

Title	Drawing Guidance for Freehand Sketching with Interactive Diffusion Model
Author(s)	Chen, Chuang; Xie, Xiaoxuan; Zhang, Yongming; Zhang, Tianyu; Xie, Haoran
Citation	研究報告コンピュータグラフィックスとビジュアル情報学 (CG), 2024-CG-195(2): 1-6
Issue Date	2024-09-02
Type	Journal Article
Text version	publisher
URL	http://hdl.handle.net/10119/19346
Rights	<p>社団法人情報処理学会, Chuang Chen, Xiaoxuan Xie, Yongming Zhang, Tianyu Zhang, Haoran Xie, 情報処理学会研究報告. コンピュータグラフィックスとビジュアル情報学 (CG), 2024-CG-195 (2), 2024, pp.1-6. ここに掲載した著作物の利用に関する注意: 本著作物の著作権は(社)情報処理学会に帰属します。本著作物は著作権者である情報処理学会の許可のもとに掲載するものです。ご利用に当たっては「著作権法」ならびに「情報処理学会倫理綱領」に従うことをお願いいたします。 Notice for the use of this material: The copyright of this material is retained by the Information Processing Society of Japan (IPSJ). This material is published on this web site with the agreement of the author (s) and the IPSJ. Please be complied with Copyright Law of Japan and the Code of Ethics of the IPSJ if any users wish to reproduce, make derivative work, distribute or make available to the public any part or whole thereof. All Rights Reserved, Copyright (C) Information Processing Society of Japan.</p>
Description	コンピュータグラフィックスとビジュアル情報学研究会第195回研究発表会



Drawing Guidance for Freehand Sketching with Interactive Diffusion Model

CHUANG CHEN^{1,2} XIAOXUAN XIE^{1,2} YONGMING ZHANG¹ TIANYU ZHANG¹ HAORAN XIE¹

Abstract: Drawing is a universal form of human expression that transcends language barriers. However, the common users may have difficulty to convey their ideas independently due to the necessity of consistent practice and professional guidance. In this work, we proposed our system for freehand sketching. The proposed system adopts an interactive diffusion model, StreamDiffusion for providing real-time user guidance to assist users within the drawing process. We use a diffusion model to generate RGB images based on hand-drawn sketches and prompts, then convert RGB images into sketches using informative drawings line detector to assist users with creative expression in the drawing practice. Our method can generate images in real time that accurately align with the user's intent and provide shadow guidance in the background. This process aids users in their drawing practice by allowing them to create sketches based on prompts and their own coarse drawings. A user study evaluates our method's effectiveness and gathers feedback on interface design.

Keywords: Drawing guidance, diffusion model, freehand sketch

1. Introduction

Drawing is an innate human instinct that allows us to convey compelling stories without words. Regardless of the subjects, a good drawing may require advanced drawing skills and extensive experience. Taking portrait drawing as an example, the users need not only skillful experience in observing facial features but also the ability to draw details, such as expressions, facial proportions, age-specific features, and the effects of light and shadow.

Drawing guidance systems aim to help individuals without advanced drawing skills achieve their creative goals while improving their drawing capabilities. Previous works [3, 14] utilized data retrieval approach to dynamically retrieve sketches from a database that closely matched the user's hand-drawn sketch. Additionally, PortraitSketch [20] corrected the user's hand-drawn sketches by allowing users to upload images and make adjustments based on the traced lines, resulting in more aesthetically pleasing drawings. In recent years, generative models have been adopted extensively in user drawing guidance [4]. However, these previous systems typically rely on pre-trained datasets and the user's drawing ability, which often limits users within the scope of the dataset.

To address these issues, we propose a drawing guidance system based on the diffusion model that generates guidance sketches in real-time. This proposed approach allows users creatively drawing. For novices, they can describe their desired sketch through prompts, enabling more accurate generation of the content they envision but lack the ability to create themselves. During the




Fig. 1: The proposed system can generate guidance sketches based on the user's hand-drawn sketches and prompts in real-time. When the user selects the satisfied guidance image, the shadow guidance is adopted for tracing.

drawing process, we provide four dynamically changing guidance sketches positioned separately from the canvas to avoid disrupting users' creative freedom. Users can place these sketches beneath the canvas for tracing or remove them to continue their creation. Additionally, we offer drawing tools such as pens and erasers and include undo and redo functions to facilitate the creative process.

In summary, our contributions are listed as follows:

- We propose a real-time drawing guidance system where users can input incomplete sketches and prompts. The system generates the intended content and places it in the background layer to aid further creation.
- We introduce a new guidance method, in which the user can freely control guidance sketches. A user study demonstrates that the proposed method can help users easily achieve their desired sketches.

