

## Visual Diffusion Models are Geometric Solvers

Nir Goren<sup>1,\*</sup> Shai Yehezkel<sup>1,\*</sup> Omer Dahary<sup>1</sup> Andrey Voynov<sup>2</sup> Or Patashnik<sup>1</sup> Daniel Cohen-Or<sup>1</sup>

<sup>1</sup>Tel Aviv University <sup>2</sup>Google DeepMind

\*Equal contribution

### Abstract

*In this paper we show that visual diffusion models can serve as effective geometric solvers: they can directly reason about geometric problems by working in pixel space. We first demonstrate this on the Inscribed Square Problem, a long-standing problem in geometry that asks whether every Jordan curve contains four points forming a square. We then extend the approach to two other well-known hard geometric problems: the Steiner Tree Problem and the Simple Polygon Problem. Our method treats each problem instance as an image and trains a standard visual diffusion model that transforms Gaussian noise into an image representing a valid approximate solution that closely matches the exact one. The model learns to transform noisy geometric structures into correct configurations, effectively recasting geometric reasoning as image generation. Unlike prior work that necessitates specialized architectures and domain-specific adaptations when applying diffusion to parametric geometric representations, we employ a standard visual diffusion model that operates on the visual representation of the problem. This simplicity highlights a surprising bridge between generative modeling and geometric problem solving. Beyond the specific problems studied here, our results point toward a broader paradigm: operating in image space provides a general and practical framework for approximating notoriously hard problems, and opens the door to tackling a far wider class of challenging geometric tasks.*

### 1. introduction

Diffusion models have emerged as a transformative force in generative AI. Initially developed for image synthesis, they have quickly proven to be among the most powerful and versatile generative models across a wide range of media, including audio, video, and 3D content. Their ability to progressively denoise random signals into coherent and high-fidelity samples has enabled breakthrough applications, from photorealistic image generation to controllable editing and cross-modal translation. Beyond their re-




Figure 1. We introduce a visual diffusion approach to solving hard geometric problems directly in pixel space. Shown here on the Inscribed Square Problem, where we task the model with finding a square such that all of its four vertices lie on a given curve. Our method uncovers diverse approximate solutions, corresponding to different random seeds.

markable empirical success, diffusion models are increasingly recognized as a general framework for modeling complex, multimodal distributions.

In this work, we take a different perspective on diffusion models: rather than focusing on their creative generative capacity, we demonstrate their potential as solvers of hard geometric problems. We show that the sampling process of diffusion can be harnessed to directly reason about and discover geometric structures, guided only by pixel-level formulations of the problem. This visual diffusion approach allows us to treat abstract geometric challenges as image generation tasks, bridging the gap between visual synthesis and mathematical problem-solving.

Diffusion models have been used in various contexts

to tackle optimization and reasoning problems, including combinatorial tasks such as the traveling salesman problem [27, 38, 42]. These approaches typically formulate the problem in symbolic or graph-based representations, leveraging the probabilistic nature of diffusion to search solution spaces. In contrast, our method operates purely in the visual domain. By representing geometric problems as images and reasoning directly in pixel space, we exploit the intrinsic strength of diffusion models in handling multimodal distributions and ambiguous solutions. This visual formulation makes our approach fundamentally distinct from prior problem-solving applications of diffusion.

To ground our approach, we begin with the Inscribed Square Problem, a long-standing problem that asks whether every simple closed curve in the plane admits an inscribed square. The problem is still unsolved in the general case. Furthermore, a given curve may admit multiple and often very different inscribed squares, and enumerating them is non-trivial even in restricted settings [44]. This multiplicity naturally forms a distribution, which makes the problem especially well suited to diffusion models. Our method addresses it in an unexpected way, operating directly in image space on a visual representation of the geometric challenge. By starting from different random seeds, the model can uncover diverse valid squares, each corresponding to a distinct solution of the problem. Figure 1 illustrates the setting and highlights both the complexity and the variability of possible solutions.

We next illustrate the competence of our diffusion-in-image-space approach on two additional hard geometric problems. The first is the Steiner Tree Problem, which asks for the shortest possible network connecting a given set of points. Its solution may introduce auxiliary nodes, known as Steiner points, and finding the optimal configuration is NP-hard [17]. The second is the problem of connecting a set of points into a simple polygon of maximum area, which was featured in the CG:SHOP global optimization challenge of 2019 [9]. This task is known for its combinatorial complexity and strict geometric constraints, and is also NP-hard. As with the inscribed square problem, our method addresses these problems directly in image space, operating in the pixel domain. While discretization at finite resolution imposes limitations, it nonetheless enables valid approximations to problems that are otherwise extremely hard, with solutions that can be naturally refined. We evaluate this aspect rigorously in the paper.

Training follows a deliberately simple yet effective strategy. We generate a large distribution of valid solutions directly in image space and train a diffusion process to denoise random Gaussian samples into this distribution. This strategy builds on the observation that, in many cases, constructing examples of valid solutions is far easier than deriving one for a specific input instance. Our image-space




Figure 2. Example of a curve (black) with three inscribed squares. Note that the inscribed squares are defined only by having all four vertices on the curve: they need not be fully contained within the curve, and they may overlap with each other.

formulation therefore offers notable simplicity: it requires no elaborate encodings or specialized representations. At the same time, it provides a natural entry point to a wide spectrum of hard geometric problems that admit natural visual representations, beyond the three studied in this work.

