

SVGBuilder: Component-Based Colored SVG Generation with Text-Guided Autoregressive Transformers

Zehao Chen, Rong Pan*

School of Computer Science and Engineering, Sun Yat-sen University
Guangzhou, Guangdong, China
chenzh593@mail2.sysu.edu.cn, panr@sysu.edu.cn

Abstract

Scalable Vector Graphics (SVG) are essential XML-based formats for versatile graphics, offering resolution independence and scalability. Unlike raster images, SVGs use geometric shapes and support interactivity, animation, and manipulation via CSS and JavaScript. Current SVG generation methods face challenges related to high computational costs and complexity. In contrast, human designers use component-based tools for efficient SVG creation. Inspired by this, SVGBuilder introduces a component-based, autoregressive model for generating high-quality colored SVGs from textual input. It significantly reduces computational overhead and improves efficiency compared to traditional methods. Our model generates SVGs up to 604 times faster than optimization-based approaches. To address the limitations of existing SVG datasets and support our research, we introduce ColorSVG-100K, the first large-scale dataset of colored SVGs, comprising 100,000 graphics. This dataset fills the gap in color information for SVG generation models and enhances diversity in model training. Evaluation against state-of-the-art models demonstrates SVGBuilder’s superior performance in practical applications, highlighting its efficiency and quality in generating complex SVG graphics.

Project — <https://svgbuilder.github.io>

Extended version — <https://arxiv.org/abs/2412.10488>

Introduction

Scalable Vector Graphics (SVG) is a widely used XML-based vector image format for defining graphics. Unlike raster graphics, which are composed of a fixed grid of pixels, SVGs are defined by geometric shapes such as points, lines, curves, and polygons, which makes them resolution-independent and infinitely scalable without loss of quality. This inherent scalability offers significant advantages over bitmap images, particularly in web and mobile applications where varying display resolutions and device capabilities are common. Furthermore, SVG supports interactivity and animation, and because it’s text-based, it is easily searchable, compressible, and can be manipulated through CSS and JavaScript.

*Corresponding author.

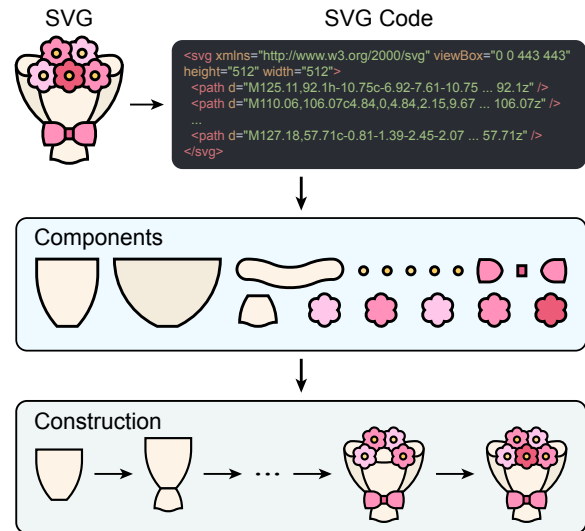


Figure 1: Illustration of component-based SVG generation, with each path in the SVG code considered as a component. The process demonstrates the construction of an SVG from individual components.

In the current methods for generating SVGs, one of the most straightforward approaches involves using text-to-image generation models (Ho, Jain, and Abbeel 2020; Rombach et al. 2022) to first create a bitmap, which is then converted into an SVG using vectorization techniques (Ma et al. 2022). While this method can produce high-quality results, it is computationally intensive and time-intensive. Another approach is optimization-based (Jain, Xie, and Abbeel 2023; Xing et al. 2024b), starting with a random SVG path and iteratively refining it to match a target SVG. Despite their ability to generate SVGs, these methods have notable drawbacks, such as prolonged processing times and high computational costs (Tang et al. 2024). Additionally, some research focuses on unguided random generation of SVG paths (Carrier et al. 2020), which poses challenges for practical applications due to the lack of directed output. Recently, there has been interest in leveraging autoregressive language models, such as GPT (Brown et al. 2020), to generate SVGs directly from textual descriptions (Wu et al. 2023a). However, SVGs




Figure 2: Random samples generated by SVGBuilder.

can contain extensive code, which poses a significant challenge for autoregressive language models when generating complex SVGs. The length and complexity of the code increase the difficulty of producing accurate and efficient SVG outputs, often leading to suboptimal performance in intricate design tasks. Additionally, generating SVGs based on individual tokens can easily result in incomplete or invalid paths, further complicating the generation process. Consequently, the creation of sophisticated SVGs using such models remains a significant computational and methodological challenge. For a more detailed discussion on the background of autoregressive language models and SVGs, please refer to the supplementary material in the extended version.

