



(a) Exemplar SVG

(b) Text-Guided Vector Graphics Customization

Figure 1: Text-guided vector graphics customization results. Given an exemplar vector graphic in SVG format (a), our model can generate customized vector graphics (b) based on diverse text inputs while keep the visual identity (V^*) of the exemplar SVG. Exemplar SVGs: the 1st row is from ©iconfont; the 2nd row is from illustac creator ©Zarame.

ABSTRACT

Vector graphics are widely used in digital art and valued by designers for their scalability and layer-wise topological properties. However, the creation and editing of vector graphics necessitate creativity and design expertise, leading to a time-consuming process. In this paper, we propose a novel pipeline that generates high-quality customized vector graphics based on textual prompts while preserving the properties and layer-wise information of a given exemplar SVG. Our method harnesses the capabilities of large pre-trained text-to-image models. By fine-tuning the cross-attention layers of the model, we generate customized raster images guided by textual prompts. To initialize the SVG, we introduce a semantic-based path alignment method that preserves and transforms crucial paths from the exemplar SVG. Additionally, we optimize path parameters using both image-level and vector-level losses, ensuring smooth shape deformation while aligning with the customized raster image. We

*Corresponding author

SA Conference Papers '23, December 12–15, 2023, Sydney, NSW, Australia

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-0315-7/23/12...\$15.00 https://doi.org/10.1145/3610548.3618232 extensively evaluate our method using multiple metrics from vectorlevel, image-level, and text-level perspectives. The evaluation results demonstrate the effectiveness of our pipeline in generating diverse customizations of vector graphics with exceptional quality. The project page is https://intchous.github.io/SVGCustomization.

CCS CONCEPTS

• Computing methodologies → Graphics file formats; Image processing; • Applied computing → Arts and humanities.

KEYWORDS

Vector Graphics, SVG, Diffusion Model, Image Vectorization, Image Editing

ACM Reference Format:

Peiying Zhang, Nanxuan Zhao, and Jing Liao. 2023. Text-Guided Vector Graphics Customization. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23), December 12–15, 2023, Sydney, NSW, Australia.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3610548.3618232

1 INTRODUCTION

Vector graphics is a type of art composed of geometric shapes, such as lines, curves, and polygons. It plays essential roles under various practical scenarios (*e.g.*, animations, cartoons, and graphic designs) due to its scalability, precision, and versatility for creating high-quality visuals and designs [Quint 2003]. With the feature of being resolution-independent, vector graphics ensures consistent

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.

image quality across different media and devices, and it also enables easy manipulation, making it efficient to create intricate designs. However, creating high-quality vector graphics requires expertise and can be time-consuming [Diebel 2008]. Inspired by the recent success of text-to-image (T2I) generation [Ramesh et al. 2021; Rombach et al. 2022; Ruiz et al. 2022], one question raises whether the model can assist in automating the creation of vector graphics.

One approach for adapting from text-to-image to text-to-vector (T2V) is to first generate raster images from text and then convert them into vectors using existing image vectorization methods such as Potrace [Selinger 2003], Effective Vectorization [Yang et al. 2015], Photo2Clipart [Favreau et al. 2017], or Layer-Wise Vectorization [Ma et al. 2022]. However, image vectorization methods often generate undesired vector elements (e.g., irregular shapes or overcomplicated paths) and layer relationships, complicating further graphic manipulation. Recently, a new category of T2V methods (e.g., VectorFusion [Jain et al. 2022]) has emerged that directly optimizes the paths of vector graphics using some pre-trained vision-language models such as the CLIP model[Radford et al. 2021] or the Diffusion model[Rombach et al. 2022]. These methods can generate diverse vector graphics given the same text input. However, restricting the appearance of generated subjects and ensuring path regularity can be challenging with these methods.

To address the aforementioned issues in T2V, we propose a novel task called text-guided vector graphics customization. Following the previous works [Frans et al. 2022; Ma et al. 2022], we use Scalable Vector Graphics (SVGs) with a set of parametric shape primitives (e.g., cubic Bézier curves) to represent vector graphics. Our objective is to generate customized SVGs that maintain the visual identity, vector properties, and layer relationships of an exemplar SVG, based on a given text prompt. Such a method can benefit a lot of practical use cases, imagining that you have created a new character or logo (e.g., a cute bird), but want to adapt it into a different context (e.g., the cute bird reading a book) or generate more variations, as shown in Fig. 1. However, this task poses several unique challenges. First, the visual identity of the input SVG should be well preserved in the customized results, which is a difficult task given only a single SVG input. Second, to fit into diverse contexts, the geometry and appearance features of the customized SVG can be really different, how to maintain the correct path relationships and incorporate newly added contents/paths at the same time remains unsolved. Lastly, the method must generate visually pleasing results alongside valid vector paths, further increasing the complexity of the task.