## 2. Related Work

Several works have attempted to use diffusion models in order to learn to solve problems for which no known efficient algorithm exists, mostly of combinatorial nature. Most existing diffusion works operate directly on the parameter space of the problem. DIFUSCO [42], a graph-based diffusion framework, casts a broad family of NP-complete problems into  $\{0, 1\}^N$  indicator vectors and learns a denoiser over graphs. They compare continuous (Gaussian) and discrete (Bernoulli) noise processes, and show strong results on traveling salesman problem (TSP) and maximum independent set (MIS). Complementing the supervised approach, the T2T [27] line of work learns a distribution of high-quality solutions during training and then performs gradient-guided optimization at test time by iteratively noising the current solution and denoising while guiding towards lower energy solutions of some relaxed objective, achieving competitive quality–efficiency trade-offs on TSP and MIS. Fast T2T [28] significantly speeds this up via an optimization-consistency objective, matching or surpassing multi-step diffusion solvers performance with single-step generation plus a single gradient step.

Closer to our approach, some methods solve constrained problems in pixel-space representation. [20] train an unconditional diffusion model on pixel-space representations of TSP instances. They then solve new instances with stochastic optimization using a differential renderer, utilizing the prior of the learned model. Differently from us, they optimize a parametric representation of the solution to a given instance, while we generate solutions with DDIM sampling of a conditional model. [50] train a noise-prediction UNet

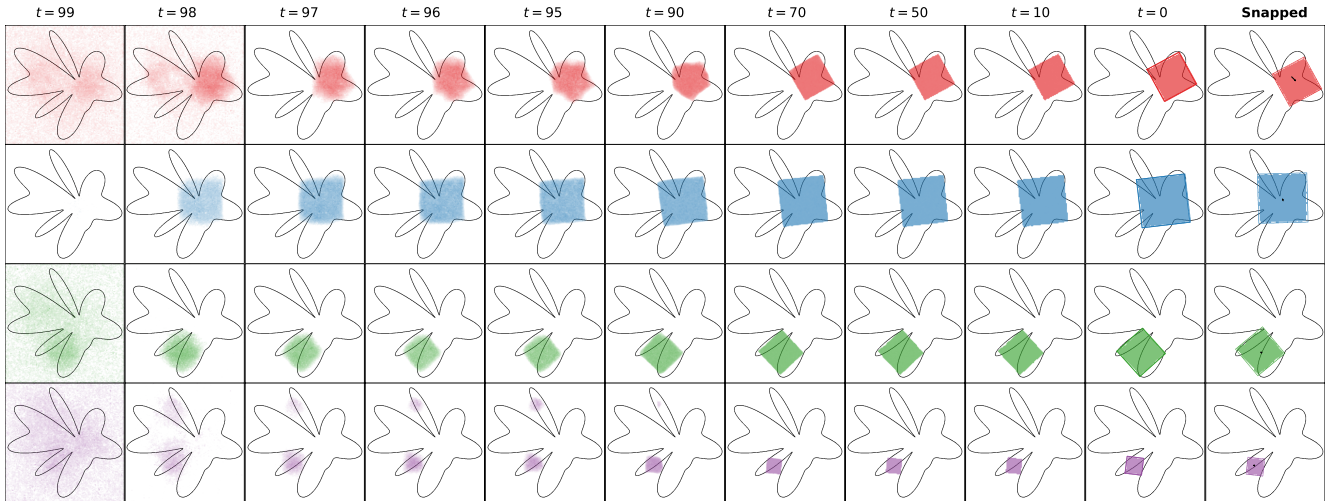


Figure 3. **Inscribed square  $x_0$  predictions across denoising steps.** Each row corresponds to a different seed (inscribed square). Columns show selected  $x_0$  predictions for decreasing timesteps  $t$  from left to right (leftmost:  $t=T$ ; penultimate:  $t=0$ ). For  $t \neq 0$  we render only the filled mask; at  $t=0$  we also draw square edges and the minimum-area bounding box. The rightmost column (“snapped”) shows the rigidly snapped version of the  $t=0$  prediction on the curve, with an arrow from the original centroid to the snapped centroid.

in pixel space on a visual representation of Sudoku, which is another NP-hard combinatorial problem. They depart from fully-parallel image diffusion by (i) assigning individual noise levels to patches, and (ii) sampling patches in a learned order or hand-crafted order. They demonstrate that sampling order matters, and can substantially outperform a conditional DDPM baseline that operates fully in parallel.

### 3. Inscribed Square Problem

We present our visual diffusion approach as a solver for hard geometric problems by structuring the paper around a set of case studies on well-known challenges. Each case study begins with the problem statement and its mathematical context, followed by a brief review of existing methods. We then describe how our image-space diffusion formulation addresses the task, highlighting both its capabilities and limitations. The first and central case we study is the inscribed square problem, which serves as an illustrative entry point into our approach.

**Problem Statement** The *Inscribed Square Problem*, also known as Toeplitz’s *Square Peg Problem* first posed in 1911, asks whether every Jordan curve in the Euclidean plane contains four points that form a square. Formally, the conjecture states that for every Jordan curve  $C \subset \mathbb{R}^2$ , there exist four points  $\{p_1, p_2, p_3, p_4\} \subset C$  such that  $p_1, p_2, p_3, p_4$  are the vertices of a non-degenerate square.

Figure 2 illustrates this setting, showing a curve with several inscribed squares.

The problem has since been resolved in several restricted settings. Some works show the conjecture holds for convex and piecewise analytic curves [11–13] and later results

prove it for  $C^1$ -smooth curves and for curves of finite total curvature or generic  $C^1$  curves [4, 41]. More recent work demonstrates validity under additional low-regularity assumptions [43], yet it remains open for the general Jordan curve case.

**Existing Methods** Algorithmic approaches for finding inscribed squares exist primarily for discrete or polygonal cases. For convex polygons, efficient procedures have been developed to detect or enumerate inscribed squares in sub-quadratic time [5, 40]. Discrete relaxations on grid polygons have likewise been studied and verified computationally for bounded cases [35]. All these methods represent the curve symbolically or combinatorially, relying on exact geometric descriptions. In contrast, our work formulates the problem directly in pixel space.