2. Related Works

2.1 Drawing Guidance System

The drawing guidance system not only enables users to expe-

¹ Japan Advanced Institute of Science and Technology, Nomi, Ishikawa 923-1292, Japan

² Authors with the same contribution.

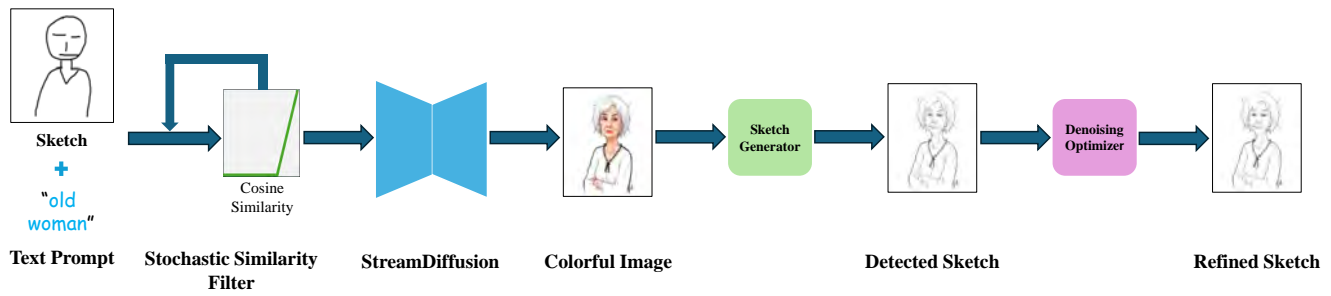


Fig. 2: The pipeline of our interface. The drawing process begins with a hand-drawn sketch and a text prompt. The stochastic similarity filter first evaluates the similarity between the input sketches to avoid generating a similar image. The Sketch Generator then extracts the detected sketch from the RGB image, which often contains some noise. The denoising optimizer finally refines the detected sketch to produce a clean, refined sketch.

rience the joy of drawing and boosts their confidence, but also effectively enhances their drawing abilities. Dualface [8] introduces a two-stage drawing guidance interface specifically tailored for freehand portrait sketching. Sketchhelper [2] generates shadow guidance for users by retrieving strokes from the dataset in real-time. EZ-Sketching [18] automatically corrects hand-drawn sketch lines roughly traced over an image by a three-level optimization method. The drawing interface features with grid guidance [11] and variable canvas sizes [19] may have influences on the drawing results. Moreover, systems like AniFaceDrawing [9] propose an unsupervised stroke-level decoupled training strategy that enables sparse sketches to be automatically matched to corresponding localizations of anime portraits. However, these systems predominantly focus on facial drawing, thereby limiting the user's creativity. To address this limitation, we propose a system capable of drawing objects in real-time. Our system significantly expands users' creative freedom and enhances interactive experiences.

2.2 Sketch-based Applications

Sketches can quickly and intuitively express creative ideas and design concepts, playing a crucial role in fields such as art design and engineering modeling. Sketch2Model [22] converts hand-drawn sketches into 3D models, providing users with a convenient tool for 3D design. Additionally, Pix2Pix [10] transforms sketches into realistic images, which are widely utilized across various design and creative disciplines. Moreover, Sketch-RNN [5] introduces machine learning techniques to intelligently and automatically generate sketches. In summary, these methods not only enhance the efficiency of sketch utilization in the design process but also broaden their application scope, positioning sketches as pivotal tools in numerous fields.

2.3 Diffusion Models

The advancement of diffusion models has propelled the field of image generation, demonstrating remarkable image quality and detail. Stable Diffusion [16] has excelled in image generation with its realistic resolution and flexible input methods, swiftly becoming a popular choice in the field. ControlNet [21] enhances the image generation quality of models like Stable Diffusion by introducing spatial conditioning, overcoming spatial layout con-

straints, and enabling users to create more complex and detailed images. Additionally, StreamDiffusion [13] achieves real-time, efficient image generation through optimized computational processes and parallel processing techniques. LoRA [7] enhances the practicality and application scope of diffusion models through its lightweight network structure and efficient training methods. These methods enable sketch-based drawing guidance systems to offer users a smoother, higher-quality creative experience, paving the way for interactive drawing support.