To this end, we propose a pipeline for text-guided vector graphics customization. We leverage a powerful pre-trained diffusion model [Rombach et al. 2022] that has been proven to be efficient in multiple T2I generation and editing tasks. By fine-tuning the cross-attention layers of the pre-trained diffusion model with exemplar SVG and input text prompt, we first generate a customized raster image that preserves the visual identity of the exemplar. Inspired by experts' design process that reuses essential semantic elements across a series of SVGs to achieve coherence with the concept[Ambrose et al. 2019], we propose a semantic-based path alignment method that keeps important paths based on semantic correspondence between the exemplar and the customized raster image. We apply rigid transformations to these kept paths to provide a good initialization of the output customized SVG. Path parameters (e.g., control point positions and fill colors) of the SVG are further optimized with a set of image-level loss and vector-level loss to reconstruct the customized raster image while preserving path regularity. The image-level loss measures the distance between the rendered SVG and customized image in the CLIP latent space, promoting fidelity to the customized image. The vector-level loss incorporates a local Procrustes loss, enforcing path regularity by minimizing local deformation of control points. After optimization, a customized SVG corresponding to the text prompt and preserving the visual identity of the exemplar SVG is generated.

We assess our method using vector-level, image-level and textlevel metrics, confirming its efficacy in creating high-quality textguided vector graphics customizations. Our key contributions are:

- We propose the first pipeline for vector graphics customization that generates customized SVGs based on text prompts while preserving the visual identity of the exemplar SVG.
- We present a semantic-based path alignment method that allows for the reuse of paths in the exemplar SVG, thus better maintaining the valid path properties and layer relationships of the exemplar SVG in customized SVGs.
- We propose a path optimization method for image vectorization that uses both image-level loss and vector-level loss to reconstruct the target image while retaining the path regularity.

2 RELATED WORK

2.1 Text-to-Image Generation and Manipulation

Rapid progress of deep learning has driven advancements in T2I tasks. Early methods, such as text-conditional GANs [Li et al. 2019; Reed et al. 2016], offer end-to-end architectures but are limited to specific domains. Autoregressive models like DALL-E [Ramesh et al. 2021] use transformers and large-scale datasets for open-domain generation. Recently, diffusion models become state-of-the-art in T2I generation. GLIDE [Nichol et al. 2021] innovatively employs a conditional diffusion model to guide the noise diffusion process of images with text. By further introducing image latent features, Stable Diffusion [Rombach et al. 2022] overcomes resolution constraints and generates amazing high-res results. Diffusion models also support flexible text-guided image editing and customization. By embedding images and text into a unified latent space, diffusion models can generate customized images of specific concepts based on limited user-provided images and text prompts as guidance [Gal et al. 2022; Kumari et al. 2022; Ruiz et al. 2022]. However, these diffusion methods are primarily designed for image customization, and a further vectorization step is required to convert them into vector graphics.

2.2 Image Vectorization

Image vectorization aims at representing an image using parametric shape primitives, such as control points and color parameters of vector paths [Quint 2003]. Traditional vectorization methods follow a pipeline that involves segmenting images into regions based on color similarity [Kopf and Lischinski 2011] and fitting curves to the region boundaries [Favreau et al. 2017; Selinger 2003; Yang et al. 2015]. Furthermore, Perception [Hoshyari et al. 2018] and PolyFit [Dominici et al. 2020] refine the curve-fitting process

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia



Figure 2: Text-Guided vector graphics customization pipeline. Given an exemplar SVG and a text prompt as input (a), the method consists of three stages. (b) Concept Fine-tuning: By fine-tuning a pre-trained T2I model, a customized raster image is generated based on the text prompt. (c) Path Alignment: Important paths from the exemplar SVG are adapted based on semantic correspondences with the customized raster image, providing an initial customized SVG. (d) Path Optimization: The path parameters are optimized using both image-level and vector-level losses to produce the final customized SVG (e). The exemplar SVG is from envatoelements creator ©Telllu.

to better align with perceptions. However, these traditional vectorization approaches ignore the topological information inherent in raster images and are thus incapable of obtaining layer-wise paths in SVGs. The advent of differentiable rendering techniques [Li et al. 2020] has improved image vectorization by enabling the optimization of path parameters using loss functions defined in image space. LIVE [Ma et al. 2022] uncovers layer-wise topology by directly optimizing path parameters under the reconstruction supervision of the input image. SketchRNN [Ha and Eck 2017] combines differentiable rendering with a recurrent neural network (RNN) to predict vector paths in order. Subsequently, Sketchformer [Ribeiro et al. 2020] uses a transformer network to recover sketch strokes from raster images. Im2vec [Reddy et al. 2021] introduces a VAE to predict ordered vector paths from raster images. Despite their success, inadequate path initialization and optimization can result in messsy and redundant paths.