**Method** We address the inscribed square problem by formulating it entirely in image space and training a conditional diffusion model to recover inscribed squares from noisy inputs. To this end, we generated a dataset of synthetic samples. Each sample consists of a smooth Jordan curve  $C$  together with one to five inscribed squares. The curves were constructed procedurally such that they pass through the square vertices while avoiding self-intersections, and both curves and squares were rasterized into  $128 \times 128$  binary images. Each training example pairs a curve image with one inscribed square. Full details of the curve-generation process are provided in the supplementary material, under Section 7.1.

Our model follows the standard diffusion framework with a U-Net [37] backbone with self-attention layers [45], similar to those employed in text-to-image diffusion models

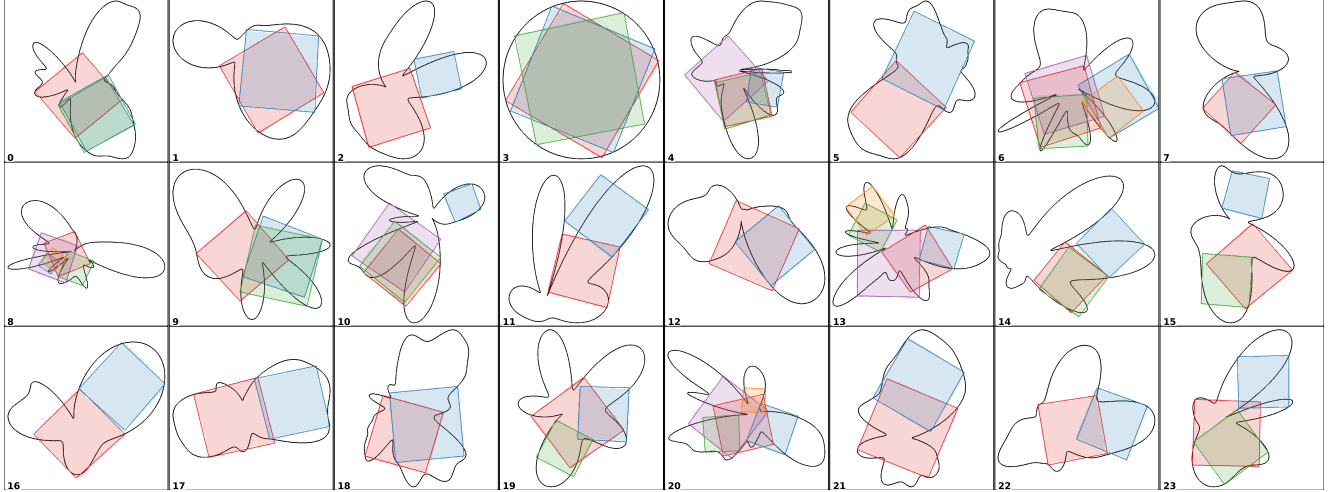


Figure 4. Solutions produced by our model. Each Jordan curve (black) is accompanied by predicted inscribed squares (colored).

such as Stable Diffusion [36]. The conditioning signal is the clean binary image of the curve, while the ground truth  $x_0$  is the clean image of the square, both represented in two-dimensional pixel space. The conditioning image is concatenated as an additional channel to the noisy input  $x_t$  at each timestep, effectively treating the curve as a non-noisy color channel. We train with 100 denoising steps using a standard noise schedule and mean-squared error objective, enabling the model to transform noisy square images into valid inscribed squares consistent with the given curve.

During sampling, at each step the network predicts an  $x_0$  estimate of the square conditioned on the curve. To inspect this behavior, we visualize  $x_0$  predictions at a subset of timesteps arranged left-to-right from  $t=T$  to  $t=0$  (penultimate column) in Fig. 3.

**Square Enhancement (Snapping)** As a final post-processing step, we refine each predicted square  $\hat{S}$  by snapping it to the conditioning curve  $C$  of the respective problem instance. Let  $V(S) = \{p_1, p_2, p_3, p_4\}$  denote the set of four vertices of a square  $S$ . To quantify how well  $S$  aligns with  $C$ , we define the negative average corner-to-curve distance as the alignment score:

$$\mathcal{A}(S, C) = -\frac{1}{4} \sum_{p \in V(S)} \text{dist}(p, C), \quad (1)$$

where  $\text{dist}(p, C)$  is the Euclidean distance from vertex  $p$  to the curve  $C$ .

We then apply a small rigid transformation  $(R_\theta, t)$  to  $\hat{S}$  and select the configuration that maximizes this alignment score. In practice, we approximate this optimization with a coarse grid search over small rotations and translations (see Section 7.2 for details). This procedure nudges the predicted square so that its corners sit more closely on  $C$ , as visualized in the rightmost column of Fig. 3.

Table 1. **Evaluation Results.** We report alignment  $\mathcal{A}(S, C)$  and squareness  $\mathcal{Q}$  under three conditions: before snapping, after snapping, and ground truth (GT).

w/o snapping		w/ snapping		Data (GT)	
Align $\uparrow$	Square $\uparrow$	Align $\uparrow$	Square $\uparrow$	Align $\uparrow$	Square $\uparrow$
-1.60	0.892	-0.90	0.891	-0.14	0.924

**Evaluation** Since multiple valid squares could potentially fit the curves beyond the ground truth squares used in construction, we avoid comparing distances to the closest ground truth square and instead focus on evaluating the geometric properties of our generated solutions. We evaluate our method along two complementary axes: *alignment* and *quality*. For alignment, we report the score  $\mathcal{A}(S, C)$  defined in Eq. 1, which directly measures how well the vertices of a predicted square lie on the conditioning curve. For quality, we introduce a *squareness metric* that captures how close a predicted shape is to a valid square. Given a predicted square  $S$ , let  $\text{area}(S)$  denote its contour area, and let  $(w, h)$  denote the side lengths of its minimum-area enclosing rectangle. We define:

$$\mathcal{Q}(S) = \frac{\text{area}(S)}{w \cdot h} \cdot \exp\left(-2 \left| \frac{\max(w, h)}{\min(w, h)} - 1 \right| \right). \quad (2)$$

This produces a score in  $[0, 1]$  that is high only when  $S$  tightly fills a nearly equilateral rectangle, i.e., when it closely resembles a true square.

