Check for updates

Text-to-Vector Generation with Neural Path Representation

PEIYING ZHANG, City University of Hong Kong, China NANXUAN ZHAO, Adobe Research, USA JING LIAO^{*}, City University of Hong Kong, China



Fig. 1. Examples of text-guided vector graphics generated by our framework, with clear and valid layer-wise vector paths.

Vector graphics are widely used in digital art and highly favored by designers due to their scalability and layer-wise properties. However, the process of creating and editing vector graphics requires creativity and design expertise, making it a time-consuming task. Recent advancements in text-to-vector (T2V) generation have aimed to make this process more accessible. However, existing T2V methods directly optimize control points of vector graphics paths, often resulting in intersecting or jagged paths due to the lack of geometry constraints. To overcome these limitations, we propose a novel neural path representation by designing a dual-branch Variational Autoencoder (VAE) that learns the path latent space from both sequence and image modalities. By optimizing the combination of neural paths, we can incorporate geometric constraints while preserving expressivity in generated SVGs. Furthermore, we introduce a two-stage path optimization method to improve the visual and topological quality of generated SVGs. In the first stage, a pre-trained text-to-image diffusion model guides the initial generation of complex vector graphics through the Variational Score Distillation (VSD) process. In the second stage, we refine the graphics using

*Corresponding author

Authors' Contact Information: Peiying Zhang, City University of Hong Kong, Hong Kong, China, zhangpeiying17@gmail.com; Nanxuan Zhao, Adobe Research, San Jose, USA, nanxuanzhao@gmail.com; Jing Liao, City University of Hong Kong, Hong Kong, China, jingliao@cityu.edu.hk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

@ 2024 Copyright held by the owner/author (s). Publication rights licensed to ACM. ACM 1557-7368/2024/7-ART36

https://doi.org/10.1145/3658204

a layer-wise image vectorization strategy to achieve clearer elements and structure. We demonstrate the effectiveness of our method through extensive experiments and showcase various applications. The project page is https://intchous.github.io/T2V-NPR.

CCS Concepts: • Computing methodologies \rightarrow Graphics file formats; Image processing; • Applied computing \rightarrow Arts and humanities.

Additional Key Words and Phrases: Vector Graphics, SVG, Diffusion Model, Text-Guided Generation

ACM Reference Format:

Peiying Zhang, Nanxuan Zhao, and Jing Liao. 2024. Text-to-Vector Generation with Neural Path Representation. *ACM Trans. Graph.* 43, 4, Article 36 (July 2024), 13 pages. https://doi.org/10.1145/3658204

1 INTRODUCTION

Vector graphics, specifically in the Scalable Vector Graphics (SVG) format, play an essential role in digital arts such as clipart, animation, and graphic design. Benefiting from their composition of geometric shapes, vector graphics are widely favored by designers due to the ease of manipulation, with the nature of resolution independence and compact file sizes. However, crafting high-quality vector graphics requires both professional expertise and considerable time investment. With the success of text-to-image (T2I) generation models [Rombach et al. 2022; Ruiz et al. 2022], recent works have started to explore text-to-vector graphics generation (T2V), aiming to make the creation more accessible to users with text prompts.

One common T2V approach is conducting existing image vectorization [Dominici et al. 2020; Ma et al. 2022; Selinger 2003] based on T2I results. However, T2I models often generate raster images with photographic and realistic styles, with intricate textures, and complex color variations. This poses a challenge when transitioning to vector graphics, where the goal is to achieve smooth geometric shapes and uniform colors. The conversion process often results in excessively complex vector elements, introducing complications for subsequent graphic manipulations and deviating from the simplicity and clarity expected in vector graphics. In recent developments, a new category of T2V methods has emerged (*e.g.*, CLIPDraw [Frans et al. 2022] and VectorFusion [Jain et al. 2022]). These methods directly optimize vector graphics paths using large pre-trained visual-textual models [Rombach et al. 2022]. They represent vector graphics through parametric shape primitives (*e.g.*, cubic Bézier curves) and refine path parameters, such as control points. However, directly optimizing control points often leads to intersecting or jagged paths, due to their high degrees of freedom and the lack of geometry constraints.

Therefore, in addition to representing vectors as the explicit parametric paths, it becomes imperative to develop an efficient representation to capture both geometric relationship and shape perception. Previous works have explored learning such latent representations based on sequences of SVGs using different network architectures, such as Recurrent Neural Networks (RNN) [Ha and Eck 2017; Lopes et al. 2019; Wang and Lian 2021] and Transformers [Carlier et al. 2020; Wang et al. 2023b]. However, given the inherent complexity and diverse nature of SVGs, learning a comprehensive global SVGlevel representation is challenging. As a result, previous methods are often restricted to generating SVGs within specific categories, such as fonts [Lopes et al. 2019; Wang and Lian 2021; Wang et al. 2023b], sketches [Ha and Eck 2017; Ribeiro et al. 2020], and simple icons [Cao et al. 2023; Carlier et al. 2020; Wu et al. 2023].

To overcome the above limitations, instead of learning global SVG features, we propose a novel neural path representation that effectively learns valid geometric properties of paths. This is motivated by the finding that complex vector graphics are often composed of simple paths [Chen et al. 2023; Liu et al. 2023]. This compact representation ensures both simplicity and expressivity. To learn a neural path representation, we design a dual-branch Variational Autoencoder (VAE) to learn from both sequence and image modalities. They jointly optimize the shared path representation, and the sequence modality provides supervision for learning geometry properties, while the image modality helps to learn rendering visual features.