In contrast, human designers often create SVGs using software tools that facilitate component-based design. These tools allow designers to construct images by assembling, scaling, moving, and coloring reusable components. This modular approach simplifies the design process and enhances efficiency, as it enables the reuse of predefined elements and reduces the need to generate complex paths from scratch. Inspired by this observation, our model adopts a component-based strategy for SVG generation. By leveraging predefined graphical components, our approach aims to streamline the SVG creation process, reducing computational overhead and improving the feasibility of generating intricate designs programmatically. As illustrated in Figure 1, an SVG file consists of SVG code written in XML format, containing tags such as `<path>`. We consider each `<path>` as a distinct component. These components are then assembled to construct a complete SVG. This method not only mirrors the intuitive workflows of human designers but also offers a structured and scalable framework for automat-

ing SVG generation, thereby addressing some of the limitations inherent in current generation techniques.

Moreover, due to the significant differences between stroke-based SVGs and SVGs with path fills, previous SVG-related datasets primarily focus on stroke-based SVGs. In contrast, datasets containing colored SVGs are mostly stored in bitmap formats (Lu et al. 2021) or consist of small-scale collections, which are insufficient for training models to learn diverse variations and generation patterns. To address this gap, we develop the first large-scale colored SVG dataset, ColorSVG-100K, aimed at enhancing the capability of models to generate SVGs with rich color information.

To address the aforementioned issues, we present SVGBuilder, a novel model that leverages a component-based, autoregressive approach guided by textual input for colored SVG generation. Our method significantly reduces the generation time and computational expense associated with traditional techniques while ensuring practical applicability by providing meaningful, text-driven guidance for the SVG output. This approach effectively overcomes the challenges of random generation and the difficulties inherent in generating lengthy SVG sequences using autoregressive models. To the best of our knowledge, we are the first to propose a component-based autoregressive model for generating colored SVGs. As shown in Figure 2, these are some examples of SVGs generated by SVGBuilder.

Compared to state-of-the-art (SOTA) SVG generation models, our approach demonstrates significantly higher efficiency. Specifically, our method generates SVGs up to 604 times faster than optimization-based models currently available. Additionally, our approach outperforms other SOTA models on the FID metric, indicating superior performance.

Furthermore, in our case study, it is evident that the graphics produced by our method are of higher quality compared to those generated by other language-based models, showcasing the advantages of our approach in practical applications.

Our contribution lies in three key aspects:

- **Development of SVGBuilder:** We present SVGBuilder, the first component-based autoregressive model for generating colored SVGs guided by textual input. This method significantly reduces the generation time and computational expense associated with traditional techniques, ensuring practical applicability and overcoming the challenges of random generation and lengthy SVG sequences.
- **Introduction of ColorSVG-100K Dataset:** We develop the first large-scale dataset of colored SVGs, ColorSVG-100K, aimed at enhancing the capability of models to generate SVGs with rich color information. This dataset addresses the limitations of previous datasets that focused primarily on stroke-based SVGs or small-scale collections.
- **Superior Performance and Efficiency:** Our approach not only mirrors the intuitive workflows of human designers by leveraging predefined graphical components but also significantly outperforms SOTA SVG generation models on the FID metric. Our method generates SVGs much faster and produces higher-quality graphics, demonstrating superior efficiency and practical applicability in real-world scenarios.

Related Work

Text-to-Image Generation

The field of generating images from text has progressed significantly, evolving through several important phases involving GANs and diffusion models. GANs (Goodfellow et al. 2014) have been instrumental, employing a generator to produce images and a discriminator to evaluate them, refining the output iteratively to align with textual descriptions. While text-conditioned GANs (Kang et al. 2023; Qiao et al. 2019; Xu et al. 2018) have made considerable strides, they often encounter challenges related to scalability and stability when dealing with complex datasets, despite recent advancements (Kang et al. 2023). More recently, diffusion models have emerged as the preferred method for text-to-image generation. These models start with Gaussian noise and progressively refine it to generate coherent images. Text-guided diffusion models (Nichol et al. 2021; Rombach et al. 2022) incorporate text embeddings either directly or through cross-attention mechanisms to direct the image generation process. Previous research in this domain has predominantly focused on producing raster images with fixed resolutions. In contrast, our work targets the generation of vector icons from text, enabling arbitrary scaling and providing greater flexibility and potential for various applications.