2.3 Text-to-Vector Generation

Besides combining T2I generation and image vectorization, another line of work for T2V generation is to directly optimize SVG paths guided by pre-trained vision-language models, such as the CLIP model [Radford et al. 2021] or the diffusion model [Rombach et al. 2022]. CLIPDraw [Frans et al. 2022] optimizes the image-text similarity metric in CLIP latent space to generate vector graphics from text prompts. StyleCLIPDraw [Schaldenbrand et al. 2022] augments CLIPDraw by introducing a style loss, enabling both artistic and textual control over the synthesized drawings. ES-CLIP [Tian and Ha 2022] adopts an evolution strategy to optimize triangle-based vector elements. To vectorize and edit the input image based on text prompts, CLIPVG [Song et al. 2022] progressively enhances SVG details by adding additional paths. Building upon the strong visual and semantic priors of pre-trained diffusion models, VectorFusion [Jain et al. 2022] optimizes an SVG consistent with a text prompt using Score Distillation Sampling (SDS) loss. However, challenges persist in applying these methods to vector graphics customization. Ensuring that the resulting SVG maintains the characteristics of the exemplar SVG, even with initialization, can be difficult due to

diverse optimization processes. Without regularization, resulting SVGs may have messy and stacked paths, complicating further editing and modification.

3 OVERVIEW

Given an exemplar SVG and a text prompt, our goal is to generate a customized SVG that conforms to the semantics of the given text prompt while preserving the visual identity of the exemplar SVG. Additionally, the customized SVG should maintain the valid path properties and layer relationships of the exemplar SVG. Since an SVG is composed of a set of paths, we define each path as a piecewise Bézier spline composed of several cubic Bézier curves connected end-to-end and filled with a uniform color *c*. Therefore, we can represent a path as $Path = (p_1, p_2, ..., p_d, c)$, where $\{p_j\}_{j=1}^d$ are *d* control points used to define the cubic Bézier curves. Given an exemplar SVG consisting of *n* parametric paths, denoted as $SVG_E = \{Path_1, Path_2, ..., Path_n\}$, and a text prompt *T*, our goal is to obtain a customized SVG composed of *m* optimized paths, denoted as $SVG_C = \{Path_1^*, Path_2^*, ..., Path_m^*\}$.

Our framework, as shown in Fig. 2, is inspired by experts' design process. When designing a series of SVGs for a specific concept, designers initially select essential semantic elements to reuse across SVGs, place them in appropriate regions, and progressively adjust the paths concerning their position, shape, color, and hierarchical order to generate a set of SVGs that share the same visual identity but have different contexts [Ambrose et al. 2019]. Inspired by this process, our text-guided vector graphics customization framework can be decomposed into three stages: concept fine-tuning, path alignment, and path optimization.

Concept Fine-tuning (Sec. 4.1). To avoid the need for training on large vector graphics datasets, our first step is to leverage the powerful visual and semantic priors induced by a pre-trained T2I generation model for image customization. Specifically, we use a differentiable rasterizer [Li et al. 2020], denoted as \mathcal{R} , to obtain a rasterized exemplar SVG, denoted as $I_E = \mathcal{R}(SVG_E)$. We then fine-tune the pre-trained Stable Diffusion model [Rombach et al. 2022] based on the exemplar image I_E to generate a customized raster image, denoted as I_C , that reflect the desired modifications indicated by the text prompt *T* (see Fig. 2 (b)).

Path Alignment (Sec. 4.2). To preserve path properties and layerwise information from the exemplar SVG, we use pre-trained DINO-ViT features [Caron et al. 2021] to establish semantic correspondences between the exemplar image I_E and the customized image I_C . We then transform reusable paths from SVG_E to their corresponding connected component in I_C , adjusting their positions and scales as necessary. For unmatched components in I_C , we add new paths. As a result, we obtain $SVG_C^0 = \{Path_1^0, Path_2^0, ..., Path_m^0\}$ as an initialization (see Fig. 2 (c)) for the customized SVG.

Path Optimization (Sec. 4.3). After initialization, we optimize the geometric and color parameters of paths based on both image-level loss and vector-level loss. The differentiable rasterizer connects the path parameters SVG_C and the customized image I_C . Path optimization refines the path parameters to more accurately fit I_C , while maintaining smooth path deformation. This process ultimately yields the customized SVG with optimized paths $SVG_C = \{Path_1^*, Path_2^*, ..., Path_m^*\}$ (see Fig. 2 (d)).

4 METHOD

4.1 Concept Fine-tuning

Instead of training a T2V model directly for SVG generation and editing, which requires a large dataset of vector graphics that may not be readily available, we utilize a pre-trained T2I generation model like Stable Diffusion [Rombach et al. 2022]. We first generate a customized image using this model and then proceed to vectorize it. However, since the Stable Diffusion model is solely trained on text prompts, it cannot maintain the visual identity of our exemplar SVG. To overcome this limitation, we introduce a Concept Fine-tuning approach, where the model is fine-tuned on the rasterized exemplar SVG. This fine-tuning enables the model to generate customized images that reflect the desired modifications specified by the text prompt while preserving the visual identity of the exemplar.

Given an exemplar image I_E , which is rendered from SVG_E , we aim to fine-tune the Stable Diffusion model to generate an image that resembles I_E under a specific text prompt denoted as 'A clipart of V^* '. V^* is initialized with a rarely occurring token embedding and optimized concurrently with the fine-tuning. However, finetuning all the weights of the Stable Diffusion model may lead to overfitting on I_E , resulting in reduced editing variations. Thus, we adopt an effective concept fine-tuning strategy that only fine-tunes the cross-attention blocks between text and image, following the approach in [Kumari et al. 2022]. Specifically, latent image features form the query, while text features are projected into the key and value using the projection matrices W_k and W_v . The cross-attention operation is conducted between the query, key and value, and we only fine-tune the parameters W_k and W_v , as shown in Fig. 2 (b).

