [ZPVBG01], enabled high-quality image synthesis free from aliasing. The emergence of differentiable visual computing [Li19, SZR\*23] has further advanced this field, with soft reconstruction kernels facilitating the solution of inverse problems through differentiable point splatting [WGSJ20, YSW\*19, LZ21]. A notably efficient implementation of this approach is 3D Gaussian Splatting (3DGS) [KKLD23], which marks the current state of the art in scene reconstruction from images. We adapt their splatting framework to recover 2D fractals and pair it with a stochastic optimization routine.

Similar to our work, 3DGS-based reconstruction has been treated as a stochastic process [KRS\*24], akin to Stochastic Gradient Langevin Dynamics [BDM18]. However, we demonstrate that mere gradient perturbations are insufficient to obtain high-quality fractal inversions.

## 3. Background

Here, we briefly review the concept of Iterated Function Systems (IFS) [Hut81, BD85] that implicitly define fractal geometry through the repeated application of geometric transformations. Since this

work focuses on *image* fractals, we restrict our exposition to the 2D domain. For a comprehensive introduction, we refer to Barnsley [Bar88].

An IFS fractal code  $\mathcal{F} = \{(f_i, p_i)\}_{i=1}^N$  is a set of  $N$  continuous functions  $f_i \in \mathbb{R}^2 \rightarrow \mathbb{R}^2$ , each with an associated probability  $p_i \in [0, 1]$ , where  $\sum_i p_i = 1$ . Each function  $f_i$  deterministically moves a point at location  $\mathbf{x}$  to a new location  $f_i(\mathbf{x})$  and is required to be contractive, i.e.,

$$\|f_i(\mathbf{x}_1) - f_i(\mathbf{x}_2)\|_2 < \|\mathbf{x}_1 - \mathbf{x}_2\|_2 \quad \forall \mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^2. \quad (1)$$

Based on this construction and starting from an arbitrary initial point  $\mathbf{x}_0$ , a sequence of points  $\mathbf{x}_k$  can be generated using the ‘‘chaos game’’

$$\mathbf{x}_{k+1} = f_{\pi_k}(\mathbf{x}_k), \quad (2)$$

where the index  $\pi_k$  is randomly selected from the set  $\{i\}_{i=1}^N$ , with the probability of choosing index  $i$  given by  $p_i$ . It can be shown [Hut81, BD85] that the sequence  $\mathbf{x}_k$  obtained through this Markov process converges to a compact set  $\mathcal{A} \subset \mathbb{R}^2$ , the *attractor* of  $\mathcal{F}$ . The attractor  $\mathcal{A}$  acts as a unique fixed point solely determined by the functions  $f_i$ , i.e., once a point  $\mathbf{x}_k$  enters  $\mathcal{A}$ , subsequent applications of  $\mathcal{F}$  will cause the points to jump around within  $\mathcal{A}$ , but they will never leave it again. Accumulating all points  $\mathbf{x}_k$  effectively leads to an increasingly dense coverage of  $\mathcal{A}$  (Fig. 3). The probabilities  $p_i$  do not affect the shape of  $\mathcal{A}$ ; they only influence the convergence rate of the chaos game by modulating local point densities and are typically derived from the determinants of  $f_i$  [TCCC23, AF22]. Importantly,  $\mathcal{A}$  is a fractal by construction, as it consists of a union of (distorted) copies of itself [BD85]. Henceforth, we will refer to sequences  $\mathbf{x}_k$  as *trajectories*.

The above exposition holds for any continuous and contractive functions  $f_i$  [CLV<sup>+</sup>95]. However, in this work we follow standard practice [BEHL86, GCI00, GD21, TCCC23] and restrict ourselves to *affine* functions of the form

$$f_i(\mathbf{x}) = M_i \mathbf{x} + \mathbf{b}_i, \quad (3)$$

with  $M_i \in \mathbb{R}^{2 \times 2}$  and  $\mathbf{b}_i \in \mathbb{R}^2$ .


The fractal inverse problem can now be stated more precisely: Given an input image  $I_{\text{ref}}$  depicting a (fractal) shape, find an IFS fractal code  $\mathcal{F}$  whose attractor  $\mathcal{A}$  matches that shape. Once  $\mathcal{F}$  is obtained, the shape can be re-synthesized at virtually infinite resolution by rendering points generated through the chaos game.

## 4. Method

*Chaos is a friend of mine.*

– Bob Dylan

The input to our algorithm is a binary or grayscale image  $I_{\text{ref}}$ , representing a (fractal) shape of finite resolution or bandwidth, e.g., a raster image. Without loss of generality, we assume the shape is centered within  $I_{\text{ref}}$  and padded by 25% of the image’s dimensions. Our goal is to find an IFS code  $\mathcal{F}$ , such that when the chaos game from Eq. 2 is applied, the resulting attractor  $\mathcal{A}$  closely approximates the shape depicted in  $I_{\text{ref}}$ . Since  $\mathcal{A}$  is a fractal, we can use  $\mathcal{F}$  to



**Figure 4:** Overview of our model. The fractal point generator uses functions in  $\mathcal{F}$  and employs a parallel stochastic scheme to efficiently and robustly execute the chaos game (Eq. 2). The generated points are then rendered into an image using differentiable splatting.

render highly detailed, zoomed-in views at any scale, effectively surpassing the resolution limitations of  $I_{\text{ref}}$  by an arbitrary amount.

We frame our task as an optimization problem over a set of affine functions  $\{f_i\}_{i=1}^N$  (Eq. 3), where the number  $N$  of functions is set to a fixed conservative estimate. As shown in Fig. 4, our model consists of two components. The first component is a fractal point generator, which runs the chaos game to generate a large point set  $\mathcal{P} = \{\mathbf{x}_k\}$  that densely covers and thereby defines  $\mathcal{A}$ . To effectively incorporate this point generator into an optimization framework, we need to implement measures to improve efficiency and avoid overfitting to specific point trajectories, as detailed in Sec. 4.1. The second component is a point splatting module, which differentially rasterizes the points from the first component into an image. This process is described in Sec. 4.2. Our model enables the formulation of an objective function in image space, which we minimize using a hybrid iterative optimizer that specifically addresses local minima via additional stochasticity, as outlined in Sec. 4.3.

### 4.1. Fractal Point Generator

The purpose of this stage is the generation of points on the attractor  $\mathcal{A}$  via the chaos game. We observed that achieving high-quality and stable results requires a large number of points, typically  $|\mathcal{P}| \approx 500\text{k}$  in practice. Therefore, evaluating Eq. 2 recursively for each optimization iteration is (i) highly inefficient because of the lengthy computation sequence, and (ii) unstable due to the need for gradients to flow through such an ultra-deep pipeline.

To address this issue, we use a generator that runs  $b$  trajectories of length  $l$  in parallel (rows in Fig. 4, left). Each trajectory is initialized with a uniformly sampled random point  $\mathbf{x}_0$ . Then, following Eq. 2, we apply a sequence of transformations  $f_i$ , randomly sampled from an uniform distribution to generate the point trajectories. Since  $\mathbf{x}_0$  does not generally lie on  $\mathcal{A}$ , we treat the first  $m$  iterations as a warm-up phase, allowing the trajectory to converge onto  $\mathcal{A}$ . After this warm-up, we begin concatenating the generated points from all trajectories into a large buffer.

In each optimization iteration, our parallel pipeline samples  $b$

**Table 1:** Sources of stochasticity in our framework, along with their granularities, i.e., the levels at which random numbers are drawn.

Component	Source of Stochasticity	Granularity
Model	Point Initialization Function Selection	Trajectory Point
Optimizer	Simulated Annealing	Opt. Iteration

from the  $N^l$  possible combinations of  $f_i$ , producing trajectories that start from randomly initialized positions  $\mathbf{x}_0$ . This approach greatly helps the optimization focus on matching the attractor shape using the IFS rather than overfitting it using specific trajectories. Tab. 1 provides a list of all sources of stochasticity in our point generation pipeline, along with their respective sampling granularities.

After generating all the points, we normalize them by applying a global uniform scaling and offset to center the point cloud and achieve the 25% padding assumed for the shape in the input  $I_{\text{ref}}$ .

**Parameterization** Our affine functions  $f_i$ , as defined in Eq. 3, must satisfy two key properties: (i) They must be contractive, i.e., Eq. 1 needs to hold, and (ii) they should be appropriately constrained to prevent divergence during the highly stochastic optimization process. To achieve this, we design a custom parameterization for the matrices  $M_i$  and offsets  $\mathbf{b}_i$ .