SVG Generation

SVG uses XML code to generate graphics. Currently, a widely used approach is optimization-based, which involves

randomly initializing some SVG paths and then optimizing them using a differentiable rasterizer (Jain, Xie, and Abbeel 2023; Xing et al. 2024b). However, this method is highly time-intensive, as creating an SVG with 24 SVG paths can take over 20 minutes (Tang et al. 2024). Some methods generate SVGs purely through random initialization without text guidance (Carrier et al. 2020). Recent research explores autoregressive methods for generating SVGs (Wu et al. 2023a), but the lengthy code of SVG graphics poses challenges for generating complex SVGs using this approach. To address this issue, we adopt a component-based autoregressive generation method for SVGs, significantly improving the generation speed and overcoming the difficulties of using autoregressive techniques for complex SVGs, while achieving high-quality results.

Methodology

In this section, we provide a comprehensive explanation of our proposed method for SVG generation. Our system framework consists of two primary parts. The first part involves the construction of the component library, which serves as the foundation for the SVG generation process. The second part is our SVG generation model, which utilizes the established component library to produce high-quality SVGs. The overall system framework, highlighting the interaction between these parts, is illustrated in Figure 3.

Component Library Construction

In our system, the generation of models is component-based, making the establishment of a component library the critical first step, as illustrated in Figure 3 (1). In Figure 3 (1a), we begin by decomposing all SVGs in the training set into individual paths, with each path forming a unique component. These isolated components exhibit a wide variety of appearances, necessitating a thorough reorganization.

Next, as shown in Figure 3 (1b), we remove the colors from these components and normalize them. Initially, the geometric centers of these components do not align with the origin, complicating the scaling and normalization process. To address this issue, we calculate the geometric center of each component and translate it to the origin. We then scale the component such that its longest dimension measures 100 units, ensuring that the extents along both the positive and negative axes are 50 units. This process reorganizes and normalizes the components, making them ready for subsequent use in the model.

The initial component library contains many duplicate components. For example, flower shapes in a bouquet may be identical except for their colors. Therefore, it is necessary to merge these redundant components. As depicted in Figure 3 (1c), we start by merging paths that are identical in the SVGs. This initial merging reduces the number of components but does not eliminate all redundancies. To further refine the component library, we convert all components into 100×100 pixel images and then transform these images into grayscale. In these grayscale images, the areas enclosed by the paths are represented by 1 (white), and the background is represented by 0 (black). We then compute