We report both alignment and quality metrics under three conditions: (i) predictions before snapping, (ii) predictions after snapping, and (iii) the ground-truth squares from the dataset (Tab. 1). This evaluation disentangles the intrinsic generative ability of the diffusion model from the gains achieved by the geometric snapping refinement.




Figure 5. Example of a Steiner Minimal Tree. Left: The input terminal nodes, colored in red. Right: The Steiner Minimal Tree, where auxiliary Steiner points are colored in dark gray.

**Interpretation of Evaluation Results** The evaluation demonstrates that our model consistently produces shapes that closely approximate true inscribed squares with strong accuracy. Even though the alignment of predicted squares with the conditioning curve is not always pixel-perfect, the snapping step leads to a substantial refinement, bringing the results impressively close to the ground truth. In practice, this shows that the diffusion process is highly effective at capturing both the structural regularity and the correct placement of squares, with only minimal residual deviation that often manifests at the sub-pixel level. Importantly, such deviations are expected given the inherent discretization of the pixel domain, which naturally puts an upper bound even on the ground truth results. Within this setting, our method reliably recovers high-quality approximations of valid inscribed squares. The reported numbers confirm that the model does not merely suggest plausible candidates but in fact achieves precise and robust approximations, validating the strength of the visual diffusion framework as a solver for this classical geometric challenge.

#### 4. Steiner Tree Problem

The second problem we cover in our case study is that of the Steiner Tree Problem. The Steiner tree problem asks, given a set of terminal points, to find a network of minimum total length that connects all terminals, where the construction is allowed to introduce additional points (Steiner points) to reduce length. In the Euclidean variant, which we focus on, Steiner points may be placed anywhere in the plane. The Steiner formulation is central to many applications where minimizing connection cost is critical. Some typical use cases for it include telecommunication [47], PCB routing [6], and in infrastructure layout (roads, pipelines) [8, 39].

**Problem statement.** Formally, given a finite set of terminals  $P = \{p_1, \dots, p_n\} \subset \mathbb{R}^2$ , the Euclidean Steiner Tree (EST) problem asks for a straight-line embedded tree  $T = (V, E)$  with  $P \subseteq V$  that minimizes total Euclidean length  $L(T) = \sum_{uv \in E} \|u - v\|_2$ . Vertices in  $V \setminus P$  are *Steiner points*. An optimal solution is called a *Steiner minimal tree* (SMT) and its length is denoted  $L^*(P)$  [18]. Fig-




Figure 6. **Steiner tree  $x_0$  predictions across denoising steps.** Each row corresponds to a different seed. Columns show selected  $x_0$  predictions for decreasing timesteps  $t$  from left to right (leftmost:  $t=T$ ; rightmost:  $t=0$ ). Input points are overlaid in red.

ure 5 illustrates an instance of the problem and its optimal solution.

**Complexity and Approximability.** For general  $n$  (the number of initial given points), the EST problem is NP-hard already in the plane [17]. However, there is a polynomial-time approximation scheme (PTAS) for Euclidean Steiner tree in fixed dimensions: for any fixed  $\varepsilon > 0$  one can compute a  $(1 + \varepsilon)$ -approximate tree in time polynomial in  $n$  (for fixed  $\varepsilon$  and dimension).

**Structure and Properties** SMTs in the plane are highly structured: (i) no two edges cross; (ii) each Steiner point has degree exactly 3 and the three incident edges meet at  $120^\circ$ ; (iii) all angles in the tree are at least  $120^\circ$  (at terminals that have degree  $> 1$ ); (iv) the number of Steiner points is at most  $n - 2$ ; and (v) all Steiner points lie in the convex hull of  $P$  [3, 22]. These properties provide strong geometric constraints that are exploited by both exact and approximate methods, as discussed next.

**Algorithmic Approaches.** Exact algorithms typically rely on generating locally optimal full Steiner trees and then selecting a minimum-length subset, with implementations such as GEOSTEINER solving large 2D instances [24, 49]. On the approximation side, PTASes based on dissection or guillotine methods provide near-optimal guarantees in fixed dimensions [2, 32].

**Learning-based Approaches** While classical exact/approximate algorithms dominate practice, there have been several attempts to use learning-based solvers for the problem. For the Euclidean case, Wang et al. propose *Deep-Steiner*, which casts SMT construction as a sequential decision process: it discretizes the continuous search space, prunes candidates via KNN/MST neighborhoods, and then uses an attention-based policy trained with REINFORCE to add Steiner points iteratively [48]. For non-Euclidean variants (rectilinear SMT and