We parameterize  $M_i$  using the singular value decomposition (SVD)  $M_i = U_i S_i V_i^T$ , with  $U_i, S_i, V_i \in \mathbb{R}^{2 \times 2}$ . The diagonal matrix  $S_i$  contains the non-negative singular values of  $M_i$ . Each singular value is represented by a freely optimizable parameter to which a sigmoid function is applied, ensuring contractivity. The orthonormal matrices  $U_i$  and  $V_i$  are represented by four optimizable parameters each, corresponding to their individual matrix entries. After each optimization iteration,  $U_i$  and  $V_i$  are projected onto the nearest orthonormal matrices using the Gram-Schmidt process [CBG\*17], ensuring that the decomposition remains a valid SVD. Note that this represents an over-parameterization of  $M_i$  – with 10 optimizable parameters for a  $2 \times 2$  matrix. The offset  $\mathbf{b}_i \in \mathbb{R}^2$  is parameterized using a freely optimizable 2D vector, with the parameters passed through the hyperbolic tangent function. This ensures that offsets remain within the range  $[-1, 1]$ . We refer to the optimizable parameters of our model collectively as  $\theta \in \mathbb{R}^{12N}$ .

In contrast to recent observations in previous work [BN24, Sco24], we discovered that our pipeline is insensitive to parameter initialization. Therefore, we use a straightforward uniform random initialization of  $\theta$ .

## 4.2. Differentiable Point Splatting

To enable learning fractals from images, we differentially rasterize the point set  $\mathcal{P}$  generated in the previous stage into an image  $I$ . Each point is splatted as an isotropic 2D Gaussian [ZPVBG01] using standard alpha blending [PD84]. For efficiency, we employ a tile-based rasterization pipeline [KKLD23].

Splatting a stochastic approximation of an infinite-resolution fractal onto a finite-resolution image introduces challenges related to

aliasing. We address this issue with super-sampling. Specifically, we splat the points into an image with a spatial resolution that is  $5 \times 5$  higher than the input image  $I_{\text{ref}}$ , using Gaussian splatting kernels with a standard deviation of 2.5 pixels. This is followed by averaging downsampling to the target resolution to arrive at our final output  $I$ . Through a parameter grid search, we found this combination of splat size and super-sampling factor to be an effective balance between detail preservation, anti-aliasing, gradient flow, and efficiency.

## 4.3. Optimization

With our differentiable forward model established, we design an optimization algorithm to recover the fractal codes. Specifically, we optimize for the parameters  $\theta$  of the IFS to generate an attractor image  $I$  that closely matches the input image  $I_{\text{ref}}$ .

**Objective Function** Our image-space objective function consists of four terms, addressing individual aspects of the problem setting:

$$\mathcal{L} = \lambda_{\text{MSE}} \mathcal{L}_{\text{MSE}} + \lambda_{\text{SSIM}} \mathcal{L}_{\text{SSIM}} + \lambda_{\text{LPIPS}} \mathcal{L}_{\text{LPIPS}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}, \quad (4)$$

where  $\lambda_i \in \mathbb{R}_+$  are corresponding weighting factors. The first term,

$$\mathcal{L}_{\text{MSE}} = \sum_k \|\text{mip}_k(I) - \text{mip}_k(I_{\text{ref}})\|_2^2, \quad (5)$$

evaluates the reconstruction across a range of scales using image pyramids [Bur81, Wil83], where  $\text{mip}_k$  denotes the  $k$ 'th linear MIP map level. The terms  $\mathcal{L}_{\text{SSIM}}$  and  $\mathcal{L}_{\text{LPIPS}}$  assess the error between  $I$  and  $I_{\text{ref}}$  using the D-SSIM [WBSS04] and LPIPS [ZIE\*18] metrics, respectively, promoting the alignment of overall shape structure. Although the LPIPS metric has been specifically trained to evaluate *natural* images, we found that  $\mathcal{L}_{\text{LPIPS}}$  significantly improves inversion results. This aligns with recent findings suggesting that pre-training neural networks with fractals is effective [KOM\*20, AF22].

Finally, we found that an additional regularization in the form of

$$\mathcal{L}_{\text{reg}} = \sum_{i=1}^N \sigma_{i,1}^2 + \sigma_{i,2}^2 + \|\mathbf{b}_i\|_2^2 + \lambda_{\text{cond}} \left( \frac{\sigma_{i,1} + 1}{\sigma_{i,2} + 1} - 1 \right)^2 \quad (6)$$

is essential for stable optimization, where  $\sigma_{i,j}$  is the  $j$ -th singular value of  $M_i$  that can be read off directly from  $S_i$ . While the singular values  $\sigma_{i,j}$  and offsets  $\mathbf{b}_i$  are already constrained using saturating functions, we found that the optimization often drives them into the saturation region, leading to unfavorable dynamics. The first three summands of Eq. 6 address this issue by encouraging these values to remain as small as possible. The last summand in Eq. 6 utilizes a stable approximation of the condition number of  $M_i$ , helping to prevent singular matrices, which appear as streaks in the rendered fractal.

We set  $\lambda_{\text{MSE}} = 10$ ,  $\lambda_{\text{SSIM}} = 1$ ,  $\lambda_{\text{LPIPS}} = 2$ ,  $\lambda_{\text{reg}} = 10^{-2}$ , and  $\lambda_{\text{cond}} = 10$  in all our experiments.

**Hybrid Optimizer** While our pipeline allows to compute  $\nabla_{\theta} \mathcal{L}$  using automatic differentiation, we found that a pure gradient-based optimizer does not lead to satisfactory results. This is due to the complex, recursive generation process, which is driven by a relatively small set of variables  $\theta$  that all have a global impact. As a solution, we employ a hybrid iterative optimization strategy that alternates between gradient descent and simulated annealing [KGJV83]. The

gradient descent component provides a stable path toward lower-energy states, while the simulated annealing component helps escape local minima.

Throughout the optimization, we maintain a “temperature” parameter  $\tau$  that decreases linearly from 1 to 0. Gradient descent is implemented using the Adam optimizer [KB15] with a fixed learning rate of  $10^{-2}$  and default parameters otherwise. After every 250 iterations of gradient descent, we save the current parameters, denoted as  $\theta_{\text{curr}}$ , along with the current energy, denoted as  $\mathcal{L}_{\text{curr}}$ , and switch to a phase of simulated annealing. In this phase, we successively sample 10 new parameter candidates  $\theta_{\text{cand}} = \theta_{\text{curr}} + \Delta\theta$ , where perturbations  $\Delta\theta \sim \mathcal{N}(\mathbf{0}, \tau/s)$  are drawn from an isotropic normal distribution. The corresponding energy  $\mathcal{L}_{\text{cand}}$  is then evaluated. A candidate  $\theta_{\text{cand}}$  is accepted as the new state  $\theta_{\text{curr}}$  with an acceptance probability

$$p(\theta_{\text{cand}}) = \min\left(\exp\left(-10 \cdot \frac{\mathcal{L}_{\text{cand}} - \mathcal{L}_{\text{curr}}}{\tau}\right), 1\right). \quad (7)$$

This approach allows the optimization to occasionally move to a higher-energy state, helping to avoid getting trapped in local minima. After each simulated annealing phase, we switch back to a gradient descent phase, resetting the momentum of the Adam optimizer. This alternating scheme continues for the first half of the total optimization iterations, after which we stay in the gradient descent phase until completion.

#### 4.4. Implementation Details

We implemented our framework in PyTorch [PGC\*17], using custom CUDA kernels for the point splatting module. Our source code, fractal data, and supplemental materials are available at <https://chaotic-fractals.mpi-inf.mpg.de>.

Our point generator creates point sets  $\mathcal{P}$  using  $b = 2000$  trajectories in parallel, each with a length of  $l = 250$  points. Using a warm-up phase of  $m = 10$  iterations, this results in 480k points per optimization iteration. We found that this set of hyperparameters generates sufficient points for our optimization; using more points did not enhance inversion results. For rendering, we adapted the differentiable rasterization pipeline from Kerbl et al. [KKLD23] to enable efficient splatting of isotropic 2D kernels. As our approach does not require any sorting of primitives, the asymptotic complexity of splatting reduces from  $\mathcal{O}(|\mathcal{P}| \cdot (c + \log|\mathcal{P}|))$  to  $\mathcal{O}(|\mathcal{P}| \cdot c)$  compared to the original implementation, where  $c$  is the number of pixels in  $I$ . We use a resolution of  $1024 \times 1024$  pixels for  $I_{\text{ref}}$  and  $I$  in all our experiments. A full optimization requires 15k iterations. We provide the pseudocode of our optimization algorithm in the Appendix A.

#### 5. Evaluation

We evaluate our approach by comparing it against a wide range of baselines (Sec. 5.1), presenting additional results (Sec. 5.2), and performing ablation studies (Sec. 5.3). Finally, we discuss the limitations of our method (Sec. 5.4).

By default, fractals for visualization and evaluation are generated with 50 million points per view and a supersampling rate of  $8 \times 8$ . To accelerate convergence, our point generator samples  $f_i$  based on

**Table 2:** Timing breakdown per iteration of our approach on two different GPU models. During optimization, we employ differentiable point splatting, whereas for evaluation, we use hardware-accelerated point rendering with 100x more points.