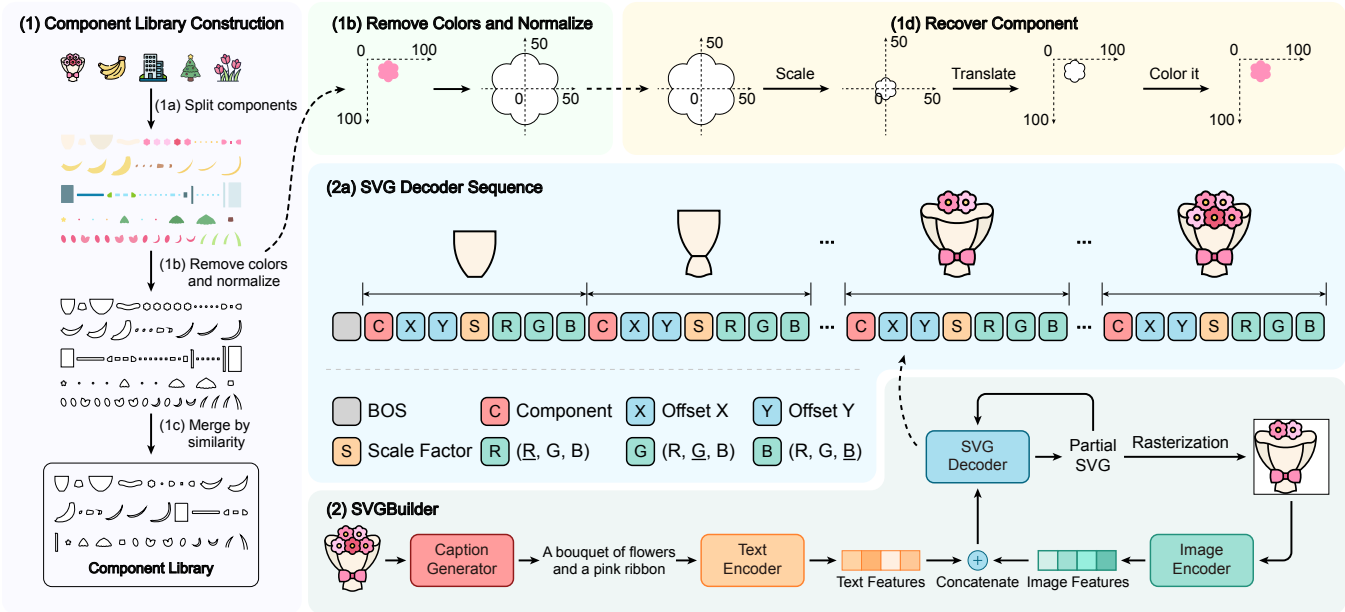


Figure 3: System framework overview. The framework comprises two primary steps. The first step is Component Library Construction, which includes the sub-process of Remove Colors and Normalize. The second step involves the SVGBuilder’s operations, which encompass the SVG Decoder Sequence and the Recover Component processes.

the Jaccard similarity index between each pair of images, sorting the similarity scores in descending order. Using a union-find data structure, we set a merging threshold and examine the similarity scores sequentially. If the similarity between two components meets or exceeds the threshold, we merge them in the union-find structure. The root nodes of these merged groups are considered the final components in the library. Subsequently, we replace occurrences of the child nodes with their corresponding root nodes in the training set. This process results in a refined and consolidated component library, ensuring minimal redundancy and optimal component organization for the model.

When the normalized components need to be restored, as shown in Figure 3 (1d), it is necessary to place them in the correct positions, similar to how a human would arrange elements when drawing in software. This involves scaling the components according to a scale factor and translating them based on offset values on the X and Y axes (Offset X and Offset Y). This ensures accurate placement of the components. Finally, the generated RGB tokens are used to colorize the component, restoring its original appearance.

SVGBuilder

Our SVG generation model comprises three main components: the Text Encoder, the Image Encoder, and the SVG Decoder, as illustrated in Figure 3 (2). During training, the original dataset typically provides minimal text prompts, usually limited to category labels like “flowers” without further description of the SVG. To address this limitation, we introduce a Caption Generator, which is employed solely during training and not during inference.

Initially, we rasterize the SVG into an image and feed this

image into the Caption Generator, which produces a caption, such as “A bouquet of flowers and a pink ribbon.” This textual description is then input into the Text Encoder to generate text features. At the same time, we generate image features by inputting a blank image into the Image Encoder. The Image Encoder outputs image features, which are concatenated with the text features. These combined features are then input into the SVG Decoder. The SVG Decoder also takes an SVG Decoder Sequence as input to generate the next token. The partially generated SVG is rasterized into an image, which is subsequently fed back into the Image Encoder. This process repeats until an end token is generated or the maximum sequence length is reached. The role of the Image Encoder in this process is to allow the model to understand semantics not only from textual or path-based perspectives but also from visual perspectives. By incorporating the Image Encoder, the model can recognize and comprehend the components it selects at each step, including their layout and color. This dual learning mechanism ensures that the model grasps the specific visual attributes and spatial arrangements of the elements, enhancing its overall semantic understanding and performance in tasks involving both textual and visual data.

Specifically, as shown in Figure 3 (2a), the SVG Decoder Sequence consists of several tokens. Each complete component is represented by seven tokens. The sequence begins with a BOS (beginning of sequence) token, followed by a token representing the component. The next tokens indicate the translation offsets (Offset X and Offset Y) and the scaling factor (Scale Factor), specifying how the component should be placed. Finally, the sequence includes tokens for color (R, G, and B tokens), representing the RGB values.

This sequence is repeated for each component until the entire SVG is generated.