We further propose a two-stage path optimization method for conducting text-to-vector generation based on the learned path latent space. The underlying motivation is that obtaining high-quality vectors within one stage can be hard. In the first stage, we rely on the generation power of a large pre-trained diffusion model [Rombach et al. 2022] for generating an initial SVG. Rather than using Score Distillation Sampling (SDS) loss [Poole et al. 2022] as in VectorFusion [Jain et al. 2022], which may suffer from over-saturated, over-smooth, and low-diversity issues, we borrow the idea of Variational Score Distillation (VSD) [Wang et al. 2023a] to optimize a combination of neural paths given the text prompt. In the second stage, to further improve the geometry clarity and layer-wise structure of the initially generated SVG, we apply a path simplification and layer-wise optimization strategy to hierarchically enhance the paths. We evaluate our approach through extensive experiments using vector-level, image-level, and text-level metrics. The results demonstrate the effectiveness of our model in generating high-quality and diverse vector graphics with valid paths and layer properties, given the input text prompt. Fig. 1 shows examples of text-to-vector generation results produced by our framework. Our model enables various applications except for T2V generation, such as vector graphics customization, image-to-SVG generation, and SVG animation. Our key contributions are:

- We introduce a novel T2V generation pipeline, innovated by the idea of optimizing local neural path representation for highquality vector graphics generation.
- We propose a dual-branch VAE for learning neural path representation from both sequence and image modalities.
- We develop a two-stage text-driven neural path optimization method to guide the creation of vector graphics with valid and layer-wise SVG paths.
- We demonstrate several practical applications enabled by our pipeline.

2 RELATED WORK

2.1 Diffusion Models for T2I Generation

Recently, diffusion models (DM) have become state-of-the-art in T2I generation. It is a family of generative models, involving a forward process of perturbing the data with noise and a reverse process that gradually adds structure from noise [Ho et al. 2020; Nichol et al. 2021; Sohl-Dickstein et al. 2015]. Stable Diffusion [Rombach et al. 2022] further introduces image latent to the diffusion model, overcoming resolution limitations and enabling the generation of amazing high-resolution results. The diverse type of user-defined condition modalities enable numerous downstream applications, such as sketch-guided image generation [Voynov et al. 2023; Zhang et al. 2023b], multi-modal conditions in image synthesis [Zhan et al. 2023], and various image-to-image generation tasks [Saharia et al. 2022; Zhang et al. 2023a]. Diffusion models also support flexible text-guided image editing [Hertz et al. 2023; Meng et al. 2021] and customization tasks [Kumari et al. 2022; Ruiz et al. 2022]. To generate personalized images, the pioneering work DreamBooth [Ruiz et al. 2022] refines the weights of a diffusion model with a unique token to capture distinctive characteristics within a set of images. Subsequent works focus on fine-tuning only particular parts of the network, including low-rank weight residuals [Hu et al. 2021] and crossattention layers [Kumari et al. 2022]. Rather than aiming at raster image generation, our method leverages the powerful pre-trained T2I model as priors to guide the generation of vector graphics.

2.2 Text-to-Vector Generation

One approach in T2V generation combines T2I generation with image vectorization methods. Traditional vectorization methods involve segmenting images into regions based on color similarity [Kopf and Lischinski 2011] and fitting curves to the region boundaries [Dominici et al. 2020; Favreau et al. 2017; Hoshyari et al. 2018; Selinger 2003; Yang et al. 2015]. The advent of differentiable rendering techniques [Li et al. 2020] enhances image vectorization by enabling the use of loss functions in image space. LIVE [Ma



Fig. 2. Our system pipeline starts with learning a neural representation of paths by training a dual-branch VAE. Next, we optimize the SVG, represented with neural paths, to align with the provided text prompt, which is achieved through a two-stage path optimization process.

et al. 2022] optimizes path parameters guided by the reconstruction of the input image. Recent developments have explored training end-to-end neural networks for image vectorization [Chen et al. 2023; Reddy et al. 2021; Rodriguez et al. 2023]. However, this kind of method relies on images generated by pre-trained T2I models, which struggle to produce high-quality SVG-style images featuring clear geometric primitives and flat colors. With the development of generative AI, several community-made tools and commercial products such as Adobe Illustrator [Illustrator 2023], Illustroke [Illustroke 2024], and Kittl [Kittl 2024], offer capabilities to generate vector graphics from text prompts by employing vectorization techniques with T2I methods. To produce visually appealing SVG-style images, it is essential to fine-tune pre-trained T2I models [Hu et al. 2021; Ruiz et al. 2022]. These fine-tuning methods heavily rely on the training dataset, requiring a labor-intensive collection of images and textual tags.

Another approach for T2V generation involves directly optimizing the geometric and color parameters of SVG paths under the guidance of pre-trained vision-language models, such as the CLIP model [Radford et al. 2021] or the diffusion model [Rombach et al. 2022]. CLIP-based methods [Frans et al. 2022; Schaldenbrand et al. 2022; Song et al. 2022] utilize the image-text similarity metric within CLIP latent space to create vector graphics from text prompts. CLI-Passo [Vinker et al. 2022] converts images to object sketches using a CLIP perceptual loss, preserving the semantics and structure of the subject. Apart from the CLIP distance, Score Distillation Sampling (SDS) loss based on a T2I diffusion model is used for optimizing SVG to align with text prompts across various applications such as fonts [Iluz et al. 2023], vector graphics [Jain et al. 2022; Xing et al. 2023b], and sketches [Gal et al. 2023; Xing et al. 2023a]. However, these methods often lead to SVGs with cluttered and irregular paths, as the direct optimization of control points in parametric paths like cubic Bézier curves lacks essential geometric relationships.

2.3 Neural Representation for SVG

Previous works have explored learning various representations of SVG, and designing different network architectures to understand the geometric information and global perception inherent in SVG data. SketchRNN [Ha and Eck 2017] leverages a recurrent neural network (RNN) to generate vector paths for sketches. SVG-VAE [Lopes et al. 2019] uses an image autoencoder to capture font style features and employs LSTMs to generate vector fonts. Sketchformer [Ribeiro et al. 2020] uses a Transformer to recover sketch strokes in a vector form based on raster images. To maintain hierarchical relationships in SVGs, DeepSVG [Carlier et al. 2020] employs a hierarchical Transformer-based network to generate vector icons composed of multiple paths. Dual-modality learning framework [Liu et al. 2023; Wang and Lian 2021; Wang et al. 2023b] utilizes both sequence and image features to synthesize accurate vector fonts. A Transformer-based framework is designed to vectorize line drawings with dual-modal supervision [Liu et al. 2022], but it lacks smooth interpolation characteristics in the latent space. While these techniques have not yet supported text-guided SVG creation, IconShop [Wu et al. 2023] achieves T2V generation by representing SVGs as token sequences combined with text tokens.