Component	Optim. (500k pts.)		Eval. (50M pts.)	
	A40	RTX 3090	A40	RTX 3090
Generation	99.5 ms	158.9 ms	4.7 s	5.2 s
Rendering	1.2 ms	1.4 ms	0.1 s	0.1 s
Sim. ann.	2.0 ms	2.7 ms	–	–
Backward	116.6 ms	128.2 ms	–	–

probabilities proportional to their respective determinants [AF22, TCCC23]. Since differentiability is not necessary for evaluation, we use a basic hardware-accelerated point renderer implemented with PyOpenGL. For certain combinations of zoomed-in views and IFS codes, achieving a point count of 50 million can be difficult or impossible, as it may be highly unlikely for points to fall within the small view window. In such cases, we limit the generation time to 20 minutes per view. The corresponding images are labeled with a  $\blacktriangle$ . Tab. 2 presents a breakdown of the synthesis time for our fractal images across two GPU models.

#### 5.1. Comparisons

**Setup** For our evaluation, we use a test suite of 100 randomly generated IFS fractals, following Anderson et al. [AF22]. In addition to randomly sampling the IFS parameters, we also randomly select the number of functions  $N$  from the range  $\{3, 10\}$ .

**Baselines** We compare our approach against a wide spectrum of existing approaches for fractal inversion or high-resolution image synthesis.

First, we consider the recent method Learning Fractals [TCCC23]. Like our approach, this baseline uses differentiable point rendering to optimize IFS parameters but focuses on low-resolution reconstruction within a purely gradient-based framework. Since their default settings – 300 points and an image resolution of  $32 \times 32$  – produced unsatisfactory results, we also present results with an increased point count of 600 and a higher image resolution of  $256 \times 256$ . Since the original implementation is not specifically optimized for efficiency, we were unable to run experiments with higher-quality settings.

Second, we evaluate two stochastic optimization approaches. Evolutionary programming [NG94] optimizes a population of solutions by randomly perturbing them with Gaussian noise to create offspring, and then selecting the fittest through competitive survival. We carefully re-implemented this method. A more recent Cuckoo search approach [QIG17], on the other hand, uses Lévy flight random walks and employs reproductive rules inspired by the egg-laying behavior of cuckoos. For this baseline, we adopted the code provided by the original authors.

Further, we present best-effort results for directly regressing IFS parameters from images using a neural network. In contrast to the approach in the seminal work [GD21], we found that a residual design – regressing deviations from identity transformations – yielded more

stable results. Training of the convolutional network is supervised on 8000 random fractals with IFS codes sorted according to their singular values.

Finally, to further contextualize the capabilities of our approach, we present results from a recent single-image super-resolution method. Specifically, we evaluate SwinIR [LCS\*21], a neural super-resolution architecture that integrates the spatially invariant filters of convolutional layers with Swin transformer [LLC\*21] layers, enabling the model to learn long-range dependencies across the image. We train this model on 10k fractal images to upsample  $1000 \times 1000$  images to a resolution of  $8000 \times 8000$  pixels, following the setup used by the original authors.

For all IFS-based baselines, we assume a conservative number of functions, with  $N = 10$ . To ensure fair comparisons, we utilized tuned hyperparameters for all baselines in cases where they produced better results.

**Metrics** To evaluate reconstruction quality across different scales, we synthesize 6 randomly sampled patches per test fractal, along with the full fractal, using zoom-in factors of up to 8x. All numerical evaluations are based on this extended set of 700 images. To assess shape fidelity, we use the F1 score and the mean intersection over union (IoU). Additionally, we report the standard image quality metrics PSNR, SSIM [WBSS04], and LPIPS [ZIE\*18]. We observe consistent inversion quality across GPU models (Nvidia A40 and RTX 3090), with only a slight increase in optimization time with the RTX 3090 (Tab. 2).

**Results** We present the quantitative results of our analysis in Tab. 3 and corresponding qualitative results in Fig. 5. Please refer to our supplemental material for a complete gallery of qualitative results for the entire test set across methods.

Our approach significantly outperforms all baselines. The original version of Learning Fractals struggles to recover reasonable fractals due to its low-resolution  $32 \times 32$  canvas. When we increase the resolution to  $256 \times 256$  pixels – the maximum supported by our modern hardware – this baseline can recover coarse shapes of the target fractals but fails to capture finer structural details. Despite our extensive efforts and several hours of optimization, stochastic methods using evolutionary programming and Cuckoo search did not converge to accurate results in our test setup. Similarly, neural network-based regression of IFS codes from images failed to produce coherent results, confirming previous findings [GD21] and suggesting that generalizable solutions for the fractal inverse problem may not be feasible at this point. As expected, the super-resolution baseline performs well at coarser scales but ultimately produces structureless images when zoomed in due to its limited upsampling capacity.

## 5.2. Additional Results

In Fig. 6 we show fractal inversion results of real images. Fig. 7 presents our reconstruction of the famous Sierpinski triangle [Sie15] using two different values for the number of functions  $N$ . The ground-truth fractal uses  $N = 3$ . Our reconstructions closely match the ground truth, with only minor shifts, regardless of the  $N$  value, provided it remains conservative. This happens because, when more functions are provided than needed, our pipeline converges to so-

lutions where the extra functions are mapped to zero, resulting in negligible sampling probabilities.

## 5.3. Ablations

Here, we analyze individual components of our method through ablation studies, examining variations in the model, objective function, and optimization routine. The results of the quantitative evaluations, based on 50 fractals, are presented in Tab. 4, with corresponding qualitative examples shown in Fig. 8.

**Model** We first explore an alternative to our SVD-based parameterization, where the affine functions  $f_i$  are parameterized directly by their matrix and vector entries (Naïve Parameterization). We found this approach to be prone to generating non-contractive functions, requiring us to significantly increase the regularization weight  $\lambda_{\text{reg}}$  to obtain meaningful results. Second, we omit our multisampling stage to investigate the importance of anti-aliasing in our framework (w/o Multisampling). We observe that both components are important for high-quality inversions.

**Objective function** We omit individual components of our objective function to better understand their contributions to the final result. Specifically, we omit the MIP map in  $\mathcal{L}_{\text{MSE}}$ , as well as the terms  $\mathcal{L}_{\text{SSIM}}$ ,  $\mathcal{L}_{\text{LPIPS}}$ , and  $\mathcal{L}_{\text{reg}}$  individually. Additionally, we replace our entire objective function with a classical alternative that minimizes the distance between moments [VR89, RZ94]. Each term in our objective function plays a significant role in achieving our results, whereas the moments-based approach performs noticeably worse.