Experiments

In this section, we present a series of experiments conducted to substantiate the efficacy of our proposed SVGBuilder methodology. This section elucidates our evaluation metrics, describes our baseline comparisons, and offers a comprehensive analysis of the results derived from our experimental investigations. We also introduce our proposed ColorSVG-100K dataset, which serves as a critical resource for our experiments. Additionally, we conducted SVG complexity analysis, ablation experiments, and a case study to further validate our approach.

Dataset

The creation of the ColorSVG-100K dataset arises from the significant limitations observed in existing SVG datasets. Available SVG datasets are predominantly either black and white or stored as raster images, lacking the essential vector properties necessary for many applications. Such limitations pose substantial challenges for tasks requiring high-quality, scalable, and color-rich SVGs. The primary motivation behind developing the ColorSVG-100K dataset is to address these gaps and provide a robust foundation for research and applications in the field of colored SVG generation. This dataset, comprising 100,000 richly colored SVGs, is the first of its kind on a large scale. It not only fills the void of color information in existing datasets but also leverages the advantages of the SVG format, such as scalability and editability. For more information about the ColorSVG-100K dataset, including the motivation for its creation, the specific construction process, dataset statistics, and examples, please refer to the supplementary material in the extended version.

Experimental Settings

Evaluation Metrics We adopt evaluation metrics from SkexGen, IconShop (Wu et al. 2023a), and StrokeNUWA. Specifically, we use the Fréchet Inception Distance (FID) (Heusel et al. 2017) to quantify the distance between the image features of the generated and ground-truth SVGs. Additionally, we utilize two types of CLIPScore (Radford et al. 2021): CLIPScore-T2I, which measures the similarity between the rasterized images of the generated SVGs and the textual prompts used for generation, and CLIPScore-I2I, which measures the similarity between the rasterized images of the generated SVGs and the rasterized images of the ground-truth SVGs. Furthermore, we assess the “Uniqueness” and “Novelty” of the generated SVGs, which are derived from SkexGen. “Uniqueness” indicates the proportion of generated data occurring only once among all generated results, while “Novelty” refers to the proportion of generated data absent from the training set. We also incorporate the Human Preference Score (HPS) (Wu et al. 2023b) to evaluate user satisfaction with the generated outputs. These metrics collectively provide a comprehensive evaluation of the quality, originality, and performance of our SVG generation methodology. Lastly, we measure the generation speed per SVG to assess the efficiency of our method.

Baselines In this paper, we compare our method with various SOTA open-source models. VectorFusion (Jain, Xie, and Abbeel 2023) utilizes diffusion models trained on pixel representations of images to generate SVGs, without requiring large datasets of captioned SVGs. By optimizing a differentiable vector graphics rasterizer (Li et al. 2020), VectorFusion extracts semantic knowledge from a pretrained diffusion model, resulting in high-quality vector graphics across various styles. CLIPDraw (Frans, Soros, and Witkowski 2022) employs a pretrained CLIP language-image encoder to synthesize novel drawings by optimizing the similarity between the generated drawing and the given description. It operates over vector strokes rather than pixel images, producing simple, human-recognizable shapes. DiffSketcher (Xing et al. 2024a), based on a pretrained diffusion model, optimizes a set of Bézier curves to create vectorized free-hand sketches that retain the structure and visual details of the subject. We re-implemented IconShop using Flan-T5 (Chung et al. 2024) as the backbone, employing autoregressive transformers to sequentialize and tokenize SVG paths and textual descriptions, resulting in improved icon synthesis capabilities. Additionally, we explored the capabilities of GPT-3.5 and GPT-4 (Achiam et al. 2023) in SVG generation, demonstrating advancements in large language models (LLMs) for this task. All these methods are text-guided, allowing us to evaluate the performance and effectiveness of the proposed SVGBuilder methodology against the current SOTA methods in the field of SVG generation.

Implementation Details In our experiments, when rasterizing SVGs, we place them on a white background. For constructing the ColorSVG-100K dataset, “Category Calibration” corrects misclassified SVGs. Both the Text Encoder and Image Encoder in this stage, as well as during model training and FID calculation, are CLIP models (Radford et al. 2021). During SVG processing, we utilize the DeepSVG Library. The Caption Generator is BLIP-2 (Li et al. 2023), and the SVG Decoder is based on GPT-2 (Radford et al. 2019). The batch size for training is set to 16, with a learning rate of 6×10^{-4} . We use the AdamW optimizer for 50 epochs of training. The experiments are conducted on a single NVIDIA A800 GPU. For more implementation details, please refer to the supplementary material in the extended version.