We fine-tune the Stable Diffusion model "SD-V1-5" on I_E , then we can concatenate V* with the input text prompt T to form a new prompt such as "V* running, holding a trophy," to generate a customized raster image I_C that reflects the intended edits. To further promote the generation of customized images with flat vector style, we add a suffix to the prompt during inference: "minimal flat 2d



Figure 3: Our path alignment module for generating initial paths for customized SVG based on deep DINO-ViT features.

vector. no outlines. trending on artstation." As sampling results may not always align well with text prompts, we generate K images and select the sample that exhibits the highest consistency with the text prompt based on similarities in CLIP ViT-B/32 features [Radford et al. 2021]. For generated images containing backgrounds, we use U2-Net [Qin et al. 2020] to automatically mask out background.

4.2 Path Alignment

After obtaining a customized image I_C , our goal is to convert it into a vectorized version SVG_C . However, the quality of image vectorization strongly relies on initialization. Random initialization often leads to unsuccessful topological extraction and generates redundant shapes. To address this issue, LIVE [Ma et al. 2022] attempts to determine the initial path location based on the color and size of each connected component. However, color and size features are inadequate for retrieving topological path information due to the complex occlusion relationships among elements.

To tackle this challenge, our proposed solution is to reuse suitable paths from the exemplar SVG_E as initialization. To accomplish this, we need to establish correspondences between paths in SVG_E and connected components in I_C . We can then transform and scale selected paths of SVG_E to match their corresponding components in I_C . As customization may require appearance changes, we use semantic features from the pre-trained deep DINO-ViT model "dinovitb8" [Caron et al. 2021] instead of relying on lower-level features such as color and shape for matching. This approach enables us to establish more reliable semantic correspondences.

Dual Semantic Matching. Deep DINO-ViT features have been explored and applied to various visual tasks, such as image cosegmentation and object part correspondence [Amir et al. 2021; Hadjivelichkov et al. 2023], providing fine-grained local semantic features of object parts. Therefore, we adopt DINO-ViT features to represent paths in SVG_E and components in I_C in order to find semantic correspondences between them. In DINO-ViT, each layer contains a set of features, including a query, key, value, and token. Following the analysis in DeepViT [Amir et al. 2021], we utilize the key features to represent the semantic information of images. Therefore, both I_E and I_C are fed into a pre-trained DINO-ViT model to obtain the corresponding key feature maps F_E and F_C , where each point in the feature map represents semantic feature of an 8×8 image patch, as illustrated in Fig. 3 (a).

To perform semantic matching, we compute corresponding features for each path. Specifically, we segment I_E according to the paths in SVG_E , where each segment corresponds to a path with no overlaps. We then down-sample the segmentation map with the same resolution as F_E and calculate the average features of F_E within each segment. This process generates the feature representation for the *n* paths, forming a path feature vector $\{F_{Path_i}\}_{i=1}^n$. Similarly, to obtain a feature vector for the *m* components in I_C , we segment the image into connected components with a uniformlyfilled color and accumulate feature points in F_C . This feature vector is denoted as $\{F_{Comp_j}\}_{j=1}^m$. We then compute the pairwise cosine similarity between the feature vectors F_{Path_i} and F_{Comp_i} to construct a similarity matrix Sim, where Sim(i, j) denotes the similarity between the *i*-th path and the *j*-th component. To build robust correspondence between paths and components, we perform dual semantic matching. This involves applying a dual softmax to Sim by normalizing it in both row and column dimensions, resulting in a dual similarity matrix Sim₂ as follows:

 $Sim_2 = \text{softmax}(Sim, dim = row) \odot \text{softmax}(Sim, dim = col)$ (1)

where \odot indicates the element-wise product.

As shown in Fig. 3 (b), for each component with index j, we match it with the *i*-th path if $Sim_2(i, j)$ is the largest value in the *j*-th column. However, we only consider the match to be valid if $Sim_2(i, j)$ exceeds a predefined threshold $\tau_{th} = 0.0625$. Scores falling below this threshold indicate low similarity and are discarded. By leveraging the dual semantic matching mechanism, we can obtain more accurate semantic correspondence.

Path Pre-align. As shown in Fig. 3 (c), to initialize paths that align with I_C , we transform and scale the matched paths from SVG_E to their corresponding components in I_C . Specifically, we compute the convex hull for matched image components that have a corresponding SVG path, estimate the affine transformation matrix between the hull and path boundary points using the coherent point drift (CPD) algorithm [Myronenko and Song 2010], and accordingly transform the path. For unmatched components in I_C , we use a curve fitting method [Selinger 2003] to fit their boundaries with new paths composed of piecewise cubic Bézier curves. This results in a new path set $SVG_C^0 = \{Path_{1}^0, Path_{2}^0, \dots, Path_{m}^0\}$, which consists of the matched and transformed paths from the exemplar SVG, along with the fitted paths for the unmatched components.