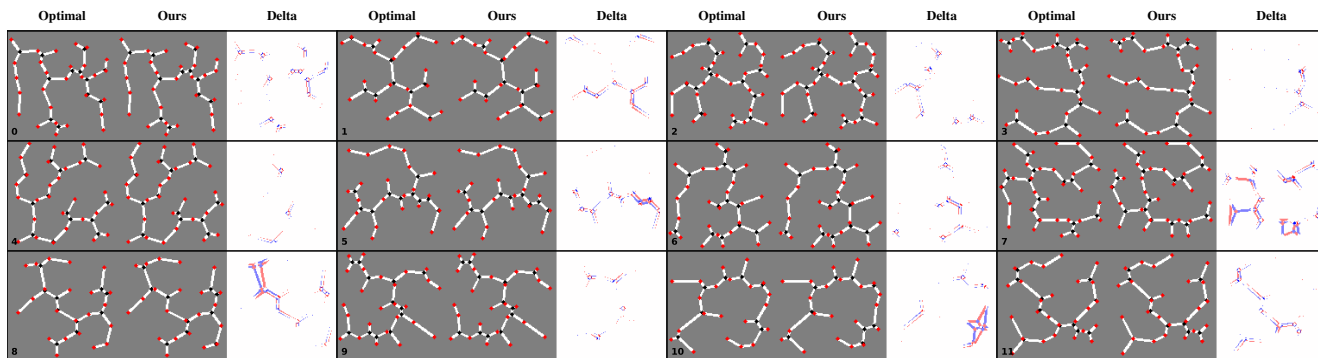


Figure 7. **Visual comparison of Steiner tree solutions.** Each example is shown as a triplet of images: the optimal solution (**Optimal**), our model’s solution (**Ours**), and the difference between them (**Delta**). Input points are overlaid as red circles.

Table 2. **Steiner Tree Evaluation Results.** Comparison of our method, MST, and random solutions. Reported are valid tree rates and mean Euclidean length ratios ( $\pm$  std) relative to the optimal solution across input point ranges.

Input Points	Valid Rate	Ours Ratio Mean $\pm$ Std	MST Ratio Mean $\pm$ Std	Random Ratio Mean $\pm$ Std
10-20	0.996	1.0008 $\pm$ 0.0005	1.036 $\pm$ 0.012	1.834 $\pm$ 0.236
21-30	0.986	1.0018 $\pm$ 0.0011	1.041 $\pm$ 0.009	1.904 $\pm$ 0.182
31-40	0.834	1.0044 $\pm$ 0.0035	1.047 $\pm$ 0.007	1.898 $\pm$ 0.165
41-50	0.334	1.0092 $\pm$ 0.0055	1.052 $\pm$ 0.007	1.860 $\pm$ 0.142

graph STP), several learning-based methods have been explored, including RL and GNN-driven solvers and mixed neural–algorithmic pipelines [1, 10, 25, 29, 30, 34].

**Method** We generate a synthetic dataset by sampling random points (between 10 and 20 for each instance), and finding their SMT using the GeoSteiner solver [24]. Each solution is then rasterized into a grayscale image with different values for edges, nodes and background. The full details for data generation are found in Section 7.3. We employ the same U-Net backbone as in the inscribed square experiment. The conditioning input consists of the rasterized terminal points, concatenated as an additional channel to the noisy input  $x_t$  at each denoising step.

Recovering a graph structure from the generated image is performed in two stages. We begin by node detection, where we binarize the output with a threshold and detect centers of connected components. The centroid of each blob is then taken as a node position, with nodes falling within a small radius of an input terminal being snapped to that terminal’s location. In the second stage, we extract edges by considering the complete graph over the detected nodes. For each candidate edge, we compute the fraction of pixels along the straight line segment that are marked as foreground. If this fraction exceeds a threshold (70% in our implementation), the edge is retained. If two vertices are very close to each other, we assume they are connected via an edge. In cases of ambiguity where multiple potential edges

with a shared node overlap, we retain the shortest one and discard the rest.

**Evaluation** We evaluate our trained model on a test set containing instances with 10-20 input points, matching the number of points seen during training, and four other test sets containing 11-20, 21-30, 31-40 and 41-50 input points. After extracting the graph from the generated solution, we check the validity of the solution by verifying that the resulting graph is a tree and that it contains all of the input points. If the solution is valid, we then measure the total Euclidean length of the tree. For each instance, we generate in parallel 10 solutions from different noise seeds and select the one with minimal total edge length that is also valid. In Table 2 we report for each test set the rate of valid solutions as well as the mean ratio between the total Euclidean length of the best solution produced by our model compared to that of the optimal solution ( $L^*(P)$ ). For comparison, we also report the ratio between the total Euclidean length of a random planar tree and the optimal solution and that of the solution produced by the minimum spanning tree of the full graph and  $L^*(P)$ .

Our model is able to successfully produce high quality solutions even for instances with markedly more input points than were seen during training, and often produces solutions that align with the optimal ones (see Figure 7).

While there is generally a one-to-one coupling between an input and the optimal solution, for some instances the model still produces variations, often of similar quality, for different noise initializations (see Figure 6). However, some of these variations can happen to be invalid, especially for instances with a large number of input points. This is evident in the third row, where the solution produced by the noise initialization contains a loop and is not a tree.

## 5. Maximum Area Polygon Problem

The third problem we attempt to tackle with our approach is the Maximum Area Polygonization Problem (MAXAP), a well-established problem in computational geometry.