Performance Comparison

We evaluate the performance of different models on the ColorSVG-100K dataset, as illustrated in Table 1. The models are categorized into two types: optimization-based and language-based. The arrows in the table indicate the preferred direction, with down arrows indicating lower values are better, and up arrows indicating higher values are better.

In the Table 1, CLIPScore is divided into T2I, which represents the CLIPScore between the generated SVG image and the text description, and I2I, which represents the CLIPScore between the generated SVG image and the real SVG image. From Table 1, it is evident that the optimization-based models have lower FID scores, indicating that the images they generate are closer to the distribution of real im-

Model Type	Model	FID (\downarrow)	CLIPScore		HPS (\uparrow)	Uniqueness (\uparrow)	Novelty (\uparrow)	Generation Speed per SVG (\downarrow)
			T2I (\uparrow)	I2I (\uparrow)				
Optimization Based	VectorFusion	20.62	27.86	<u>77.84</u>	17.65	100	100	\approx 13.69 min (1.0 \times)
	DiffSketcher	<u>20.28</u>	27.04	77.70	17.62	100	100	\approx 4.27 min (3.2 \times)
	CLIPDraw	30.11	28.46	73.13	17.51	100	100	\approx 2.92 min (4.7 \times)
Language Based	GPT-3.5	33.19	22.05	77.19	16.96	96.2	100	10.8 s (76.1 \times)
	GPT-4	31.85	22.57	77.88	16.99	98.4	100	23.8 s (34.5 \times)
	IconShop	41.03	21.96	73.53	16.57	97.8	100	8.52 s (96.4 \times)
	Ours	15.93	<u>22.76</u>	79.28	<u>17.10</u>	<u>99.2</u>	100	1.36 s (604.0\times)

Table 1: Performance comparison between SVGBuilder and the baseline models on the ColorSVG-100K dataset. Underlined values indicate the best performance within each model type, while **bold** values indicate the best overall performance across all models.

ages in the test set. In contrast, the language-based models, except for ours, generally have higher FID scores, suggesting that previous methods fail to accurately fit the dataset’s distribution.

Regarding CLIPScore, optimization-based methods generally achieve higher T2I scores. This might be because these models generate images that align better with the distribution of images on which the CLIP model is originally trained. The CLIP model might not have been extensively trained on SVG datasets, making it less effective at capturing the flat 2D nature of SVG graphics in CLIPScore-T2I scores. In contrast, the CLIPScore-I2I measures the similarity between generated and real SVG images, and hence is not affected by this limitation. Notably, our model also achieves the highest CLIPScore among the language-based models.

Similarly, since HPS uses the same CLIP model, it tends to favor SVG generation models that optimize from bitmap images. Regarding Uniqueness, the random nature of optimization-based models results in a 100% score, while language-based models exhibit slightly lower uniqueness due to potential similar outputs from similar inputs. Our model, however, achieves the highest uniqueness score among the language-based models. All models achieve a 100% score in Novelty.

In terms of Generation speed per SVG, the values in parentheses indicate the speed-up factor relative to VectorFusion. It is clear that optimization-based models take a significantly longer time to generate an SVG, with VectorFusion taking approximately 13.69 minutes and even the fastest, CLIPDraw, taking 2.92 minutes. In contrast, language-based models take only a few seconds, drastically reducing generation time. Our model achieves a generation time of 1.36 seconds, representing a 604.0 \times speed-up compared to VectorFusion, which underscores the potential and advantages of language-based models in SVG generation.

SVG Complexity Analysis

We further analyze the impact of varying SVG complexities on model performance. We sort the SVGs in ascending order based on the number of paths and divide them into four levels of complexity: 0-25%, 25-50%, 50-75%, and 75-100%. The 0-25% range represents the simplest SVGs with

Model	SVG Complexity			
	0-25%	25-50%	50-75%	75-100%
VectorFusion	32.70	30.26	<u>29.82</u>	31.34
DiffSketcher	<u>29.09</u>	<u>29.24</u>	30.62	<u>30.96</u>
CLIPDraw	42.86	40.43	38.48	39.37
GPT-3.5	32.18	38.59	42.52	47.00
GPT-4	32.28	37.62	40.46	43.88
IconShop	40.80	45.34	50.00	53.60
Ours	20.47	23.36	26.93	30.55