Previous works focus on learning SVG-level latent representations from SVG command sequences. However, the vast diversity of path combinations poses a challenge in learning a universal global SVGlevel representation. As a result, existing methods predominantly generate SVGs in specific categories like fonts [Lopes et al. 2019; Wang and Lian 2021; Wang et al. 2023b], sketches [Ha and Eck 2017; Ribeiro et al. 2020], and simple icons [Cao et al. 2023; Carlier et al. 2020; Wu et al. 2023]. In contrast, our neural path representation captures valid path properties, while enabling the generation of diverse and complex vector graphics from text prompts.

3 OVERVIEW

Given a text prompt, our goal is to generate an SVG that aligns with the semantics of the text prompt while exhibiting desirable path properties and layer-wise structures consistent with human perception. Since an SVG is composed of a set of paths, denoted as $SVG = \{Path_1, Path_2, ..., Path_m\}$, our objective is to optimize a set of *m* paths based on a text prompt *T* through:

Neural Path Representation Learning (Sec. 4). The path geometry consists of connected cubic Bézier curves. Our objective is to learn a neural path representation by mapping each path into a latent code denoted as *z*, which captures valid geometric properties. To achieve this, we propose a dual-branch VAE that learns from both the image and sequence modalities of paths (refer to Fig. 2 (a)).

Text-driven Neural Path Optimization (Sec. 5). With the learned neural path representation, an SVG can be represented by a set of latent codes, along with color and transformation parameters associated with each path, denoted as $SVG = \{\theta_1, \theta_2, ..., \theta_m\}$, where $\theta_i = (z_i, C_i, Tr_i)$. Here, z_i, C_i , and Tr_i represent the latent code, color parameter, and transformation parameter for the *i*-path (*Path_i*), respectively. Next, we use the text prompt *T* to optimize the combination of *m* paths through two stages. In the first stage, we utilize a pre-trained diffusion model as a prior to optimize the combination of neural paths aligned with *T*. In the second stage, we apply a layerwise vectorization strategy to optimize path hierarchy, ensuring clear visual elements and layer structures in the generated SVG. Finally, after decoding the optimized latent codes into Bézier curves, we obtain the resulting SVG (refer to Fig. 2 (b)).

4 NEURAL PATH REPRESENTATION LEARNING

Previous T2V approaches directly optimize the control points of parametric paths. However, this often leads to intersecting or jagged paths due to the high degrees of freedom and the lack of geometric constraints. To address this issue, we propose a novel approach that learns a neural path representation within a latent space, capturing the valid geometric properties of paths and enabling the optimization of paths while ensuring geometric regularity. For this purpose, we design and train a dual-branch VAE to learn a latent space for path representation.

4.1 Dual-branch VAE

A parametric path (ignoring the color parameter) can be defined as a series of cubic Bézier curves connected end-to-end, denoted as $Path = (p_1, p_2, ..., p_k)$, where p_j for j = 1 to k represents the k control points used to define the cubic Bézier curves. By utilizing the differentiable rasterizer \mathcal{R} [Li et al. 2020], we can render the vector path and obtain the rasterized path image $I = \mathcal{R}(Path)$. The sequence representation of a path contains rich geometric information, such as the ordering and connections between points in the path. Training a VAE based on the point sequence enables effective learning of geometric properties, but it struggles to precisely reconstruct the rendered shape. On the other hand, the image representation is better at capturing visual features after rendering, but it cannot represent the relationships between control points, as different sequences of control points may result in the same rendered shape. To address these challenges, we propose a dual-branch encoder-decoder VAE that learns a shared latent space from path data in both vector and image modalities, as shown in Fig. 2 (a). This approach allows for the incorporation of both geometric and visual information, enabling a more comprehensive and accurate learning of paths.

Encoders. The encoder of our dual-branch VAE is composed of a sequence encoder and an image encoder. The **sequence encoder** takes the control point sequence as input and employs a transformer architecture with an attention mechanism to exploit the geometric relationships between control points. First, we normalize the control points sequence to the range of [0, 1] and then use a learnable embedding layer to project each normalized control point into a d_h -dimensional vector. Similar to DeepSVG [Carlier et al. 2020], we



Fig. 3. Training loss of the dual-branch VAE.

use positional encoding to embed the index of each point in the sequence. The sequence embedding $e_{seq} \in \mathbb{R}^{d_h \times k}$ is then fed into the six transformer layers. Each layer consists of masked multi-head attention and feed-forward layers [Vaswani et al. 2017]. Finally, a linear projection is applied to obtain the aggregated output sequence features $z_{seq} \in \mathbb{R}^{d_s}$. To further learn the shape perception of the path, we also adopt an **image encoder** composed of six convolutional layers. It takes the rasterized path image *I* as input and outputs the feature $z_{img} \in \mathbb{R}^{d_I}$. In our implementation, we set $d_h = 64$, $d_S = 32$, and $d_I = 64$.

Modality Fusion. We fuse the sequence and image features to create a comprehensive representation of paths. Specifically, the sequence features z_{seq} and image features z_{img} are concatenated and passed through a linear projection layer, obtaining a latent code $z \in \mathbb{R}^{d_F}$, where $d_F = 24$, that is shared with both modalities.

Decoder. The latent code *z* is passed through the two decoding branches to reconstruct the vector path and path image, respectively. The **sequence decoder** has a similar Transformer-based architecture to the sequence encoder. It takes *z* as input and outputs the decoded points sequence $P\hat{a}th = (\hat{p_1}, \hat{p_2}, \dots, \hat{p_k})$. The **image decoder** is a deconvolutional neural network that utilizes *z* to generate the reconstructed path image \hat{l} .

4.2 Training

Loss Function. The dual-branch VAE is trained end-to-end to reconstruct the input using dual-modality losses. However, simply using mean squared error (MSE) loss between the input and output control point sequences for sequence reconstruction can lead to overfitting, as different control point sequences can produce similar shapes. To address this, we introduce a shape-level loss to capture shape features. As depicted in Fig. 3, we first sample *n* auxiliary points (in our implementation, n = 4) from each cubic Bézier curve of a vector path, resulting in the auxiliary points set P_{aux} . Next, we calculate the Chamfer Distance between the auxiliary points set of the input path *Path* and that of the reconstructed path *Pâth* by summing the distances between the nearest correspondences of the two point sets. This calculation yields the Chamfer loss $\mathcal{L}_{cfr}(Path, Path)$. The



Fig. 4. Path samples and latent interpolations. (a) Path examples from the dataset. (b) Path samples decoded from random vectors in the latent space. (c) Latent interpolations among four given samples (marked in red).

image-level loss combines the L_1 loss and perceptual loss between the input image I and reconstructed images \hat{I} . Specifically, it can be expressed as $\mathcal{L}_{img} = |I - \hat{I}|_1 + \mathcal{L}_{percep}(I, \hat{I})$. In addition to the reconstruction losses in both modalities, the latent space is also regularized by the KL divergence loss \mathcal{L}_{kI} , which encourages the latent code z to follow a Gaussian distribution $\mathcal{N}(0, I)$.