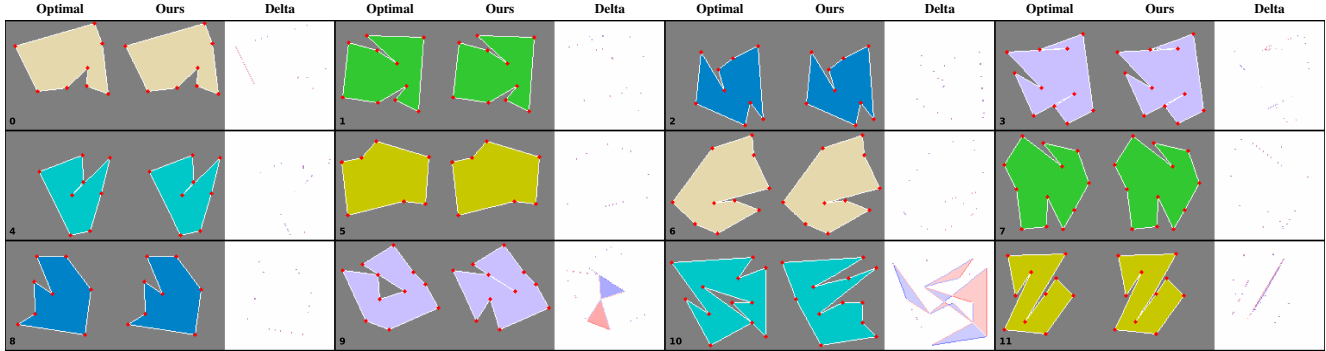


Figure 8. **Qualitative comparison of maximum-area polygon solutions.** Each example is presented as a horizontal triplet: the optimal polygon (**Optimal**), the polygon generated by our model (**Ours**), and the difference (**Delta**). Red and blue regions in the difference maps indicate areas unique to the optimal and generated polygons, respectively. Even when visual discrepancies occur, the exclusive regions are typically of similar total area, resulting in solutions of comparable area. Input points are overlaid as red circles on both Optimal and Ours.

Given a set of vertices in the plane, the problem asks to find a simple polygon (a polygon that does not intersect itself and has no holes) that passes through all the vertices and has the largest possible area.

MAXAP is known to be NP-complete [14, 15] and difficult to solve both in theory and in practice [16], with no known algorithm that provides better than  $\frac{1}{2}$ -approximation factor in polynomial time. Furthermore, deciding whether there exists a simple polygon that contains strictly more than  $\frac{2}{3}$  of the area of the convex hull is also NP-complete [14]. Exact approaches based on integer programming are able to solve instances with up to 25 points [16], while a recent mixed-integer-programming approach is able to solve instances with up to 50 points [21]. Heuristic approaches are commonly applied to larger instances, including constrained triangulations [26], simulated annealing [19, 26], and divide-and-conquer strategies for very large point sets of up to 1,000,000 points [7, 19].

**Method** To address the maximum area polygonization problem, we adopt the same visual diffusion architecture as in the previous two tasks, using an identical U-Net backbone. As with the previous methods, we generate a synthetic dataset of examples to train the model on. For each training example we sample input points randomly on the grid, and compute their optimal polygonizations through exhaustive search over all valid simple polygons, which is feasible at this scale using a DFS procedure. Each polygon is rasterized to an image, while the input points are rasterized into a separate image that is concatenated as an additional conditioning channel. The data generation procedure is described in full in Section 7.4

At inference time, we recover the polygon structure from the generated image by testing candidate edges between all point pairs. Each edge is retained if more than 70% of its pixels align with foreground edge pixels in the output. The resulting set of edges is then validated to ensure that no intersections occur, with mild tolerance for nearly parallel

Table 3. **Maximum Area Polygon Evaluation Results.** Comparison of our method, random polygons, and optimal solutions. Metrics include polygon validity rate, mean area ratio ( $\pm$  std), and optimal solution rate for different input point ranges.

Input Points	Valid Rate	Ours Ratio Mean $\pm$ Std	Random Ratio Mean $\pm$ Std	Opt. Rate
7-12	0.953	0.988 $\pm$ 0.020	0.771 $\pm$ 0.136	0.574
13-15	0.620	0.962 $\pm$ 0.041	0.477 $\pm$ 0.271	0.062

overlaps. Finally, we search for a simple cycle that passes through all input vertices, which we consider the recovered polygonization produced by our model.

**Evaluation** We evaluate the trained model on a test set containing 7–12 points, matching the range used during training. To further assess generalization, we also test the model on a set with 13–15 points. For each instance, we generate 10 candidate solutions from different random seeds and select the one that achieves the largest area while remaining valid. In Table 3, we report the mean and standard deviation of the ratio between the best solution found for each instance and the corresponding optimal solution, and show a comparison against random simple polygons on the same point sets. We also report the rate of valid polygons (the proportion of instances for which a valid polygon was produced), as well as the frequency with which the recovered polygon coincides exactly with the optimal one.

Figure 8 shows qualitative examples of polygons produced by our model. In many cases, the generated polygons align almost perfectly with the optimal ones. When discrepancies occur, the differences often balance out, with areas lost in one region largely compensated elsewhere (see instances 9, 10 in the figure). Owing to the non-local nature of the problem, instances with a larger number of points are substantially more challenging, and the rate of valid solutions drops noticeably for the 13–15 point test set. Typical failure cases include polygons that do not pass through all




Figure 9. **Maximum area polygon  $x_0$  predictions across denoising steps.** Each row corresponds to a different seed. Columns show selected  $x_0$  predictions for decreasing timesteps  $t$  from left to right (leftmost:  $t=T$ ; rightmost:  $t=0$ ). Input points in red.

input points or contain holes (see third and fourth rows in Figure 9). Nevertheless, whenever a valid polygon is produced, its area is typically very close to the optimal solution. We note that for this problem, like the last one, there is generally only a single optimal solution per problem instance. Therefore the benefit of training conditional diffusion models which are generally used to learn conditional distributions in order to solve these instances is not immediately clear. In Section 8.2 we demonstrate on the MAXAP problem that there is a performance advantage over using a regression model even for problems of this kind.

## 6. Discussion and Conclusions

In this work, we presented visual diffusion as a general framework for approximating solutions to notoriously hard geometric problems. Through three case studies, the *Inscribed Square* problem, the *Steiner Tree* Problem, and the *Simple Polygon* Problem, we demonstrated that diffusion models can operate in image space to uncover valid geometric structures.

We do not claim that our method outperforms specialized solvers tailored to any single problem. Indeed, for each of these problems, carefully designed algorithms may yield more efficient or more accurate solutions. Instead, our contribution is to reveal a paradigm: visual diffusion provides a single, simple framework that applies across a diverse set of problems without requiring custom formulations. Specifically, each task uses the same diffusion architecture without modification, varying only in task-specific training data and minimal decoding to recover the symbolic structure.