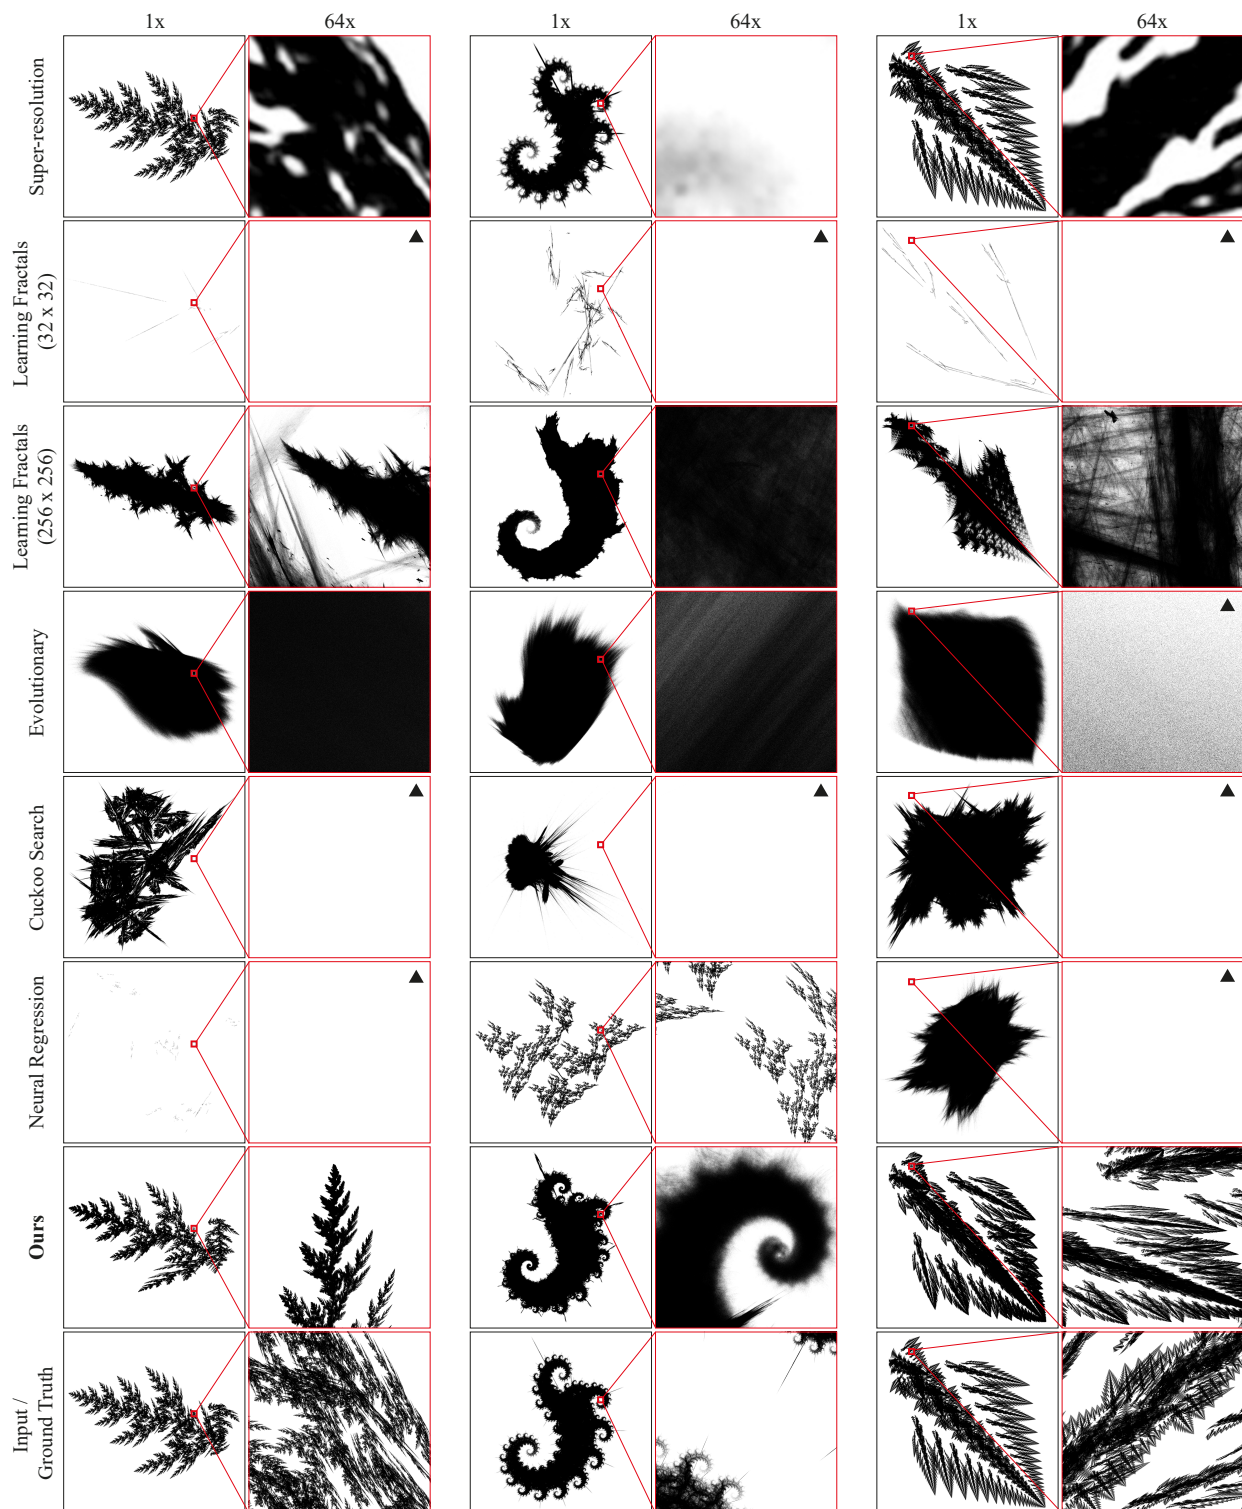
**Optimizer** To analyze our hybrid optimizer design, we evaluate our method using pure gradient-based optimization (w/o Simulated Annealing) and pure simulated annealing (w/o Gradients). Additionally, we explore a gradient-based approach in which we introduce carefully tuned Gaussian noise to the gradients in each optimization iteration to help escape local minima (Noisy Gradients). Simulated annealing clearly aids in capturing the global shape, whereas a solution without gradients fails to converge to a meaningful result. Adding noise to the gradients provides only marginal improvement over a pure gradient-based solution.

## 5.4. Limitations and Discussion

While our analysis shows that our approach marks the new state of the art in fractal inversion, it is important to acknowledge some limitations.

Our results do not perfectly align with the ground truth; rather, the generated details only *statistically* resemble the expected patterns. A single global solution exists and the extremely complex optimization landscape—characterized by numerous local minima and plateaus—makes perfect reconstruction of the ground truth challenging. At present, our hybrid optimization strategy seems best suited to effectively navigate this intricate space regardless of the initialization.

Our system, constrained by its analytical approach, *always* results in perfectly self-similar structures. However, the self-similarity in natural images typically extends over only a few scales and is generally approximate due to the repeated influence of various physical

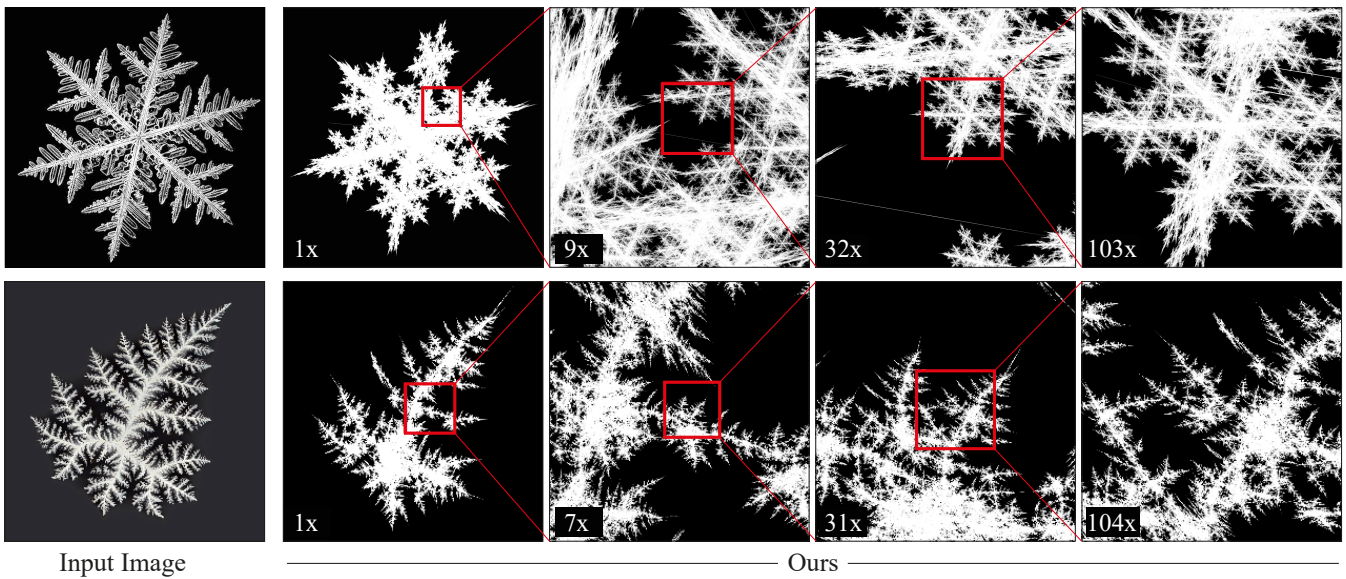


**Figure 5:** Qualitative fractal inversion results for various methods (rows). For each instance (columns), the full fractal is shown on the left, with a 64x zoomed-in view on the right. The bottom row presents the input image  $I_{ref}$  used for all methods (left) next to a zoomed-in view of the ground-truth fractal (right). Note that the zoomed-in views are consistent across rows. Images marked with a ▲ denote views that do not contain the full point count (see Sec. 5). More visuals can be found in our supplemental materials. For continuous zoom-in visuals, please refer to our supplemental video.


**Table 3:** Quantitative evaluation of our method compared to previous work. We report shape-based (F1, IoU) and image-based (PSNR, SSIM, LPIPS) quality metrics alongside the time it takes to run each method on a single GPU. We include five fractal-based baselines and a pixel-based super-resolution method. Unlike all other approaches, including ours, the super-resolution baseline is restricted to a maximum zoom-in capability of 8x.

Method	F1↑	IoU↑	PSNR↑	SSIM↑	LPIPS↓	Time
Super-resolution [LCS*21]*	.917	.854	16.30	.714	.239	12 sec
Learning Fractals (32 × 32) [TCCC23]	.121	.080	6.59	.436	.656	8 min
Learning Fractals (256 × 256) [TCCC23]	.553	.430	6.78	.348	.648	35 min
Evolutionary [NG94]	.621	.492	6.49	.289	.745	180 min
Cuckoo Search [QIG17]	.490	.351	3.43	.140	.825	540 min
Neural Regression [GD21]*	.161	.106	5.33	.358	.709	16 ms
<b>Ours</b>	.691	.572	8.91	.449	.428	55 min

\* Generalizable method that requires pre-training.



**Figure 6:** Inversion results for real images. Our representation enables infinite zoom-ins, continuously generating self-similar details at any scale while capturing intricate structural patterns. In contrast, natural structures exhibit self-similarity only within a limited scale range.