The overall loss function is:

$$\mathcal{L} = \lambda_{cfr} \mathcal{L}_{cfr} + \lambda_{img} \mathcal{L}_{img} + \lambda_{kl} \mathcal{L}_{kl}, \tag{1}$$

where λ_{cfr} , λ_{img} , λ_{kl} are the weights of each loss term. We set $\lambda_{cfr} = 1$, $\lambda_{img} = 0.1$, $\lambda_{kl} = 0.01$ in our implementation.

Training Details. We train our dual-branch VAE on the FIGR-8-SVG dataset [Clouâtre and Demers 2019], which consists of blackand-white vector icons. To preprocess the data, we follow the same steps as IconShop [Wu et al. 2023] to obtain valid SVG data. We extract paths from the SVGs and remove duplicate shapes. In Fig. 4, we showcase several examples of paths from the dataset. For the raw input, since the control point sequences in each path can have variable lengths, we pad the point sequences with zeros to a fixed length (in our implementation, k = 50) and filter out those with longer lengths. This results in 200,000 samples for model training. The image resolution for training VAE is 64×64 . To train the model, we use the Adam optimizer with an initial learning rate of 0.001. We incorporate linear warmup and decay techniques. The dropout rate in all transformer layers is set to 0.1. We train the dual-branch network for 100 epochs.

Upon completion of training, we obtain a smooth latent space shared by both modalities. In Fig. 4, we present paths decoded from random samples in the learned latent space, along with smooth path interpolations.

5 TEXT-DRIVEN NEURAL PATH OPTIMIZATION

With the learned neural path representation, an SVG can be represented by a combination of *m* paths in the latent space, denoted as $SVG = \{\theta_1, \theta_2, ..., \theta_m\}$, where $\theta_i = (z_i, C_i, Tr_i)$. Here, the latent code z_i defines the shape, C_i defines the color, and Tr_i defines the affine transformation of the *path_i*. We initialize the latent codes by randomly drawing from a zero-mean Gaussian distribution. Our aim is to optimize these parameters based on the given text prompt *T*. Unlike previous methods [Jain et al. 2022; Xing et al. 2023b] that optimize control points explicitly, we optimize the latent code *z* within the learned space to guarantee the regularity of each path in the result. After optimization, we can obtain control points of a path by decoding z_i with our sequence decoder and transforming points with Tr_i , as $Path_i = Tr_i \cdot SeqDec(z_i)$.

Moreover, previous works [Frans et al. 2022; Jain et al. 2022] simultaneously optimize all paths of an SVG in a single stage, resulting in cluttered paths and a messy layer structure. To tackle this challenge, we have developed a two-stage neural path optimization process, as shown in Fig. 2 (b). In the first stage, we employ VSD defined on a pre-trained text-to-image diffusion model to optimize paths, resulting in an initial SVG that aligns with *T*. In the second stage, starting from the initial SVG, we apply a layer-wise vectorization strategy to hierarchically optimize path combinations, ensuring clear visual elements and well-defined layer-wise structures in the result SVG. Next, we introduce these two stages.

5.1 Optimization with Variational Score Distillation

In this stage, we optimize an initial SVG guided by the text prompt *T*. We draw inspiration from VectorFusion [Jain et al. 2022], which leverages a pre-trained text-to-image diffusion model as a prior to optimize path parameters through a Score Distillation Sampling (SDS) process. The process begins with an *SVG* with randomly initialized path parameters θ , and in each iteration, the *SVG* is rendered using the differentiable rasterizer \mathcal{R} to obtain a raster image $I_{SVG} = \mathcal{R}(SVG)$. The pre-trained encoder of the diffusion model encodes I_{SVG} into latent features $\mathbf{x} = \mathcal{E}(I_{SVG})$, and a Gaussian noise $\epsilon \in \mathcal{N}(0, I)$ is added to \mathbf{x} , obtaining \mathbf{x}_t at time t of the forward diffusion process. Finally, the SDS loss is defined as the distance between the added noise ϵ and the predicted noise ϵ_{ϕ} using the pre-trained diffusion model, and its gradient to optimize θ can be estimated as follows:

$$\nabla_{\theta} \mathcal{L}_{\text{SDS}} \triangleq \mathbb{E}_{t,\epsilon} \left[w(t) \left(\epsilon_{\phi} \left(\mathbf{x}_{t}; T, t \right) - \epsilon \right) \frac{\partial \mathbf{x}}{\partial \theta} \right]$$
(2)

where w(t) is a time-dependent weighting function.

Despite its success, empirical observations have revealed that results optimized from SDS suffer from issues such as over-saturation, over-smoothing, and a lack of diversity. These issues stem from SDS treating parameter θ as a single point and using a single point to approximate a distribution output by the diffusion model. In light of this, we leverage the VSD loss proposed by ProlificDreamer [Wang et al. 2023a] to replace SDS in our optimization. Unlike SDS, VSD models the parameter θ as a distribution, and consequently, the images rendered by *SVG* with parameter θ are also a distribution. Following ProlificDreamer, we employ LoRA (Low-rank adaptation) [Hu et al. 2021] of the pre-trained diffusion model to model the distribution of the rendered images. Therefore, the VSD loss is defined as the distance between the noises predicted by the pre-trained diffusion model and the LoRA model. Its gradient to optimize θ can be formulated as follows:

$$\nabla_{\theta} \mathcal{L}_{\text{VSD}} \triangleq \mathbb{E}_{t,\epsilon} \left[w(t) \left(\epsilon_{\phi} \left(\mathbf{x}_{t}; T, t \right) - \epsilon_{\text{LoRA}} \left(\mathbf{x}_{t}; T, t \right) \right) \frac{\partial \mathbf{x}}{\partial \theta} \right] \quad (3)$$

The use of VSD helps generate SVGs with higher quality and diversity. We denote the SVG optimized from this stage as SVG^0 , which will guide the layer-wise refinement in the next stage.