4.3 Path Optimization

The path alignment step provides a good initialization for the customized vector graphic, which does not need to be highly accurate. To better reconstruct the customized image with our customized SVG, we perform a path optimization process. Given the initial path parameters SVG_C^0 , the goal is to obtain a refined set of path parameters SVG_C^0 by optimizing control point positions and fill colors to better align with I_C while preserving the regularity of the initial paths. To accomplish this, we employ both image-level and vector-level losses to guide the path optimization process (Fig. 4).

Image-Level Loss. The differentiable rasterizer [Li et al. 2020] establishes a connection between the SVG parameters and the raster image, enabling backpropagation of image-level losses to optimize path parameters such as control point positions and filling colors.

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia



Figure 4: The path optimization module of our method.

Drawing inspiration from CLIPasso [Vinker et al. 2022], we employ a CLIP-based loss to evaluate the image-level similarity between the rendered image of the customized SVG SVG_C^i at the current *i*-th iteration and the target image I_C . We compute the L_2 distance between their intermediate-level activations of CLIP as follows:

$$L_{CLIP} = \sum_{l} \|CLIP_{l}(I_{C}) - CLIP_{l}(\mathcal{R}(SVG_{C}^{i}))\|_{2}^{2}, \qquad (2)$$

where $CLIP_l$ denotes the CLIP encoder activation at layer *l*. For this purpose, we utilize layers 3 and 4 of the ResNet101 CLIP model. By employing the CLIP-based loss, we encourage the customized SVG to be faithful to the customized image.

Vector-Level Loss. To avoid messy paths during optimization while still allowing for necessary deformation, we incorporate the Procrustes distance [Wang and Mahadevan 2008] into our optimization process. The Procrustes distance ProDist(P_1 , P_2) measures the similarity between two point sets P_1 and P_2 by minimizing their dissimilarity through a sequence of rigid transformations, including translation, scaling, and rotation operations. After applying these rigid transformations, the Procrustes distance is computed by summing up the Euclidean distances between the corresponding points in the aligned two point sets.

We apply the Procrustes distance to a local window of every control point for every path. Let $p_{j,k}^i$ denote the *k*-th control point in the *j*-th path of the optimizing SVG_C^i in the current *i*-th iteration, while $p_{j,k}^0$ denotes the same control point in the initialized SVG_C^0 . Our local Procrustes loss is defined as:

$$L_{Procrustes} = \sum_{j}^{m} \sum_{k}^{d_{j}} \operatorname{ProDist}(\mathcal{W}(p_{j,k}^{0}), \mathcal{W}(p_{j,k}^{i})).$$
(3)

Here, W(p) denotes a neighboring point set with a local window centered at the control point p. The window size can be adjusted to balance local and global rigidity, with a larger window size implying a stronger constraint on global rigidity. In our experiments, we set the window size to be 4 neighboring control points. Procrustes distances are accumulated for all d_j control points, where j ranges from 1 to m, the total number of paths.

This local Procrustes loss applies geometric constraints to control points, regulating the shape deformation of paths to be as locally rigid as possible. As illustrated in Fig. 4, the local Procrustes loss is small when the shape deformation remains locally rigid, while it becomes larger in the presence of curve intersections or abrupt curvature changes. By constraining path optimization in this manner, we preserve the regularity of the paths while allowing for some degree of deformation to match the desired customization.

The overall loss function *L* is defined as a combination of both CLIP loss and local Procrustes loss:

$$\mathcal{L} = \mathcal{L}_{\text{CLIP}} + \lambda \mathcal{L}_{\text{Procrustes}} \tag{4}$$

where λ is a balance factor. As there is a trade-off between the CLIP loss and the local Procrustes loss, we prioritize the CLIP loss in the early stages of optimization to achieve a better fit to the customized image, and then increase the importance of the local Procrustes loss to regularize the path shape. Specifically, we linearly increase λ from 0.01 to 0.04 over 200 iterations. After convergence, we obtain the optimized customized SVG, denoted as $SVG_C = \{Path_1^*, Path_2^*, ..., Path_m^*\}$.

5 EXPERIMENT

Experiment Setup. To evaluate our method, we collected 100 vector graphics from *freepik*¹ and *iconfont*², including 45 characters, 35 animals, and 20 scenes. We select text prompts from the Stable Diffusion Prompts dataset [Dehouche and Kullathida 2023], which includes diverse actions and scenes. For each SVG, we randomly select 5 prompts and generate corresponding customized SVGs, resulting in 500 pairs of exemplar and customized vector graphics. The average number of paths in our customized SVGs is 28.

5.1 Evaluation Metrics

To assess the quality of text-to-SVG synthesis, we evaluate the generated SVGs at vector-level, image-level and text-level.