**Figure 7:** The Sierpinski triangle (left) alongside our reconstructions using different function counts  $N$  (center and right). Our reconstructions closely match the ground truth, with only minor misalignments, visible in the insets, independent of  $N$ .

forces acting on structures [Man82]. For instance, the fern in Fig. 2 exhibits about three scales of approximate self-similarity, while the formation of snow crystals, a subject of study for decades [Lib05], continues to present challenges. In such cases, our method may fail to find functions that represent the observed approximate self-similarity and hence inaccurately captures the ground truth (Fig. 6). Using fractals, this discrepancy appears to be the trade-off for achieving arbitrary-scale image generation, as evidenced by our significant zoom-ins.

For future work, we envision a fractal system that incorporates more expressive functions as building blocks [CLV\*95] – potentially parameterized by neural networks or enhanced by the injection of physical forces. Such an approach could enable better representation and simulation of complex self-similar geometries, while also achieving generative capabilities across scales.

A minor technical limitation of our approach is the fixed number of functions used in our model. A potential avenue for future work is the explicit optimization of sampling probabilities to remove this restriction, which likely requires an appropriate parameterization [JGP17].

## 6. Conclusion

We proposed a novel approach to solving the fractal inverse problem, which includes a differentiable forward model that implements the chaos game based on Iterated Function Systems and a point splatting module that renders the generated fractal point set into an image. Notably, we demonstrated that combining large-scale fractal generation with high-quality rendering, along with a hybrid of gradient-based and stochastic optimization, leads to state-of-the-art results in inverting image fractals. This was validated through extensive comparisons with a diverse range of classical and modern methods.

We envision fractals as optimizable graphics primitives that, due to their point-based nature, can seamlessly integrate with recent advancements in Gaussian scene representations. Fractals can serve as a fundamental tool for representing content across a wide range of scales – capabilities that are challenging to achieve with classical (e.g., meshes) or neural representations (e.g., neural radiance fields [MST\*20]). To fully realize this potential, future research must extend robust fractal optimization frameworks to operate in 3D [Nor82, SK24] and incorporate appearance [BN24]. These advancements will unlock new opportunities for fractals to contribute to the evolving landscape of scene representations and image synthesis.

## Acknowledgement

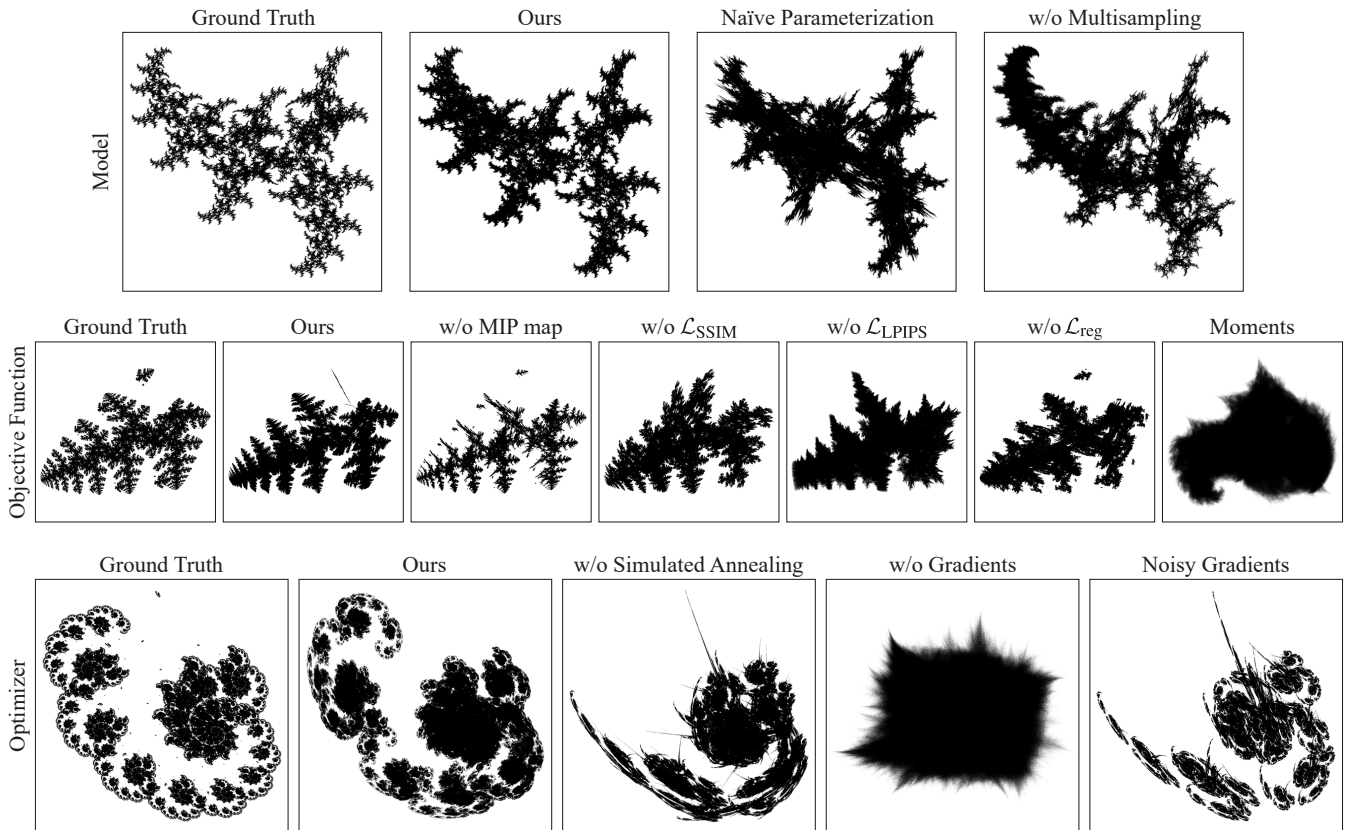
Open Access funding enabled and organized by Projekt DEAL.

## References

- [AF22] ANDERSON C., FARRELL R.: Improving fractal pre-training. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* (2022).
- [Ban91] BANDT C.: Self-similar sets 5. Integer matrices and fractal tilings of  $\mathbb{R}^n$ . *Proceedings of the American Mathematical Society* 112, 2 (1991), 549–562.
- [Bar88] BARNSLEY M.: *Fractals Everywhere*. Academic Press, 1988.
- [BD85] BARNSLEY M. F., DEMKO S.: Iterated function systems and the global construction of fractals. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 399, 1817 (1985), 243–275.
- [BDM18] BROSSE N., DURMUS A., MOULINES E.: The promises and pitfalls of stochastic gradient langevin dynamics. In *Advances in Neural Information Processing Systems (NeurIPS)* (2018).
- [BdR17] BLOEM P., DE ROOIJ S.: An expectation-maximization algorithm for the fractal inverse problem. *arXiv preprint arXiv:1706.03149* (2017).
- [BEHL86] BARNSLEY M. F., ERVIN V., HARDIN D., LANCASTER J.: Solution of an inverse problem for fractals and other sets. *Proceedings of the National Academy of Sciences* 83, 7 (1986), 1975–1977.
- [Ber97] BERKNER K.: A wavelet-based solution to the inverse problem for fractal interpolation functions. In *Fractals in Engineering: From Theory to Industrial Applications*. Springer, 1997, pp. 81–92.
- [BGF\*23] BELHE Y., GHARBI M., FISHER M., GEORGIEV I., RAMAMOORTHY R., LI T.-M.: Discontinuity-aware 2d neural fields. *ACM Transactions on Graphics* 42, 6 (2023).
- [BH93] BARNSLEY M. F., HURD L. P.: *Fractal Image Compression*. Peters A. K., 1993.
- [BJM\*88] BARNSLEY M. F., JACQUIN A., MALASSENET F., REUTER L., SLOAN A. D.: Harnessing chaos for image synthesis. In *ACM SIGGRAPH Conference* (1988).
- [BKK19] BAI S., KOLTER J. Z., KOLTUN V.: Deep equilibrium models. *Advances in Neural Information Processing Systems (NeurIPS)* (2019).
- [BKSI19] BELL-KLIGLER S., SHOCHER A., IRANI M.: Blind super-resolution kernel estimation using an internal-GAN. In *Advances in Neural Information Processing Systems (NeurIPS)* (2019).
- [BN24] BANNISTER J. J., NOWROUZEZHRAI D.: Learnable fractal flames. *arXiv preprint arXiv:2406.09328* (2024).
- [Bur81] BURT P. J.: Fast filter transform for image processing. *Computer Graphics and Image Processing* 16, 1 (1981), 20–51.
- [Car80] CARPENTER L. C.: Computer rendering of fractal curves and surfaces. In *ACM SIGGRAPH Conference* (1980).
- [CBG\*17] CISSE M., BOJANOWSKI P., GRAVE E., DAUPHIN Y., USUNIER N.: Parseval networks: Improving robustness to adversarial examples. In *International Conference on Machine Learning (ICML)* (2017).
- [CLRS00] COLLET P., LUTTON E., RAYNAL F., SCHOENAUER M.: Polar IFS + Parisian genetic programming=Efficient IFS inverse problem solving. *Genetic Programming and Evolvable Machines* 1, 4 (2000), 339–361.
- [CLV\*95] CRETIN G., LUTTON E., VÉHEL J. L., GLEVAREC P., ROLL C.: Mixed ifs: Resolution of the inverse problem using genetic programming. In *European Conference on Artificial Evolution (AE)* (1995).
- [CLW21] CHEN Y., LIU S., WANG X.: Learning continuous image representation with local implicit image function. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021).
- [CZHN22] CHEN B., ZHI T., HEBERT M., NARASIMHAN S. G.: Learning continuous implicit representation for near-periodic patterns. In *European Conference on Computer Vision (ECCV)* (2022).
- [DHN85] DEMKO S., HODGES L., NAYLOR B.: Construction of fractal objects with iterated function systems. In *ACM SIGGRAPH Conference* (1985).
- [dSNLV21] DA SILVA V., NOVELLO T., LOPES H., VELHO L.: Real-time rendering of complex fractals. In *Ray Tracing Gems II: Next Generation Real-Time Rendering with DXR, Vulkan, and OptiX*. Apress, 2021, pp. 529–544.
- [Ebe02] EBERT D.: *Texturing & Modeling: A Procedural Approach*. Morgan Kaufmann, 2002.
- [Elt87] ELTON J. H.: An ergodic theorem for iterated maps. *Ergodic Theory and Dynamical Systems* 7, 4 (1987), 481–488.
- [Fat01] FATHAUER R. W.: Fractal tilings based on kite-and dart-shaped prototiles. *Computers & Graphics* 25, 2 (2001), 323–331.
- [Fis94] FISHER Y.: Fractal image compression. *Fractals* 2, 3 (1994), 347–361.
- [FSWK20] FATHONY R., SAHU A. K., WILLMOTT D., KOLTER J. Z.: Multiplicative filter networks. In *International Conference on Learning Representations (ICLR)* (2020).
- [FV94] FORTE B., VCSCAY E. R.: Solving the inverse problem for function/image approximation using iterated function systems. *Fractals* 2, 3 (1994), 325–334.
- [GBI09] GLASNER D., BAGON S., IRANI M.: Super-resolution from a single image. In *IEEE International Conference on Computer Vision (ICCV)* (2009).