Our approach produces accurate and diverse approximations, naturally recovering multiple valid solutions through diffusion, as illustrated in the *Inscribed Square* problem. These solutions can be further refined if desired.

The most obvious limitation of our approach lies in the

discreteness of the underlying representation, yet this constraint behaves much more mildly than one might expect. At a high level, our methods operate with finite precision, but the problems we consider possess stability properties that make such approximations structurally reliable. Consequently, the approximate solutions typically retain the same structure as the exact ones. For example, Steiner Tree topology is known to be invariant to small perturbations of the terminal vertices. Thus, the solutions produced by our algorithms maintain the correct structure under general input assumptions (see [23, 31]). Similarly, because the minimal tree length is a continuous function of the input coordinates, the length calculated using quantized vertices approximates the true value and converges as the grid step decreases.

More broadly, approximate solutions are often sufficient for hard geometric problems where exact solvers are computationally expensive or unavailable [46, 51], and a fast solution preserving the correct structure can be more valuable than an intractable exact one.

Importantly, we also observe that models trained on relatively simple instances generalize to more complex inputs, such as handling a larger number of points than those seen in training. Our inference time remains constant regardless of input size, whereas classical solvers for these problems scale polynomially or exponentially, suggesting potential advantages at larger scales.

Despite the diversity of the problems they are trained to solve, the models exhibit a consistent behavior, evident in the denoising progression (Figures 3, 6 and 9). In the early steps of the sampling process, the global structure of the solution becomes apparent, suggesting that the essence of the solution lies primarily in low-frequency geometric features that can be recovered quickly. The subsequent denoising steps refine these structures to achieve high accuracy.

This observation indicates that inference time could be further optimized, with only a small trade-off in precision, by using denoising schedulers that allocate more of the sampling steps budget to earlier timesteps. More broadly, this progressive reasoning mirrors how humans intuitively reason about geometric problems: they first sketch a mental image of a coarse solution in their mind, which they can then try to translate to a concrete solution, with details settled upon after the initial structure is clear.

The key message of this work is that image diffusion models, long celebrated for their generative capacity, also serve as geometric solvers. While our experiments focus on 2D to make the case concrete, evaluating whether similar efficiency can be achieved in 3D is left for future work, potentially with volumetric or video-based models. This outlook opens the door to exploring a wide spectrum of geometric challenges under a single methodology and suggests new opportunities for bridging generative modeling with geometric problem solving.