Vector-level Evaluation. Evaluating text-to-SVG synthesis is challenging without ground truth SVGs. We directly compare the customized SVG paths with the input exemplar SVG paths based on shape similarity, since the exemplar SVGs can be regarded as well-designed SVGs. Higher similarity indicates that the customized SVG better preserves the original shape properties. In addition, we evaluate vectorization quality based on smoothness criteria derived from the prior perception and computer graphics research [Dominici et al. 2020]. Specifically, the vector-level evaluation metrics include: (a) **Shape similarity** (*Sim_{shape}*): $1 - D_H$, where D_H is Hausdorff distance between path control points of the exemplar SVG and the customized SVG. (b) **Smoothness**: Inverse of the average curvature variation of the paths in customized vector graphics.

Image-level Evaluation. As the customized image should keep the same visual identity as the exemplar image, and the customized SVG should follow the customized image after rendering, we define the image-level metrics as: (a) **Exemplar image similarity** (Sim_{exp}) : Cosine similarity between the exemplar image and the rendered customized SVG result in CLIP space [Vinker et al. 2022]. (b) **Customized image similarity** (Sim_{cus}) : 1 – *MSE*, where *MSE*

Table 1: Quantitative comparison with existing methods. Sim_s , Sim_{cus} and Sim_{exp} denote shape similarity between SVG_E and SVG_C , RGB similarity between I_C and rendered SVG_C , CLIP similarity between I_E and rendered SVG_C .

Methods	Vector Level		Image Level		Text level
	$Sim_s \uparrow$	Smooth ↑	Sim_{cus} \uparrow	Sim_{exp} \uparrow	Sim_{CLIP} \uparrow
Potrace	-	0.7919	0.9937	0.9076	0.3218
LIVE	-	0.5929	0.9906	0.8576	0.2859
CLIPDraw	0.1380	0.4613	-	0.7862	0.3175
Vectorfusion	0.1434	0.4895	-	0.7214	0.2718
Ours	0.8121	0.7931	0.9944	0.9111	0.3227



Figure 5: Qualitative comparison with text-guided SVG optimization methods. Exemplar SVGs: the 1^{st} row is from ©Freepik; the 2^{nd} row is from illustac creator ©Zarame; the 3^{rd} row is from envatoelements creator ©Telllu.

is the mean squared error between the customized image and the rendered customized SVG result in RGB space.

Text-level Evaluation. To compute whether the customized SVG is aligned with the input text prompt, we define the text-level similarity by calculating the CLIP cosine similarity between the text prompt and rendered customized SVG (*Sim_{CLIP}*).

5.2 Comparison with Existing Methods

We compare our proposed pipeline with two types of T2V works: vectorization with T2I and text-guided SVG optimization methods. For vectorization with T2I methods, we use the same customized images generated by our concept fine-tuning method, and select both traditional and deep learning-based vectorization methods for comparisons: (a) Potrace [Selinger 2003]: A traditional vectorization method that segments images into regions and then fits piecewise cubic Bézier curves to the region boundaries. (b) LIVE [Ma et al. 2022]: A deep learning method generates SVGs by initializing path locations based on region colors and then uses loss functions to optimize paths. We use the same number of paths as our method for fairness, and set the control points of each path to 12, consistent with their paper. For text-guided SVG optimization, we compare our method with CLIP-based and diffusion-based optimization approaches: (c) CLIPDraw [Frans et al. 2022]: This method optimizes CLIP's image-text similarity metric to generate vector graphics from text prompts. (d) VectorFusion [Jain et al.

¹https://www.freepik.com/

²https://icofont.com/

2022]: This approach employs SDS loss to optimize SVGs to be consistent with given text prompts. We use the exemplar SVG as initialization. For a fair comparison, if the exemplar SVG has fewer paths than our initialization, we randomly add paths to match the number of paths as our method. The quantitative result is shown in Tab. 1 and the qualitative results are shown in Fig. 5 and Fig. 9.

Vectorization with T2I Methods. Potrace [Selinger 2003] segments images into regions based solely on color, thus it can generate multiple redundant paths for a single semantic element, such as multiple paths for a single pants in the 3^{rd} row of Fig. 9 (c). This leads to higher complexity. It also loses the layer-wise relationships provided by the exemplar SVG and curve smoothness along boundaries due to the lack of pixel-level supervision. As for LIVE [Ma et al. 2022], due to the path complexity caused by the occlusion, only relying on color and size features to reconstruct the customized image fails to maintain the topological path properties, such as the broken paths in Fig. 9 (d). When using the same number of paths as our method, the similarity to the customized image reduces a lot as shown in Tab. 1. In contrast, our model employs a semantic-based path alignment module that preserves layer-wise path properties and a path optimization module to obtain smooth and accurate final paths, outperforming these two methods quantitatively.

Text-guided SVG Optimization Methods. Although these methods directly inherit the original layer-wise properties, neither the CLIP loss nor the SDS loss is sufficient for preserving path quality. The original paths may undergo complex transformations to generate text-conforming SVGs, resulting in low shape similarity with the exemplar SVG (Tab. 1). As shown in Fig. 5, there are instances of intersecting paths and abrupt curvature changes, which lead to diminished path smoothness (Tab. 1) and visually unappealing outcomes. While CLIPDraw directly optimizes CLIP loss and achieves high CLIP metrics, the rendered SVGs appear messy. Our pipeline generates much clearer vector graphics than CLIPDraw and Vector-Fusion, by finetuning the diffusion model as a prior for customized images and preserving the desirable path properties of the exemplar SVG during the vectorization process.