Table 2: Comparison of FID scores across different SVG complexity levels for various models. Underlined values indicate the best performance within each model type, while **bold** values indicate the best overall performance across all models.

the fewest paths, indicating the lowest complexity, while the 75-100% range encompasses the SVGs with the most paths, indicating the highest complexity. We evaluate the performance of different models across these complexity levels using the FID metric, as shown in Table 2. Our model achieves the best results among all models across all levels of complexity. It is observed that as the complexity of SVGs increases, the performance of optimization-based models does not deteriorate, whereas the performance of language-based models worsens. This may be because the randomly initialized bitmaps are generally more complex than SVGs.

Ablation Study

To evaluate the impact of different components on the overall performance of our model, we conduct an ablation study. The components subjected to ablation include the Caption Generator, Image Encoder, whether the Text Encoder is trained, and whether color fills are used in SVG during inference. We test each component sequentially, and the results are presented in Table 3. In these experiments, we observe both the FID and CLIPScore metrics. Since changes in the CLIPScore are not significant, we primarily focus on the FID metric. The results indicate that removing the Im-

Model	FID (\downarrow)	CLIPScore	
		T2I (\uparrow)	I2I (\uparrow)
Ours	15.93	22.76	79.28
w/o Caption Generator	17.40	21.94	78.23
w/o Image Encoder	16.90	22.58	77.97
w/o training Text Encoder	21.75	21.73	78.04
w/o SVG filling	17.39	22.64	77.79

Table 3: Performance comparison of SVGBuilder without different components.



Figure 4: Examples of SVGs generated by different models. DiffSketcher represents the optimization-based model, while the others are language-based models.

age Encoder results in the least performance loss on the FID metric, suggesting that the model does not effectively utilize the information from the Image Encoder. Removing the Caption Generator has a moderate impact on FID, as using the original SVG categories as input text lacks the data augmentation benefits provided by diverse captions. Freezing the Text Encoder leads to a significant decline in FID performance, highlighting the importance of training the Text Encoder for maintaining model performance. Lastly, not filling colors during inference and only performing outlines also results in a decrease in FID.

Case Study

In our case study, we select SVGs generated by different models for comparative analysis. We focus on language-based models, including GPT-4, IconShop, and our model. For comparison, we also include DiffSketcher, an optimization-based model with the lowest FID score. The results are illustrated in Figure 4.

DiffSketcher produces visually impressive results, accurately capturing the original shapes and colors of objects,

likely due to the application of the Diffusion model. However, some artifacts in the form of odd lines appear during the bitmap-to-SVG optimization process. The output from GPT-4 aligns well with the objects in terms of color and somewhat meets the contour requirements, but without accompanying text, the contours are hard to interpret. IconShop generates accurate outlines except for the taxi example, which performs poorly. This discrepancy might be due to the greater complexity of SVGs in our dataset compared to those used by IconShop, indicating its suitability for simpler SVG generation tasks.

Our model’s outputs match the objects in both color and contour. However, the “avocado” example is influenced by the seed, affecting the overall hue, and the “taxi” appears abstract. Overall, our model approaches the expression of the optimization-based models, demonstrating its potential.

Limitations

While SVGBuilder demonstrates strong performance in component-based autoregressive SVG generation, it does have certain limitations. Our ColorSVG-100K dataset fills a notable gap in the realm of colored SVG datasets, yet each path within the dataset currently lacks explicit semantics. This limitation forces the model to implicitly infer these semantics, which may not always be reliable. Additionally, the construction of the component library is crucial within this framework. We normalize components through translation and scaling based on similarity for merging. However, incorporating additional geometric transformations, such as rotation and reflection, could potentially further reduce the size of the component library. Although the similarity-based merging process is relatively time-consuming, it is conducted offline, representing a one-time computational cost. Furthermore, when utilizing a union-find data structure for merging, selecting the parent node is critical. Even with a similarity threshold in place, nodes within the same set may have varying degrees of similarity to the parent node due to the non-transitive nature of similarity. Additionally, the component-based approach can be more challenging to innovate compared to stroke-based methods. In our language model training, we use cross-entropy as the loss function, treating numbers as separate tokens. However, cross-entropy may not effectively capture geometric relationships. Looking ahead, we aim to address these challenges and make incremental advancements in the field of component-based autoregressive SVG generation.