**Table 4:** Ablations.

Component	Ablation	F1 $\uparrow$	IoU $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Model	Naïve Parameterization	.583	.467	9.12	.481	.431
	w/o Multisampling	.561	.445	9.16	.481	.441
Objective Fct.	w/o MIP map	.363	.286	9.08	.525	.512
	w/o $\mathcal{L}_{SSIM}$	.611	.483	8.41	.436	.454
	w/o $\mathcal{L}_{LPIPS}$	.544	.440	8.86	.466	.511
	w/o $\mathcal{L}_{reg}$	.596	.474	8.66	.449	.476
	Moments	.542	.408	4.17	.170	.818
Optimizer	w/o Simulated Annealing	.594	.468	8.72	.459	.441
	w/o Gradients	.496	.377	7.14	.296	.736
	Noisy Gradients	.578	.455	8.93	.471	.442
<b>Ours</b>		.611	.488	9.23	.482	.415

**Figure 8:** A qualitative overview of our ablation studies. We analyze different variations of our method in terms of the model (first row), the objective function (second row), and the optimizer (last row).

- [GCI00] GUTIÉRREZ J. M., COFIÑO A. S., IVANISSEVICH M. L.: An hybrid evolutive-genetic strategy for the inverse fractal problem of IFS models. In *Ibero-American Conference on Artificial Intelligence (IB-ERAMIA)* (2000).
- [GD98] GROSSMAN J. P., DALLY W. J.: Point sample rendering. In *Eurographics Workshop on Rendering Techniques* (1998).
- [GD21] GRAHAM L., DEMERS M.: Applying neural networks to a fractal inverse problem. In *International Conference on Applied Mathematics, Modeling and Computational Science (AMMCS)* (2021).
- [GHYZ20] GUO Y., HAŠAN M., YAN L., ZHAO S.: A bayesian inference framework for procedural material parameter estimation. *Computer Graphics Forum* 39, 7 (2020), 255–266.
- [GI18] GÁLVEZ A., IGLESIAS A.: Modified memetic self-adaptive firefly algorithm for 2D fractal image reconstruction. In *IEEE International Computer Software and Applications Conference (COMPSAC)* (2018).
- [GID\*21] GÁLVEZ A., IGLESIAS A., DÍAZ J. A., FISTER I., LÓPEZ J., FISTER JR I.: Modified OFS-RDS bat algorithm for ifs encoding of bitmap fractal binary images. *Advanced Engineering Informatics* 47 (2021).
- [GP11] GROSS M., PFISTER H.: *Point-Based Graphics*. Elsevier, 2011.
- [GSV\*14] GILET G., SAUVAGE B., VANHOEY K., DISCHLER J.-M., GHAZANFARPOUR D.: Local random-phase noise for procedural texturing. *ACM Transactions on Graphics* 33, 6 (2014).
- [HBG\*95] HAVLIN S., BULDYREV S., GOLDBERGER A., MANTEGNA R., OSSADNIK S., PENG C.-K., SIMONS M., STANLEY H.: Fractals in biology and medicine. *Chaos, Solitons & Fractals* 6 (1995), 171–201.
- [HD91] HART J. C., DEFANTI T. A.: Efficient antialiased rendering of 3-D linear fractals. In *ACM SIGGRAPH Conference* (1991).
- [HDR19] HU Y., DORSEY J., RUSHMEIER H.: A novel framework for inverse procedural texture modeling. *ACM Transactions on Graphics* 38, 6 (2019).
- [HHD\*22] HU Y., HE C., DESCHAIANTRE V., DORSEY J., RUSHMEIER H.: An inverse procedural modeling pipeline for svbrdf maps. *ACM Transactions on Graphics* 41, 2 (2022).
- [HMK\*19] HALLADJIAN S., MIAO H., KOUŘIL D., GRÖLLER M. E., VIOLA I., ISENBERG T.: Scale trotter: Illustrative visual travels across negative scales. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 654–664.
- [HN18] HEITZ E., NEYRET F.: High-performance by-example noise using a histogram-preserving blending operator. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 1, 2 (2018).
- [HSK89] HART J. C., SANDIN D. J., KAUFFMAN L. H.: Ray tracing deterministic 3-D fractals. In *ACM SIGGRAPH Conference* (1989).
- [Hut81] HUTCHINSON J. E.: Fractals and self similarity. *Indiana University Mathematics Journal* 30, 5 (1981), 713–747.
- [Iij59] IJIMA T.: Basic theory of pattern observation. *Papers of the Technical Group on Automata and Automatic Control* (1959).
- [Jac92] JACQUIN A.: Image coding based on a fractal theory of iterated contractive image transformations. *IEEE Transactions on Image Processing* 1, 1 (1992), 18–30.
- [JGP17] JANG E., GU S., POOLE B.: Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations (ICLR)* (2017).
- [Jul18] JULIA G.: Mémoire sur l'itération des fonctions rationnelles. *Journal de mathématiques pures et appliquées* 1 (1918), 47–245.
- [KB15] KINGMA D. P., BA J.: Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)* (2015).
- [Kes82] KESHNER M. S.:  $1/f$  noise. *Proceedings of the IEEE* 70, 3 (1982), 212–218.
- [KGJV83] KIRKPATRICK S., GELATT JR C. D., VECCHI M. P.: Optimization by simulated annealing. *Science* 220, 4598 (1983), 671–680.
- [KHE\*10] KLASHED S., HEMINGSSON P., EMMART C., COOPER M., YNNERMAN A.: Uniview - visualizing the universe. In *Eurographics Areas Papers* (2010).
- [Kim15] KIM T.: Quaternion Julia set shape optimization. *Computer Graphics Forum* 34, 5 (2015), 167–176.
- [KKLD23] KERBL B., KOPANAS G., LEIMKÜHLER T., DRETTAKIS G.: 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* 42, 4 (2023).
- [KOM\*20] KATAOKA H., OKAYASU K., MATSUMOTO A., YAMAGATA E., YAMADA R., INOUE N., NAKAMURA A., SATOH Y.: Pre-training without natural images. In *Asian Conference on Computer Vision (ACCV)* (2020).
- [KRS\*24] KHERADMAND S., REBAIN D., SHARMA G., SUN W., TSENG J., ISACK H., KAR A., TAGLIASACCHI A., YI K. M.: 3D Gaussian splatting as Markov chain Monte Carlo. In *Advances in Neural Information Processing Systems (NeurIPS)* (2024).
- [KRWM22] KARNEWAR A., RITSCHER T., WANG O., MITRA N.: 3in-GAN: Learning a 3D generative model from images of a self-similar scene. In *International Conference on 3D Vision (3DV)* (2022).
- [Lan95] LANKHORST M. M.: *Iterated function systems optimization with genetic algorithms*. Tech. Rep. CS-R 9501, University of Groningen, Department of Computing Science, 1995.
- [LCS\*21] LIANG J., CAO J., SUN G., ZHANG K., VAN GOOL L., TIMOFTE R.: Swinir: Image restoration using swin transformer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021).
- [LFS21] LICORISH C., FARAJ N., SUMMA B.: Adaptive compositing and navigation of variable resolution images. *Computer Graphics Forum* 40, 1 (2021), 138–150.
- [LHOK\*10] LIU Y., HEL-OR H., KAPLAN C. S., VAN GOOL L., ET AL.: Computational symmetry in computer vision and computer graphics. *Foundations and Trends in Computer Graphics and Vision* 5, 1–2 (2010), 1–195.
- [Li19] LI T.-M.: Differentiable visual computing. *arXiv preprint arXiv:1904.12228* (2019).
- [Lib05] LIBBRECHT K. G.: The physics of snow crystals. *Reports on progress in physics* 68, 4 (2005), 855.
- [Lin68] LINDENMAYER A.: Mathematical models for cellular interactions in development I. Filaments with one-sided inputs. *Journal of Theoretical Biology* 18, 3 (1968), 280–299.
- [LLC\*21] LIU Z., LIN Y., CAO Y., HU H., WEI Y., ZHANG Z., LIN S., GUO B.: Swin transformer: Hierarchical vision transformer using shifted windows. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021).
- [LP00] LEFEBVRE L., POULIN P.: Analysis and synthesis of structural textures. In *Graphics Interface Conference (GI)* (2000).
- [LSS\*17] LUKÁČ M., ŠYKORA D., SUNKAVALLI K., SHECHTMAN E., JAMRIŠKA O., CARR N., PAJDLA T.: Nautilus: Recovering regional symmetry transformations for image editing. *ACM Transactions on Graphics* 36, 4 (2017).
- [LVL93] LEVY-VEHEL J., LUTTON E.: Optimization of fractal functions using genetic algorithms. In *IFIP Working Conference on Fractals in the Natural and Applied Sciences* (1993).
- [LVVPW22] LINDELL D. B., VAN VEEN D., PARK J. J., WETZSTEIN G.: BACON: Band-limited coordinate networks for multiscale scene representation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022).
- [LYMF19] LI J., YUAN Y., MEI K., FANG F.: Lightweight and accurate recursive fractal network for image super-resolution. In *IEEE/CVF International Conference on Computer Vision (ICCV)* (2019).
- [LZ21] LASSNER C., ZOLLHOFER M.: Pulsar: Efficient sphere-based neural rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021).