5.3 User Study

We conduct a perceptual study to evaluate our vector graphics customization from three perspectives: overall SVG quality, alignment with text prompt, and similarity in visual identity with the exemplar SVG. We randomly select 12 pairs of exemplar SVGs and text prompts from our test set, and generate customized SVGs using the baseline methods in Sec. 5.2 and our approach. In each question, we display the results of different methods in random order and ask 30 participants to select the best result among five options for each evaluation metric. Fig. 6 demonstrates the superior performance of our method, as it achieves the highest preference in all evaluation metrics. Specifically, our method obtains 82.5% of votes for overall SVG quality, 83.7% for text alignment, and 89.1% for similarity in visual identity. The results show the effectiveness of our method in generating high-quality customized SVGs based on text prompts while preserving the visual identity of the exemplar SVG. SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia



Figure 6: User Study. We show the human preferences in %.

 Table 2: Ablation study on path alignment and path optimization modules.

Methods		$Sim_s \uparrow$	Smooth ↑	Sim _{cus} ↑
	w/o PA	0.2280	0.6217	0.9879
Path	Shape	0.7435	0.7902	0.9934
Alignment	Color	0.7271	0.7304	0.9912
(PA)	Polyfit	-	0.7922	0.9948
	Polyfit + Palette	-	0.7923	0.9946
Path	w/o L _{CLIP}	0.9998	0.8078	0.7197
Optimization	w/o L _{Procrustes}	0.4153	0.5030	0.9951
Ours		0.8121	0.7931	0.9944



Figure 7: Qualitative results on ablation study. The exemplar SVG is from Freepik creator ©Alexdndz.

5.4 Ablation Study

The Effectiveness of Path Alignment. To investigate the effectiveness of our path alignment module, we compare it with five baselines. The first one is using the input exemplar SVG directly as the path initialization (w/o PA). Due to the lack of the path alignment module, with the large visual difference between the customized image and the rendered exemplar SVG, significant distortions occur after path optimization, as shown in Fig. 7 (c). Rather than using our semantic features for finding the matched paths, the other two baselines are using Color and Shape similarities, respectively. For color, we directly compute the MSE in RGB space and for shape, we compute the distance after normalizing paths under rigid transformations. Given that many paths share similar colors and the large structure/shape deformation exists in the customized image, using shape and color similarities leads to incorrect path matches as shown in Fig. 7 (d, e). Thus the final optimized paths become messy (Tab. 2). We also ablate with standard vectorization methods: PolyFit [Dominici et al. 2020] and PolyFit with Palette compaction [Yang et al. 2023] to get initial paths. The latter method reduces noisy SVG paths by compacting the color palette of the customized image. Although their initial vectorizations fit

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia

Peiying Zhang, Nanxuan Zhao, and Jing Liao



Figure 8: Failure cases. Exemplar SVGs: the left is from ©iconfont; the right is from illustac creator ©Rico.

images well, the topological problem remains, as image segmentation cannot reflect occlusion relationships (Fig. 7 (f, g)). Instead, our model derives layer-wise paths from the exemplar SVG, preserving favorable path properties.

The Effectiveness of Path Optimization. We ablate the loss design by individually removing each of the CLIP loss and local Procrustes loss in the path optimization module. As shown in Fig. 7 (i), when the CLIP loss is removed ($w/o L_{CLIP}$), the optimization retains initialization and fails to deform paths to align with the customized image, resulting in low customized image similarity (Sim_{cus} in Tab. 2). When the local Procrustes loss is removed ($w/o L_{Procrustes}$), the deformed paths suffer from intersection issues and abrupt curvature changes, generating SVG paths of poor quality with low smoothness. Furthermore, the reduced shape similarity (Sim_s in Tab. 2) suggests that image-level loss alone is insufficient to preserve the shape properties inherited from the exemplar SVG.

6 CONCLUSION

In this paper, we investigate a novel problem called text-guided vector graphics customization, and propose a pipeline that generates customized SVGs based on input text prompts while preserving the visual identity of the input exemplar SVG. Our method reuses paths from the exemplar SVG through newly designed semanticbased path alignment and path optimization modules, maintaining the valid path properties and layer relationships in the final SVGs. Though our method can achieve high-quality SVG results, we still suffer from failure cases as shown in Fig. 8. First, our method relies on the generative capabilities of diffusion model, thus inappropriate pairs of exemplar SVGs and text prompts may result in unreasonable customizations. For example, requesting a car to "hold a trophy" could yield a nonsensical outcome. Second, if the customized image deviates significantly from the exemplar, the path alignment method may fail to find correct matches. For instance, if the exemplar depicts a side view of a running squirrel, but the customized image shows a frontal view of a smiling squirrel, the path alignment fails to find correct path matches from the exemplar SVG.