5.2 Optimization with Layer-wise Vectorization

Although VSD optimization is effective in aligning SVGs with text prompts, it often results in cluttered and stacked paths. This can introduce artifacts and lead to a poorly organized layer structure within the SVG, complicating subsequent editing and modification. To enhance the clarity of vector elements and the hierarchical structure of the generated SVG, we introduce a second-stage optimization based on SVG^0 obtained from the previous stage. This stage incorporates a path simplification step to obtain a simplified path set, and a layer-wise vectorization strategy to hierarchically optimize path combinations with the assistance of a guidance image. This stage can improve the overall quality and organization of the resulting SVG, making it easier to edit and reuse.

Path Simplification. We remove paths in SVG^0 with an opacity below 0.05 or an area smaller than 10 pixels, and merge overlapping paths that exhibit the same color into a single path using an overlap threshold of 5 pixels. These steps reduce the number of paths from *m* to *n*. The simplified path set with sequence representation is then inversely transformed into the latent space using our sequence encoder, obtaining a new set of latent codes and resulting in a new SVG, SVG^0 .

Layer-wise Optimization Strategy. After simplification, we further refine the $S\tilde{VG}^0$ with enhanced perceptual clarity and a better layerwise structure with the assistance of a guidance image. The guidance image is obtained by first rendering $S\tilde{VG}^0$ into the image I_{SVG}^0 . Then, we perturb I_{SVG}^0 with Gaussian noise, setting the noise strength to 0.4, and gradually remove the noise using a pre-trained diffusion model. This process helps reduce artifacts in the initial SVG and yields a guidance image I_g with clearer and more precise visual elements.

With the help of I_g , we introduce a layer-wise optimization strategy. Specifically, we sort the paths in SVG^0 by area and then optimize the top k paths with the largest areas in each iteration. A recursive pipeline progressively adds paths according to a path number schedule, thus optimizing the SVG from coarse to fine, as shown in Fig. 5.

In each iteration, we optimize the parameters of the top k paths in $\tilde{\theta^0}$, denoted as $SVG^k = \{\theta_1, \theta_2, ..., \theta_k\}$. The optimization is based on CLIP loss and IoU loss defined between SVG^k and the guidance image I_g . The CLIP loss is computed by summing the L_2 distances between the intermediate-level activations of the CLIP model as follows:

$$L_{CLIP} = \sum_{l} \|CLIP_{l}(I_{g}) - CLIP_{l}(\mathcal{R}(SVG^{k}))\|_{2}^{2},$$
(4)

 $CLIP_I$ denotes the CLIP encoder activation at layer l. It evaluates the image-level similarity and encourages the rendering of SVG^k to be faithful to the guidance image I_g . Moreover, to encourage a limited number of paths to cover the content of the guidance image as much as possible, we apply the intersection-over-union (IoU) loss between the rendered silhouette of SVG^k and the foreground region of I_g . The overall loss function is defined as $L_{lyr} = L_{CLIP} + \lambda_{IoU}L_{IoU}$, where λ_{IoU} is set to 0.01.



Fig. 5. Layer-wise optimization strategy.

As the iterations progress, more paths are involved in the optimization process, and the SVG gradually adds details. When all paths are optimized, we obtain the final result, SVG^n with *n* paths.

6 EXPERIMENTS

Experiment Setup. To evaluate our method, we collect 160 text prompts from the Stable Diffusion Prompts dataset [Dehouche and Kullathida 2023], including a diverse range of characters, actions, and scenes. For each prompt, we generate 5 SVGs, culminating in a total of 800 vector graphics. To encourage the generation towards a flat vector style, we append the phrase "minimal flat 2d vector" to each prompt. In VSD optimization process, we utilize the official "SD-v1-5" checkpoint¹ with a guidance scale of 10, and timestep $t \sim \mathcal{U}(50, 950)$ is uniformly sampled.

6.1 Evaluation Metrics

We evaluate the quality of our results from vector-level, image-level, and text-level perspectives.

Vector-level. Drawing upon criteria from prior perception research in vector graphics [Dominici et al. 2020; Favreau et al. 2017], a good SVG should minimize redundant paths to preserve compactness and editability. Based on this guideline, we assess vectorization quality using the following metrics: (a) **Smoothness**: Inverse of the average curvature variation of the paths in generated SVGs. (b) **Simplicity**: The average number of paths in generated vector graphics. (c) **Layer-wise semantics**: The semantics of paths are evaluated by comparing the decrease in CLIP similarity between the SVG render results and corresponding text prompts before and after randomly removing 30% of the SVG paths. A larger decrease indicates that each path has more specific semantics.

Image-level. To evaluate the visual quality and diversity of vector graphics, we collected a dataset of 800 well-designed vector graphics from *iconfont*², encompassing various categories including characters, animals, and scenes. These samples serve as the ground truth for calculating the FID metrics on rendered images of SVGs.

¹https://huggingface.co/runwayml/stable-diffusion-v1-5 ²https://iconfont.cn

ACM Trans. Graph., Vol. 43, No. 4, Article 36. Publication date: July 2024.

Table 1. Quantitative comparison with existing methods.

Methods	Smooth ↑	Simp \downarrow	Layer ↑	$FID \downarrow$	CLIP ↑
T2I + Potrace	0.7846	430	0.0226	104.92	0.2732
T2I + LIVE	0.5797	40	0.0119	97.82	0.2519
T2I + LoRA + Potrace	0.7882	160	0.0395	45.16	0.2729
CLIPDraw	0.4711	128	0.0317	123.55	0.2832
Diffsketcher	0.4829	128	0.0105	92.36	0.2607
Vectorfusion (from scratch)	0.6322	128	0.0139	65.71	0.2675
Vectorfusion (from LIVE)	0.5025	128	0.0258	76.45	0.2880
Ours	0.8012	40	0.0591	52.30	0.3015

Text-level. To compute whether the generated SVG is aligned with the input text prompt, we define the text-level similarity by calculating the CLIP cosine similarity [Radford et al. 2021] between the text prompt and rendered SVG.