## References

- [1] Reyhan Ahmed, Md Asadullah Turja, Faryad Darabi Sahneh, Mithun Ghosh, Keaton Hamm, and Stephen Kobourov. Computing steiner trees using graph neural networks, 2021. [6](#)
- [2] Sanjeev Arora. Polynomial-time approximation schemes for Euclidean TSP and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998. [5](#)
- [3] Marcus Brazil, Pawel Winter, and Martin Zachariasen. Flexibility of steiner trees in uniform orientation metrics. pages 196–208, 2004. [5](#)
- [4] Jason Cantarella, Elizabeth Denne, and John McCleary. Transversality for configuration spaces and the “square-peg” theorem, 2021. [3](#)
- [5] Bernard Chazelle. The Polygon Containment Problem. *Advances in Computing Research*, 1(1):1–33, 1983. [3](#)
- [6] Chris Chu and Yiu-Chung Wong. Flute: Fast lookup table based rectilinear steiner minimal tree algorithm for VLSI design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 27(1):70–83, 2008. [5](#)
- [7] Loïc Crombez, Guilherme D. da Fonseca, and Yan Gerard. Greedy and local search heuristics to build area-optimal polygons. *ACM J. Exp. Algorithmics*, 27, 2022. [7](#)
- [8] Ziyuan Cui, Hai Lin, Yan Wu, Yufei Wang, and Xiao Feng. Optimization of pipeline network layout for multiple heat sources distributed energy systems considering reliability evaluation. *Processes*, 9(8):1308, 2021. [5](#)
- [9] Erik D. Demaine, Sándor P. Fekete, Phillip Keldenich, Dominik Krupke, and Joseph S. B. Mitchell. Area-optimal simple polygonizations: The cg challenge 2019. *ACM J. Exp. Algorithmics*, 27, 2022. [2](#)
- [10] Haizhou Du, Zong Yan, Qiao Xiang, and Qinqing Zhan. Vulcan: Solving the steiner tree problem with graph neural networks and deep reinforcement learning. *arXiv*, 2021. [6](#)
- [11] Arnold Emch. Some properties of closed convex curves in a plane. *American Journal of Mathematics*, 35(4):407–412, 1913. [3](#)
- [12] Arnold Emch. On the medians of a closed convex polygon. *American Journal of Mathematics*, 37(1):19–28, 1915.
- [13] Arnold Emch. On some properties of the medians of closed continuous curves formed by analytic arcs. *American Journal of Mathematics*, 38(1):6–18, 1916. [3](#)
- [14] Sándor P Fekete. *Geometry and the travelling salesman problem*. University of Waterloo, 1992. [7](#)
- [15] Sándor P Fekete and William R Pulleyblank. Area optimization of simple polygons. In *Proceedings of the ninth annual symposium on computational geometry*, pages 173–182, 1993. [7](#)
- [16] Sándor P. Fekete, Andreas Haas, Phillip Keldenich, Michael Perk, and Arne Schmidt. Computing area-optimal simple polygonizations, 2021. [7](#)
- [17] M. R. Garey, R. L. Graham, and D. S. Johnson. The complexity of computing steiner minimal trees. *SIAM Journal on Applied Mathematics*, 32(4):835–859, 1977. [2](#), [5](#)
- [18] Edgar N. Gilbert and Henry O. Pollak. Steiner minimal trees. *Siam Journal on Applied Mathematics*, 16:1–29, 1968. [5](#)
- [19] Nir Goren, Efi Fogel, and Dan Halperin. Area optimal polygonization using simulated annealing. *ACM J. Exp. Algorithmics*, 27, 2022. [7](#)
- [20] Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as plug-and-play priors, 2023. [2](#)
- [21] Hipólito Hernández-Pérez, Jorge Riera-Ledesma, Inmaculada Rodríguez-Martín, and Juan-José Salazar-González. Optimal area polygonisation problems: Mixed integer linear programming models. *European Journal of Operational Research*, 2025. [7](#)
- [22] F. K. Hwang and Dana S. Richards. Steiner tree problems. *Networks*, 22(1):55–89, 1992. [5](#)
- [23] Alexander O Ivanov and Alexei A Tuzhilin. *Minimal Networks The Steiner Problem and Its Generalizations*. CRC press, 1994. [8](#)
- [24] D. Juhl, D. M. Warme, P. Winter, and M. Zachariasen. The GeoSteiner software package for computing steiner trees in the plane: An updated computational study. *Mathematical Programming Computation*, 10(4):487–532, 2018. [5](#), [6](#), [11](#)
- [25] Andrew B. Kahng, Robert R. Nerem, Yusu Wang, and Chien-Yi Yang. Nn-steiner: A mixed neural-algorithmic approach for the rectilinear steiner minimum tree problem, 2023. [6](#)
- [26] Julien Lepagnot, Laurent Moalic, and Dominique Schmitt. Optimal area polygonization by triangulation and visibility search. *ACM J. Exp. Algorithmics*, 27, 2023. [7](#)
- [27] Yang Li, Jinpei Guo, Runzhong Wang, and Junchi Yan. T2t: From distribution learning in training to gradient search in testing for combinatorial optimization. In *Advances in Neural Information Processing Systems*, 2023. [2](#)
- [28] Yang Li, Jinpei Guo, Runzhong Wang, Hongyuan Zha, and Junchi Yan. Fast t2t: Optimization consistency speeds up diffusion-based training-to-testing solving for combinatorial optimization. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. [2](#)
- [29] Jinwei Liu, Guojie Chen, and Evangeline F. Young. Rest: Constructing rectilinear steiner minimum tree via reinforcement learning. In *Proc. DAC*, pages 1135–1140, 2021. [6](#)
- [30] Ruizhi Liu, Zhisheng Zeng, Shizhe Ding, Jingyan Sui, Xingquan Li, and Dongbo Bu. Neuralsteiner: Learning steiner tree for overflow-avoiding global routing in chip design. In *Proc. NeurIPS*, 2024. [6](#)
- [31] Wouter Meulemans, Bettina Speckmann, Kevin Verbeek, and Jules Wulms. A framework for algorithm stability and its application to kinetic euclidean msts. In *LATIN 2018: Theoretical Informatics*, pages 805–819, Cham, 2018. Springer International Publishing. [8](#)
- [32] Joseph S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple PTAS for geometric TSP,  $k$ -MST, and related problems. *SIAM Journal on Computing*, 28(4):1298–1309, 1999. [5](#)
- [33] Joseph O’Rourke. *Computational Geometry in C*. Cambridge University Press, 2 edition, 1998. Section 1.3: Area of Polygon. [11](#)
- [34] Youngjoon Park, Han-Seul Jeong, Kyunghyun Lee, Sungdong Yoo, Hyungseok Song, and Woohyung Lim. Steben: Steiner tree problem benchmark for neural combinatorial optimization on graphs, 2025. [6](#)

- [35] Ville H. Pettersson, Helge A. Tverberg, and Patric R.J. Östergård. A note on toeplitz' conjecture. *Discrete & Computational Geometry*, 51(3):722–728, 2014. [3](#)
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022. [4](#)
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015. [3](#)
- [38] Sebastian Sanokowski, Sepp Hochreiter, and Sebastian Lehner. A diffusion model framework for unsupervised neural combinatorial optimization, 2025. [2](#)
- [39] Justus Schwartz and Jürg Stückelberger. Computing lower bounds for steiner trees in road network design. In *Proceedings of the 7th International Symposium on Operations Research and Its Applications (ISORA'08)*, pages 172–181, Lijiang, China, 2008. ORSC & APORC. [5](#)
- [40] Micha Sharir and Sivan Toledo. External polygon containment problems. *Computational Geometry*, 4(2):99–118, 1994. [3](#)
- [41] Walter Stromquist. Inscribed squares and square-like quadrilaterals in closed curves. *Mathematika*, 36(2):187–197, 1989. [3](#)
- [42] Zhiqing Sun and Yiming Yang. Difusco: Graph-based diffusion solvers for combinatorial optimization, 2023. [2](#), [13](#)
- [43] Terence Tao. An integration approach to the toeplitz square peg problem, 2017. [3](#)
- [44] Wouter van Heijst. The algebraic square peg problem, 2014. [2](#)
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017. [3](#)
- [46] Vijay V. Vazirani. *Approximation Algorithms*. Springer, 2001. [8](#)
- [47] Stefan Voss. *Steiner Tree Problems in Telecommunications*, pages 459–492. 2006. [5](#)
- [48] Siqi Wang, Yifan Wang, and Guangmo Tong. Deep-steiner: Learning to solve the euclidean steiner tree problem, 2022. [5](#), [14](#)
- [49] D. M. Warme, P. Winter, and M. Zachariasen. *Exact Algorithms for Plane Steiner Tree Problems: A Computational Study*, pages 81–116. Springer US, Boston, MA, 2000. [5](#)
- [50] Christopher Wewer, Bart Pogodzinski, Bernt Schiele, and Jan Eric Lenssen. Spatial reasoning with denoising models, 2025. [2](#)
- [51] David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011. [8](#)