Ethics. Our work aims to enhance, not replace, the creative flows of vector artists. We acknowledge potential ethical concerns with models trained on internet data (e.g. property rights). Solutions such as content detection systems can help mitigate these risks.

Future Work. Our text-guided vector graphics customization exhibits visual creativity. In future work, we plan to integrate humanin-the-loop methods to assist designers in achieving more precise vector graphics generation and editing.

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

- Gavin Ambrose, Paul Harris, and Nigel Ball. 2019. The fundamentals of graphic design. Bloomsbury Publishing.
- Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. 2021. Deep vit features as dense visual descriptors. arXiv preprint arXiv:2112.05814 2, 3 (2021), 4.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. 2021. Emerging properties in self-supervised vision transformers. In Proceedings of the IEEE/CVF international conference on computer vision. 9650–9660.
- 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).
- James Richard Diebel. 2008. Bayesian Image Vectorization: the probabilistic inversion of vector image rasterization. Stanford University.
- 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-todrawing synthesis through language-image encoders. Advances in Neural Information Processing Systems 35 (2022), 5207–5218.
- Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. 2022. An image is worth one word: Personalizing text-toimage generation using textual inversion. arXiv preprint arXiv:2208.01618 (2022).
- David Ha and Douglas Eck. 2017. A neural representation of sketch drawings. arXiv preprint arXiv:1704.03477 (2017).
- Denis Hadjivelichkov, Sicelukwanda Zwane, Lourdes Agapito, Marc Peter Deisenroth, and Dimitrios Kanoulas. 2023. One-Shot Transfer of Affordance Regions? AffCorrs!. In Conference on Robot Learning. PMLR, 550-560.
- 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.
- Ajay Jain, Amber Xie, and Pieter Abbeel. 2022. VectorFusion: Text-to-SVG by Abstracting Pixel-Based Diffusion Models. arXiv preprint arXiv:2211.11319 (2022).
- 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).
- Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip Torr. 2019. Controllable text-to-image generation. Advances in Neural Information Processing Systems 32 (2019).
- 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.
- 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.
- Andriy Myronenko and Xubo Song. 2010. Point set registration: Coherent point drift. IEEE transactions on pattern analysis and machine intelligence 32, 12 (2010), 2262–2275.
- 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).
- Xuebin Qin, Zichen Zhang, Chenyang Huang, Masood Dehghan, Osmar R Zaiane, and Martin Jagersand. 2020. U2-Net: Going deeper with nested U-structure for salient object detection. *Pattern recognition* 106 (2020), 107404.
- Antoine Quint. 2003. Scalable vector graphics. IEEE MultiMedia 10, 3 (2003), 99–102. 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.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In International Conference on Machine Learning. PMLR, 8821–8831.

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia

- 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.
- Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. 2016. Generative adversarial text to image synthesis. In *International* conference on machine learning. PMLR, 1060–1069.
- 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.
- 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 subject-driven generation. arXiv preprint arXiv:2208.12242 (2022).
- 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.

- 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).
- Yingtao Tian and David Ha. 2022. Modern evolution strategies for creativity: Fitting concrete images and abstract concepts. In Artificial Intelligence in Music, Sound, Art and Design: 11th International Conference, EvoMUSART 2022, Held as Part of EvoStar 2022, Madrid, Spain, April 20–22, 2022, Proceedings. Springer, 275–291.
- Yael Vinker, Ehsan Pajouheshgar, Jessica Y Bo, Roman Christian Bachmann, Amit Haim Bermano, Daniel Cohen-Or, Amir Zamir, and Ariel Shamir. 2022. Clipasso: Semantically-aware object sketching. ACM Transactions on Graphics (TOG) 41, 4 (2022), 1–11.
- Chang Wang and Sridhar Mahadevan. 2008. Manifold alignment using procrustes analysis. In Proceedings of the 25th international conference on Machine learning. 1120–1127.
- Jinfan Yang, Nicholas Vining, Shakiba Kheradmand, Nathan Carr, Leonid Sigal, and Alla Sheffer. 2023. Subpixel Deblurring of Anti-Aliased Raster Clip-Art. In Computer Graphics Forum, Vol. 42. Wiley Online Library, 61–76.
- 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.

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia

Peiying Zhang, Nanxuan Zhao, and Jing Liao



Figure 9: Qualitative comparison to vectorization with T2I methods. Exemplar SVGs: the 1^{st} and 4^{th} rows are from ©iconfont; the 2^{nd} row is from envatoelements creator ©Masastarus; the 3^{rd} and 6^{th} rows are from Freepik creator ©Alexdndz; the 5^{th} and 7^{th} rows are from ©Freepik; the 8^{th} row is from ©Vectorportal.

SA Conference Papers '23, December 12-15, 2023, Sydney, NSW, Australia



(b) Text-Guided Vector Graphics Customization

Figure 10: More results of our text-guided vector graphics customization. Exemplar SVGs: the 1st and 5th rows are from Freepik creator ©Alexdndz; the 2^{nd} row is from envatoelements creator ©Telllu; the 3^{rd} and 4^{th} rows are from ©Freepik; the 6^{th} row is from Iconduck creator ©Pavan Kondapuram.