6.2 Baselines

We compare our proposed pipeline with two types of T2V pipelines.

Vectorization with T2I. We first generate images from text prompts using a diffusion model and then convert the images into SVGs using two distinct vectorization methods: (a) **Potrace** [Selinger 2003]: A traditional vectorization method involves image segmentation and curve fitting to transform raster images into SVGs. (b) **LIVE** [Ma et al. 2022]: A deep learning method that generates SVGs by optimizing paths using loss functions defined in the image space. We use the same number of paths as our method.

Text-guided SVG optimization. We compare our method with CLIP-based and diffusion-based optimization approaches: (c) **CLIP-Draw** [Frans et al. 2022]: A method that leverages the image-text similarity metric of CLIP to optimize SVGs from prompts. (d) **Diffsketcher** [Xing et al. 2023a]: A method incorporates both image-level CLIP loss and SDS loss for text-guided sketch generation. (e) **VectorFusion** [Jain et al. 2022]: This approach employs SDS loss to optimize SVGs to align with given text prompts. It offers two optimization routes: starting from scratch or refining vectorized results produced by LIVE. We use 128 paths for text-guided SVG optimization methods, consistent with the default setting in VectorFusion.

6.3 Comparisons

We evaluate the performance of our method by comparing it with baselines qualitatively and quantitatively. The quantitative results are provided in Tab. 1 and the qualitative results are shown in Fig. 6 and Fig. 17. As shown in Tab. 1, our method outperforms other approaches from a comprehensive perspective.

Comparisons with Vectorization with T2I Methods. Though with a vector style prompt suffix, T2I models still often generate raster images with intricate textures and complex color variations. It poses a challenge when converting to vector graphics with smooth geometric shapes and uniform colors. As shown in Fig. 6, the vectorization results using Potrace exhibit overly complex vector elements. This leads to a high number of paths (indicated by a high simplicity score in Tab. 1) and a lack of layer organization among these paths, resulting in diminished path semantics (reflected by a low layer-wise semantics score in Tab. 1). Furthermore, such vectorization results



Fig. 6. Qualitative comparison to vectorization with T2I methods.

often deviate from the style of well-designed SVGs, as evidenced by the low FID scores in Tab. 1. LIVE faces similar challenges. When the path number is set the same as our method, LIVE obtains suboptimal vectorization outcomes as the constrained path number is often insufficient for accurately reconstructing the images generated by the T2I model. Furthermore, the SVGs produced by LIVE exhibit numerous irregular and broken paths, as indicated by the low smoothness score in Tab. 1.

LoRA [Hu et al. 2021] is a popular technique to fine-tune a pretrained T2I model for specific styles. To produce SVG-style images, we utilize the "Vector Illustration" LoRA³ fine-tuned on the base diffusion model "SD-v1-5", and then convert the generated images to vector graphics using Potrace. Since this LoRA model is specifically fine-tuned on well-designed vector graphics images similar to the iconfont dataset we used for evaluation, the generated results exhibit lower FID scores. However, as demonstrated in Fig. 6 (c), while this baseline achieves pleasing visual results, it still struggles with overcomplicated and disorganized paths, as this is a common issue in traditional image vectorization methods. In contrast, our methods produce visually appealing SVGs with smooth paths and layer-wise structures.

Comparisons with Text-guided SVG Optimization Methods. These methods directly optimize the control points of parametric paths. However, due to their high degrees of freedom and the lack of geometry constraints, the control points often undergo complex transformations to generate SVGs that align with the text prompts. This results in low path smoothness, as indicated by the lower Smooth score in Tab. 1. Zoom-in illustrations in Fig. 17 highlight the issues of intersecting and jagged paths, leading to visually unappealing outcomes. Moreover, the resulting SVGs often contain complex and redundant shapes, making them difficult to edit. In contrast, our neural path representation effectively captures valid geometric

³https://civitai.com/models/60132/vector-illustration

ACM Trans. Graph., Vol. 43, No. 4, Article 36. Publication date: July 2024.

36:10 • Peiying Zhang, Nanxuan Zhao, and Jing Liao



Fig. 13. Text-guided SVG customization. Exemplar SVGs: the 1^{st} row is from Envato Elements creator ©Telllu; the 2^{nd} row is from ©Freepik.



Fig. 14. Image-to-SVG generation.

on an exemplar image (rendered from exemplar SVG) and a text prompt containing a special token V*. This token learns the concept from the exemplar image. Subsequently, we apply our method to optimize the SVG based on a new prompt, such as "V* holding a laptop" or "Trees in V* style," resulting in a customized SVG that reflects the desired customization in object or style. Fig. 13 showcases the results of customizing the exemplar SVG with different text prompts, demonstrating the flexibility and creativity of our method.

7.4 Image-to-SVG Generation

Our framework enables flexible control beyond text prompts, which is particularly useful for designers seeking inspiration from natural images for SVG-style designs. For example, as shown in Fig. 14, our method can generate vector icons from natural images. This is achieved by integrating the ControlNet [Zhang et al. 2023b] into the VSD optimization process, which ensures an optimization direction



Fig. 15. SVG animation aligned with the described motions.



Fig. 16. Failure cases.

that respects both the original structure of the image and the input text prompt.

7.5 SVG Animation

Our framework can be extended to SVG animation by animating an initial SVG according to a text prompt describing the desired motion. We first generate a static SVG using our pipeline and then animate it by optimizing a temporal sequence of *m* paths into *k* video frames, represented as *Video* = {*SVG*¹, *SVG*², ..., *SVG*^k} = {*Path*_i^j}_{i \in m}^{j,k}, aligning with the motion described in the text prompt. Specifically, we employ a similar two-stage optimization process. In the VSD optimization stage, we leverage a pre-trained text-to-video diffusion model ModelScope [Wang et al. 2023c] to replace the text-to-image model. In the layer-wise vectorization stage, we generate a guidance video based on the initial SVG sequence [Geyer et al. 2023], which

video based on the initial SVG sequence [Geyer et al. 2023], which contains k guidance images. As demonstrated in Fig. 15, our method animates SVGs with smooth motions, showcasing the effectiveness of our neural path representation.