Conclusion

In this study, we present SVGBuilder, an innovative component-based autoregressive model designed for the generation of colored SVGs, alongside the introduction of the ColorSVG-100K dataset. Through rigorous evaluation, our approach demonstrates SOTA performance, surpassing both optimization-based and language-based models in the FID metric. Remarkably, SVGBuilder not only excels in visual quality but also offers a substantial 604× enhancement in generation speed compared to optimization-based models, showing great potential for future applications.

Acknowledgements

The authors are supported by the National Key Research and Development Program of China (Grant No. 2020YFA0712500), the National Natural Science Foundation of China (Grant No. 12126609), and the Pazu Lab (Huangpu) Research and Development Project, funded under Grant No. 2023K0601. The authors would also like to thank National Supercomputer Center in Guangzhou for providing high performance computational resources.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Carlier, A.; Danelljan, M.; Alahi, A.; and Timofte, R. 2020. Deepsvg: A hierarchical generative network for vector graphics animation. *Advances in Neural Information Processing Systems*, 33: 16351–16361.
- Chung, H. W.; Hou, L.; Longpre, S.; Zoph, B.; Tay, Y.; Fedus, W.; Li, Y.; Wang, X.; Dehghani, M.; Brahma, S.; et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70): 1–53.
- Frans, K.; Soros, L.; and Witkowski, O. 2022. Clipdraw: Exploring text-to-drawing synthesis through language-image encoders. *Advances in Neural Information Processing Systems*, 35: 5207–5218.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Jain, A.; Xie, A.; and Abbeel, P. 2023. Vectorfusion: Text-to-svg by abstracting pixel-based diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1911–1920.
- Kang, M.; Zhu, J.-Y.; Zhang, R.; Park, J.; Shechtman, E.; Paris, S.; and Park, T. 2023. Scaling up gans for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10124–10134.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Li, T.-M.; Lukáč, M.; Gharbi, M.; and Ragan-Kelley, J. 2020. Differentiable vector graphics rasterization for editing and learning. *ACM Transactions on Graphics (TOG)*, 39(6): 1–15.
- Lu, P.; Qiu, L.; Chen, J.; Xia, T.; Zhao, Y.; Zhang, W.; Yu, Z.; Liang, X.; and Zhu, S.-C. 2021. IconQA: A New Benchmark for Abstract Diagram Understanding and Visual Language Reasoning. In *The 35th Conference on Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks*.
- Ma, X.; Zhou, Y.; Xu, X.; Sun, B.; Filev, V.; Orlov, N.; Fu, Y.; and Shi, H. 2022. Towards layer-wise image vectorization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16314–16323.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.
- Qiao, T.; Zhang, J.; Xu, D.; and Tao, D. 2019. Mirrorgan: Learning text-to-image generation by redescription. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1505–1514.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 8748–8763. PMLR.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Tang, Z.; Wu, C.; Zhang, Z.; Ni, M.; Yin, S.; Liu, Y.; Yang, Z.; Wang, L.; Liu, Z.; Li, J.; et al. 2024. StrokeNUWA: Tokenizing Strokes for Vector Graphic Synthesis. *arXiv preprint arXiv:2401.17093*.
- Wu, R.; Su, W.; Ma, K.; and Liao, J. 2023a. IconShop: Text-Guided Vector Icon Synthesis with Autoregressive Transformers. *ACM Transactions on Graphics (TOG)*, 42(6): 1–14.
- Wu, X.; Sun, K.; Zhu, F.; Zhao, R.; and Li, H. 2023b. Human preference score: Better aligning text-to-image models with human preference. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2096–2105.
- Xing, X.; Wang, C.; Zhou, H.; Zhang, J.; Yu, Q.; and Xu, D. 2024a. Diffsketcher: Text guided vector sketch synthesis through latent diffusion models. *Advances in Neural Information Processing Systems*, 36.
- Xing, X.; Zhou, H.; Wang, C.; Zhang, J.; Xu, D.; and Yu, Q. 2024b. SVGDreamer: Text Guided SVG Generation with

Diffusion Model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4546–4555.

Xu, T.; Zhang, P.; Huang, Q.; Zhang, H.; Gan, Z.; Huang, X.; and He, X. 2018. AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1316–1324.