- [MAAB\*17] MOHAMMED H., AL-AWAMI A. K., BEYER J., CALI C., MAGISTRETTI P., PFISTER H., HADWIGER M.: Abstractocyte: A visual tool for exploring nanoscale astroglial cells. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 853–861.
- [Man80] MANDELBROT B. B.: Fractal aspects of the iteration of  $z \rightarrow \lambda z(1-z)$  for complex  $\lambda$  and  $z$ . *Annals of the New York Academy of Sciences* 357, 1 (1980), 249–259.
- [Man82] MANDELBROT B. B.: *The Fractal Geometry of Nature*. W. H. Freeman and Co., 1982.
- [Man97] MANDELBROT B. B.: *Fractals and Scaling in Finance*. Springer Science & Business Media, 1997.
- [Mar10] MARTYN T.: Realistic rendering 3D IFS fractals in real-time with graphics accelerators. *Computers & Graphics* 34, 2 (2010), 167–175.
- [MEG04] MANDELBROT B. B., EVERTSZ C. J., GUTZWILLER M. C.: *Fractals and Chaos: The Mandelbrot Set and Beyond*. Springer, 2004.
- [Mer23] MERRELL P.: Example-based procedural modeling using graph grammars. *ACM Transactions on Graphics* 42, 4 (2023).
- [MESK22] MÜLLER T., EVANS A., SCHIED C., KELLER A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics* 41, 4 (2022).
- [MI14] MICHAELI T., IRANI M.: Blind deblurring using internal patch recurrence. In *European Conference on Computer Vision (ECCV)* (2014).
- [MNT\*24] MUJKANOVIC F., NSAMPI N. E., THEOBALT C., SEIDEL H.-P., LEIMKÜHLER T.: Neural gaussian scale-space fields. *ACM Transactions on Graphics* 43, 4 (2024).
- [MP98] MELNIK O., POLLACK J.: A gradient descent method for a neural fractal memory. In *IEEE International Joint Conference on Neural Networks (IJCNN)* (1998).
- [MRF\*23] MOSER B. B., RAUE F., FROLOV S., PALACIO S., HEES J., DENGEL A.: Hitchhiker’s guide to super-resolution: Introduction and recent advances. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 8 (2023), 9862–9882.
- [MST\*20] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHY R., NG R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision (ECCV)* (2020).
- [MVN68] MANDELBROT B. B., VAN NESS J. W.: Fractional brownian motions, fractional noises and applications. *SIAM Review* 10, 4 (1968), 422–437.
- [NG94] NETTLETON D. J., GARIGLIANO R.: Evolutionary algorithms and a fractal inverse problem. *Biosystems* 33, 3 (1994), 221–231.
- [Nor82] NORTON A.: Generation and display of geometric fractals in 3-D. *ACM SIGGRAPH Computer Graphics* 16, 3 (1982), 61–67.
- [OCNG21] OUYANG P., CHUNG K. W., NICOLAS A., GDAWIEC K.: Self-similar fractal drawings inspired by MC Escher’s print square limit. *ACM Transactions on Graphics* 40, 3 (2021).
- [PD84] PORTER T., DUFF T.: Compositing digital images. In *ACM SIGGRAPH Conference* (1984).
- [Pen84] PENTLAND A. P.: Fractal-based description of natural scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6, 6 (1984), 661–674.
- [PGC\*17] PASZKE A., GROSS S., CHINTALA S., CHANAN G., YANG E., DEVITO Z., LIN Z., DESMAISON A., ANTIGA L., LERER A.: Automatic differentiation in pytorch. In *NeurIPS Workshop on Autodiff* (2017).
- [PJSF04] PEITGEN H.-O., JÜRGENS H., SAUPE D., FEIGENBAUM M. J.: *Chaos and Fractals: New Frontiers of Science*. Springer, 2004.
- [PL12] PRUSINKIEWICZ P., LINDENMAYER A.: *The algorithmic beauty of plants*. Springer Science & Business Media, 2012.
- [PNS\*22] PAZ H., NOVELLO T., SILVA V., SCHARDONG G., SCHIRMER L., CHAGAS F., LOPES H., VELHO L.: Multiresolution neural networks for imaging. In *Conference on Graphics, Patterns and Images (SIBGRAPI)* (2022).
- [PXM\*22] POLI M., XU W., MASSAROLI S., MENG C., KIM K., ERMON S.: Self-similarity priors: Neural collages as differentiable fractal representations. *Advances in Neural Information Processing Systems (NeurIPS)* (2022).
- [QIG17] QUIRCE J., IGLESIAS A., GÁLVEZ A.: Cuckoo search algorithm approach for the IFS inverse problem of 2D binary fractal images. In *Advances in Swarm Intelligence (ICSI)* (2017).
- [RHD08] REDIES C., HASENSTEIN J., DENZLER J.: Fractal-like image statistics in visual art: similarity to natural scenes. *Spatial Vision* 21, 1–2 (2008), 137–148.
- [RS80] ROZENBERG G., SALOMAA A.: *The Mathematical Theory of L Systems*. Academic Press, 1980.
- [RZ94] RINALDO R., ZAKHOR A.: Inverse and approximation problem for two-dimensional fractal sets. *IEEE Transactions on Image Processing* 3, 6 (1994), 802–820.
- [SB06] SARAFPOPOULOS A., BUXTON B.: Resolution of the inverse problem for iterated function systems using evolutionary algorithms. In *IEEE International Conference on Evolutionary Computation* (2006).
- [SBI19] SHOCHER A., BAGON S., ISOLA P., IRANI M.: InGAN: Capturing and retargeting the “DNA” of a natural image. In *IEEE/CVF International Conference on Computer Vision (ICCV)* (2019).
- [SCI18] SHOCHER A., COHEN N., IRANI M.: “Zero-shot” super-resolution using deep internal learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2018).
- [Sco24] SCOTT C. B.: Differentiable iterated function systems. In *ICML Workshop on Differentiable Almost Everything: Differentiable Relaxations, Algorithms, Operators, and Simulators* (2024).
- [SDG\*93] STRUZIK Z. R., DOOJES E., GROEN F., ET AL.: The solution of the inverse fractal problem with the help of wavelet decomposition. In *Fractals in Natural Sciences* (1993).
- [SDM19] SHAHAM T. R., DEKEL T., MICHAELI T.: SinGAN: Learning a generative model from a single natural image. In *IEEE/CVF International Conference on Computer Vision (ICCV)* (2019).
- [SHM05] SONG C., HAVLIN S., MAKSE H. A.: Self-similarity of complex networks. *Nature* 433, 7024 (2005), 392–395.
- [SI07] SHECHTMAN E., IRANI M.: Matching local self-similarities across images and videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2007).
- [Sie15] SIERPINSKI W.: Sur une courbe dont tout point est un point de ramification. *CR Acad. Sci.* 160 (1915), 302–305.
- [SK24] SCHOR A., KIM T.: Into the portal: Directable fractal self-similarity. In *ACM SIGGRAPH Conference* (2024).
- [SLH\*20] SHI L., LI B., HAŞAN M., SUNKAVALLI K., BOUBEKEUR T., MECH R., MATUSIK W.: Match: Differentiable material graphs for procedural material capture. *ACM Transactions on Graphics* 39, 6 (2020).
- [SRDT01] SHLYAKHTER I., ROZENOER M., DORSEY J., TELLER S.: Reconstructing 3D tree models from instrumented photographs. *IEEE Computer Graphics and Applications* 21, 3 (2001), 53–61.
- [Sta91] STARK J.: Iterated function systems as neural networks. *Neural Networks* 4, 5 (1991), 679–690.
- [STB\*22] SARAGADAM V., TAN J., BALAKRISHNAN G., BARANIUK R. G., VEERARAGHAVAN A.: MINER: Multiscale implicit neural representation. In *European Conference on Computer Vision (ECCV)* (2022).
- [SZR\*23] SPIELBERG A., ZHONG F., REMATAS K., JATAVALLABHULA K. M., OZTIRELI C., LI T.-M., NOWROUZEZAHRAI D.: Differentiable visual computing for inverse problems and machine learning. *Nature Machine Intelligence* 5, 11 (2023), 1189–1199.
- [TCCC23] TU C.-H., CHEN H.-Y., CARLYN D., CHAO W.-L.: Learning fractals by gradient descent. In *AAAI Conference on Artificial Intelligence* (2023).