8 CONCLUSION

In this paper, we propose a novel text-to-vector pipeline to generate vector graphics that align with the semantics of given text prompts. Our framework learns a neural path representation within the latent space to capture valid geometric properties of paths. By employing the two-stage text-driven neural path optimization, our method effectively generates vector graphics with desirable path properties and layer-wise structures.

While our method achieves high-quality SVG results, it still has some limitations, as shown in Fig. 16. First, our method relies on the generative capabilities of the diffusion model; thus, overly detailed text prompts may lead to inaccurate SVG results. For example, as depicted in Fig. 16 (a), semantic elements like the "ping pong table" and "ashtray" are missing. Second, our method tends to simplify shapes with intricate boundaries into smoother paths, as they exceed the representational capacity of our path latent space. For instance, the detailed edges of the bear's body in Fig. 16 (b) are smoothed out, resulting in the loss of the original complexity. This can be improved by collecting a larger path dataset containing more complex paths, which we leave as future work. Third, our method, similar to other text-guided SVG optimization methods such as CLIPDraw and VectorFusion, is generally slow due to the iterative optimization process. It takes approximately 13 minutes to optimize 128 paths on an NVIDIA-3090, whereas vectorization with T2I methods can be completed within a few seconds. Despite the current slowness, our neural path representation lays the groundwork for training fast feed-forward T2V networks to replace iterative optimization in future work. This approach could also bring benefits to the generation of graphic layouts, fonts, and CAD models.

ACKNOWLEDGMENTS

The work described in this paper was substantially supported by a GRF grant from the Research Grants Council (RGC) of the Hong Kong Special Administrative Region, China [Project No. CityU 11216122].

REFERENCES

- Defu Cao, Zhaowen Wang, Jose Echevarria, and Yan Liu. 2023. SVGformer: Representation Learning for Continuous Vector Graphics using Transformers. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 10093–10102.
- Alexandre Carlier, Martin Danelljan, Alexandre Alahi, and Radu Timofte. 2020. Deepsvg: A hierarchical generative network for vector graphics animation. Advances in Neural Information Processing Systems 33 (2020), 16351–16361.
- Ye Chen, Bingbing Ni, Xuanhong Chen, and Zhangli Hu. 2023. Editable Image Geometric Abstraction via Neural Primitive Assembly. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 23514–23523.
- Louis Clouâtre and Marc Demers. 2019. Figr: Few-shot image generation with reptile. arXiv preprint arXiv:1901.02199 (2019).
- Nassim Dehouche and Kullathida. 2023. What is in a Text-to-Image Prompt: The Potential of Stable Diffusion in Visual Arts Education. arXiv preprint arXiv:2301.01902 (2023).
- Edoardo Alberto Dominici, Nico Schertler, Jonathan Griffin, Shayan Hoshyari, Leonid Sigal, and Alla Sheffer. 2020. Polyfit: Perception-aligned vectorization of raster clip-art via intermediate polygonal fitting. ACM Transactions on Graphics (TOG) 39, 4 (2020), 77–1.
- Jean-Dominique Favreau, Florent Lafarge, and Adrien Bousseau. 2017. Photo2clipart: Image abstraction and vectorization using layered linear gradients. ACM Transactions on Graphics (TOG) 36, 6 (2017), 1–11.
- Kevin Frans, Lisa Soros, and Olaf Witkowski. 2022. Clipdraw: Exploring text-to-drawing synthesis through language-image encoders. Advances in Neural Information Processing Systems 35 (2022), 5207–5218.
- Rinon Gal, Yael Vinker, Yuval Alaluf, Amit H Bermano, Daniel Cohen-Or, Ariel Shamir, and Gal Chechik. 2023. Breathing Life Into Sketches Using Text-to-Video Priors. arXiv preprint arXiv:2311.13608 (2023).
- Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. 2023. Tokenflow: Consistent diffusion features for consistent video editing. arXiv preprint arXiv:2307.10373 (2023).

David Ha and Douglas Eck. 2017. A neural representation of sketch drawings. arXiv preprint arXiv:1704.03477 (2017).

- Amir Hertz, Kfir Aberman, and Daniel Cohen-Or. 2023. Delta denoising score. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 2328– 2337.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. Advances in neural information processing systems 33 (2020), 6840–6851.
- Shayan Hoshyari, Edoardo Alberto Dominici, Alla Sheffer, Nathan Carr, Zhaowen Wang, Duygu Ceylan, and I-Chao Shen. 2018. Perception-driven semi-structured boundary vectorization. ACM Transactions on Graphics (TOG) 37, 4 (2018), 1–14.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language

models. arXiv preprint arXiv:2106.09685 (2021).