- [TGSC24] TZATHAS P., GAILLETON B., STEER P., CORDONNIER G.: Physically-based analytical erosion for fast terrain generation. *Computer Graphics Forum* 43, 2 (2024).
- [TLW\*19] TAO W., LIU X., WANG Y., BATTLE L., DEMIRALP Ç., CHANG R., STONEBRAKER M.: Kyrix: Interactive pan/zoom visualizations at scale. *Computer Graphics Forum* 38, 3 (2019), 529–540.
- [TPRC00] TURIEL A., PARGA N., RUDERMAN D. L., CRONIN T. W.: Multiscaling and information content of natural color images. *Physical Review E* 62, 1 (2000), 1138–1148.
- [Tur97] TURCOTTE D. L.: *Fractals and Chaos in Geology and Geophysics*. Cambridge University Press, 1997.
- [VR89] VRSCAY E. R., ROHRIG C. J.: Iterated function systems and the inverse problem of fractal construction using moments. In *Conference on Computers and Mathematics* (1989).
- [Vrs91] VRSCAY E. R.: Iterated function systems: theory, applications and the inverse problem. In *Fractal Geometry and Analysis*. Springer, 1991, pp. 405–468.
- [VS99] VRSCAY E. R., SAUPE D.: Can one break the “collage barrier” in fractal image coding? In *Fractals: Theory and Applications in Engineering*. Springer, 1999, pp. 307–323.
- [VSLD13] VANHOEY K., SAUVAGE B., LARUE F., DISCHLER J.-M.: On-the-fly multi-scale infinite texturing from example. *ACM Transactions on Graphics* 32, 6 (2013).
- [WBSS04] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612.
- [WDJ\*24] WOLSKI K., DJEACOMMAR A., JAVANMARDI A., SEIDEL H.-P., THEOBALT C., CORDONNIER G., MYSZKOWSKI K., DRETTAKIS G., PAN X., LEIMKÜHLER T.: Learning images across scales using adversarial trainings. *ACM Transactions on Graphics* 43, 4 (2024).
- [WGSJ20] WILES O., GKIOXARI G., SZELISKI R., JOHNSON J.: Synsin: End-to-end view synthesis from a single image. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020).
- [Wil83] WILLIAMS L.: Pyramidal parametrics. In *ACM SIGGRAPH Conference* (1983).
- [Wit87] WITKIN A. P.: Scale-space filtering. In *Readings in Computer Vision*. Elsevier, 1987, pp. 329–332.
- [WWSR03] WONKA P., WIMMER M., SILLION F., RIBARSKY W.: Instant architecture. *ACM Transactions on Graphics* 22, 3 (2003), 669–677.
- [YSW\*19] YIFAN W., SERENA F., WU S., ÖZTIRELI C., SORKINE-HORNUNG O.: Differentiable surface splatting for point-based geometry processing. *ACM Transactions on Graphics* 38, 6 (2019).
- [ZII11] ZONTAK M., IRANI M.: Internal statistics of a single natural image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2011).
- [ZIE\*18] ZHANG R., ISOLA P., EFROS A. A., SHECHTMAN E., WANG O.: The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2018).
- [ZPVBG01] ZWICKER M., PFISTER H., VAN BAAR J., GROSS M.: Surface splatting. In *ACM SIGGRAPH Conference* (2001), pp. 371–378.
- [ZWSW23] ZHANG Y., WU S., SNAVELY N., WU J.: Seeing a rose in five thousand ways. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2023).
- [ZZB\*18] ZHOU Y., ZHU Z., BAI X., LISCHINSKI D., COHEN-OR D., HUANG H.: Non-stationary texture synthesis by adversarial expansion. *ACM Transactions on Graphics* 37, 4 (2018).

## Appendix A: Pseudocode

We provide pseudocode of our optimization algorithm in Alg. 1.

---

### Algorithm 1 Fractal Code Optimization

$I_{\text{ref}}$  : Input Image  
 $b$  : Batch Size of the Fractal Point Generator  
 $\theta$  : Optimizable Fractal Parameters

---

```

 $\theta \leftarrow \text{InitializeParameters}()$  ▷ Uniform Random
 $i \leftarrow 0$  ▷ Iteration Count
while not converged do
   $\tau \leftarrow \text{LinearAnneal}(i)$  ▷ Temperature
   $\mathcal{P} \leftarrow \text{GENERATEFRACTALPOINTS}(\theta, b)$ 
   $I \leftarrow \text{RENDER}(\mathcal{P})$ 
   $\mathcal{L} \leftarrow \text{Loss}(I, I_{\text{ref}}, \theta)$  ▷ Loss : Eq. 4
   $\theta \leftarrow \text{ADAM}(\nabla_{\theta} \mathcal{L})$  ▷ Backprop & Step
  if DoHybridOptimization( $i$ ) then ▷ Hybrid Optimization
     $\theta_{\text{curr}} \leftarrow \theta$ 
     $\mathcal{L}_{\text{curr}} \leftarrow \text{MSE}(I, I_{\text{ref}})$  ▷ Current Energy
    for  $j \leftarrow 0$  to 10 do
       $\Delta\theta_{\text{cand}} \leftarrow \mathcal{N}(\mathbf{0}, \tau/\delta)$ 
       $\theta_{\text{cand}} \leftarrow \theta_{\text{curr}} + \Delta\theta_{\text{cand}}$  ▷ Sample Neighbour
       $\mathcal{P} \leftarrow \text{GENERATEFRACTALPOINTS}(\theta, b)$ 
       $I \leftarrow \text{RENDER}(\mathcal{P})$ 
       $\mathcal{L}_{\text{cand}} \leftarrow \text{MSE}(I, I_{\text{ref}})$  ▷ Candidate Energy
       $p(\theta_{\text{cand}}) \leftarrow \text{AcceptanceCriteria}(\mathcal{L}_{\text{cand}}, \mathcal{L}_{\text{curr}}, \tau)$  ▷ Eq. 7
      if  $\mathcal{L}_{\text{cand}} < \mathcal{L}_{\text{curr}}$  or  $\text{Random}() < p(\theta_{\text{cand}})$  then
         $\theta_{\text{curr}} \leftarrow \theta_{\text{cand}}$ 
         $\mathcal{L}_{\text{curr}} \leftarrow \mathcal{L}_{\text{cand}}$ 
      end if
    end for
     $\theta \leftarrow \theta_{\text{curr}}$ 
    Reset ADAM momentum
  end if
end while

```

---

```

function GENERATEFRACTALPOINTS( $\theta, b$ ) ▷ Sec. 4.1
   $\mathbf{x}_0 \leftarrow \text{SampleRandomPoints}(b)$ 
   $G \leftarrow \text{ConstructFractalPointGenerator}(\theta, b)$ 
   $\mathcal{P}' \leftarrow G(\mathbf{x}_0)$ 
   $\mathcal{P} \leftarrow \text{RemoveInitPoints}()$  ▷ Warm-Up (Sec. 4.4)
  return  $\mathcal{P}$ 
end function

```

---

```

function RENDER( $\mathcal{P}$ ) ▷ Sec. 4.2
   $I' \leftarrow \mathcal{R}(\mathcal{P})$  ▷ Differentiable Rasterization
   $I \leftarrow \text{SuperSample}(I')$  ▷ Supersampling
  return  $I$ 
end function

```

---