- Adobe Illustrator. 2023. Turn ideas into illustrations with Text to Vector Graphic. https://www.adobe.com/products/illustrator/text-to-vector-graphic.html.
- Illustroke. 2024. Stunning vector illustrations from text prompts. https://illustroke.com/. Shir Iluz, Yael Vinker, Amir Hertz, Daniel Berio, Daniel Cohen-Or, and Ariel Shamir. 2023. Word-as-image for semantic typography. arXiv preprint arXiv:2303.01818 (2023).
- Ajay Jain, Amber Xie, and Pieter Abbeel. 2022. VectorFusion: Text-to-SVG by Abstracting Pixel-Based Diffusion Models. arXiv preprint arXiv:2211.11319 (2022).
- Kittl. 2024. AI Vector Generator. https://www.kittl.com/feature/ai-text-to-vector.
- Johannes Kopf and Dani Lischinski. 2011. Depixelizing pixel art. In ACM SIGGRAPH 2011 papers. 1–8.
- Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. 2022. Multi-Concept Customization of Text-to-Image Diffusion. arXiv preprint arXiv:2212.04488 (2022).
- Tzu-Mao Li, Michal Lukáč, Michaël Gharbi, and Jonathan Ragan-Kelley. 2020. Differentiable vector graphics rasterization for editing and learning. ACM Transactions on Graphics (TOG) 39, 6 (2020), 1–15.
- Hanyuan Liu, Chengze Li, Xueting Liu, and Tien-Tsin Wong. 2022. End-to-end line drawing vectorization. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 36. 4559–4566.
- Ying-Tian Liu, Zhifei Zhang, Yuan-Chen Guo, Matthew Fisher, Zhaowen Wang, and Song-Hai Zhang. 2023. DualVector: Unsupervised Vector Font Synthesis with Dual-Part Representation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 14193–14202.
- Raphael Gontijo Lopes, David Ha, Douglas Eck, and Jonathon Shlens. 2019. A learned representation for scalable vector graphics. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 7930–7939.
- Xu Ma, Yuqian Zhou, Xingqian Xu, Bin Sun, Valerii Filev, Nikita Orlov, Yun Fu, and Humphrey Shi. 2022. Towards layer-wise image vectorization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 16314–16323.
- Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. 2021. Sdedit: Guided image synthesis and editing with stochastic differential equations. arXiv preprint arXiv:2108.01073 (2021).
- Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021).
- Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. 2022. Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988 (2022).
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- Pradyumna Reddy, Michael Gharbi, Michal Lukac, and Niloy J Mitra. 2021. Im2vec: Synthesizing vector graphics without vector supervision. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 7342–7351.
- Leo Sampaio Ferraz Ribeiro, Tu Bui, John Collomosse, and Moacir Ponti. 2020. Sketchformer: Transformer-based representation for sketched structure. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 14153–14162.
- Juan A Rodriguez, Shubham Agarwal, Issam H Laradji, Pau Rodriguez, David Vazquez, Christopher Pal, and Marco Pedersoli. 2023. StarVector: Generating Scalable Vector Graphics Code from Images. arXiv preprint arXiv:2312.11556 (2023).
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 10684–10695.
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2022. Dreambooth: Fine tuning text-to-image diffusion models for subjectdriven generation. arXiv preprint arXiv:2208.12242 (2022).
- Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. 2022. Palette: Image-to-image diffusion models. In ACM SIGGRAPH 2022 Conference Proceedings. 1–10.
- Peter Schaldenbrand, Zhixuan Liu, and Jean Oh. 2022. Styleclipdraw: Coupling content and style in text-to-drawing translation. *arXiv preprint arXiv:2202.12362* (2022).

Peter Selinger. 2003. Potrace: a polygon-based tracing algorithm.

- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In International conference on machine learning. PMLR, 2256–2265.
- Yiren Song, Xning Shao, Kang Chen, Weidong Zhang, Minzhe Li, and Zhongliang Jing. 2022. CLIPVG: Text-Guided Image Manipulation Using Differentiable Vector Graphics. arXiv preprint arXiv:2212.02122 (2022).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. Advances in neural information processing systems 30 (2017).
- Yael Vinker, Ehsan Pajouheshgar, Jessica Y Bo, Roman Christian Bachmann, Amit Haim Bermano, Daniel Cohen-Or, Amir Zamir, and Ariel Shamir. 2022. Clipasso:

ACM Trans. Graph., Vol. 43, No. 4, Article 36. Publication date: July 2024.

36:12 • Peiying Zhang, Nanxuan Zhao, and Jing Liao

Semantically-aware object sketching. ACM Transactions on Graphics (TOG) 41, 4 (2022), 1-11.

- Andrey Voynov, Kfir Aberman, and Daniel Cohen-Or. 2023. Sketch-guided text-to-image diffusion models. In ACM SIGGRAPH 2023 Conference Proceedings. 1–11.
- Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang. 2023c. Modelscope text-to-video technical report. arXiv preprint arXiv:2308.06571 (2023).
- Yizhi Wang and Zhouhui Lian. 2021. DeepVecFont: Synthesizing high-quality vector fonts via dual-modality learning. ACM Transactions on Graphics (TOG) 40, 6 (2021), 1–15.
- Yuqing Wang, Yizhi Wang, Longhui Yu, Yuesheng Zhu, and Zhouhui Lian. 2023b. DeepVecFont-v2: Exploiting Transformers to Synthesize Vector Fonts with Higher Quality. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 18320–18328.
- Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. 2023a. ProlificDreamer: High-Fidelity and Diverse Text-to-3D Generation with Variational Score Distillation. arXiv preprint arXiv:2305.16213 (2023).
- Ronghuan Wu, Wanchao Su, Kede Ma, and Jing Liao. 2023. IconShop: Text-Guided Vector Icon Synthesis with Autoregressive Transformers. ACM Transactions on Graphics (TOG) 42, 6 (2023), 1–14.
- Ximing Xing, Chuang Wang, Haitao Zhou, Jing Zhang, Qian Yu, and Dong Xu. 2023a. DiffSketcher: Text Guided Vector Sketch Synthesis through Latent Diffusion Models.

arXiv preprint arXiv:2306.14685 (2023).

- Ximing Xing, Haitao Zhou, Chuang Wang, Jing Zhang, Dong Xu, and Qian Yu. 2023b. SVGDreamer: Text Guided SVG Generation with Diffusion Model. arXiv preprint arXiv:2312.16476 (2023).
- Chuan Yan, David Vanderhaeghe, and Yotam Gingold. 2020. A benchmark for rough sketch cleanup. ACM Transactions on Graphics (TOG) 39, 6 (2020), 1–14.
- Ming Yang, Hongyang Chao, Chi Zhang, Jun Guo, Lu Yuan, and Jian Sun. 2015. Effective clipart image vectorization through direct optimization of bezigons. *IEEE transactions on visualization and computer graphics* 22, 2 (2015), 1063–1075.
- Fangneng Zhan, Yingchen Yu, Rongliang Wu, Jiahui Zhang, Shijian Lu, Lingjie Liu, Adam Kortylewski, Christian Theobalt, and Eric Xing. 2023. Multimodal image synthesis and editing: A survey and taxonomy. *IEEE Transactions on Pattern Analysis* and Machine Intelligence (2023).
- Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023b. Adding conditional control to text-to-image diffusion models. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 3836–3847.
- Peiying Zhang, Nanxuan Zhao, and Jing Liao. 2023c. Text-Guided Vector Graphics Customization. In SIGGRAPH Asia 2023 Conference Papers. 1–11.
- Yuxin Zhang, Nisha Huang, Fan Tang, Haibin Huang, Chongyang Ma, Weiming Dong, and Changsheng Xu. 2023a. Inversion-based style transfer with diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 10146–10156.



Fig. 17. Qualitative comparison with text-guided SVG optimization methods.

ACM Trans. Graph., Vol. 43, No. 4, Article 36. Publication date: July 2024.