3. Proposed Method

This work aims to generate guidance sketches from the user's hand-drawn sketches in real-time. Figure 2 illustrates the pipeline of our system, which consists of three main components: stochastic similarity filter, image generation, sketch generation, and optimization. In the drawing guidance generation process, the user's current hand-drawn sketch and prompts are first input into the similarity filter (Section 3.2). The filtered input is then processed by the StreamDiffusion (Section 3.3) to generate an RGB image. This image is subsequently refined by the sketch generator (Section 3.4) to produce the final result.

3.1 User Interface

Our interface aims to assist art enthusiasts and novices in drawing practice. For novices who may struggle to effectively convey their ideas, we provide a prompt feature. Users can describe their intended sketch using prompts, which generate guidance sketches tailored to their needs. Users can select a guidance sketch to place it on the background layer of the canvas. They can then remove the background layer and proceed with free drawing. We also offer a range of tools designed to enhance the drawing process. As shown in Figure 3, our interface includes options such as undo, redo, clear canvas, pen, eraser, and save.

In contrast to traditional drawing assistance systems, our method allows users can decide whether to place the guidance sketch on the background layer of the canvas to aid their drawing. Capturing users' intentions is a challenging task. If the generated content does not meet their expectations, it can become a visual distraction. Therefore, we position guidance sketches on the right side of the canvas, allowing users to reference the generated sketches for inspiration without resorting to mere tracing. If




Fig. 3: The proposed interface includes drawing tools such as a pen, an eraser, and other essential tools. Users can flexibly control the position of the guidance sketch to suit their needs by clicking on the guidance sketch or selecting “Continue Drawing”.Users can also enter prompts to convey their intended sketch.

users struggle to draw what they envision, they can choose to fix a preferred guidance sketch on the background layer for tracing, helping to improve their drawing skills.

3.2 Stochastic Similarity Filter

During the hand-drawing process, we skip the extremely approximate hand-drawn image generation stage to minimize the negative user experience caused by the fast switching of the generated image, as the repeated sketching of the same brush stroke. The probability of skipping the image generation step is calculated as follows:

$$P(\text{skip}|x) = \begin{cases} 0 & \text{if } x < T \\ 50 \cdot x - 49 & \text{if } x \geq T \end{cases} \quad (1)$$

where x represents the cosine similarity between the hand-drawn sketch I_i and the guidance sketch I_{ref} . T denotes the threshold of similarity value. In our implementation, $T = 98$ for skipping smaller repetitive strokes and providing a better experience for the user. The system calculates the probability of skipping the image generation process accordingly. Regardless of whether the image generation process is skipped, the current input image is defined as the new guidance sketch. When the prompt changes, regardless of the similarity between hand-drawn sketches, we allow the system to regenerate the image to meet the user’s needs.

$$x = \cos(\theta) = \frac{\mathbf{I}_i \cdot \mathbf{I}_{\text{ref}}}{\|\mathbf{I}_i\| \|\mathbf{I}_{\text{ref}}\|} \quad (2)$$

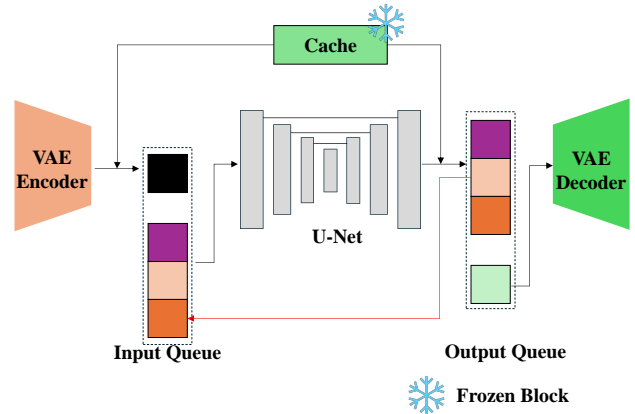


Fig. 4: The framework of StreamDiffusion including VAE Encoder, U-Net, and VAE Decoder components.

3.3 Image Generation

Real-time performance is a critical factor affecting the interactive experience in the guidance interfaces. To provide users with more options, we configured StreamDiffusion pipelines [13] that can generate images synchronously for user guidance. We address the trade-off between data input and model throughput by employing input-output queue strategies and batch processing.

Hand-drawn sketches are encoded into tensor form via a Variational Autoencoder (VAE) [12] and then enter the input queue, awaiting denoising by the U-Net [17] model. During denoising, the hand-drawn sketch tensors are processed sequentially by U-Net, with each batch undergoing a denoising cycle. The par-

allel processing capabilities of the GPU allow multiple images to be denoised simultaneously, significantly enhancing system throughput compared to traditional sequential methods. After denoising, the tensors are moved to the output queue, where they are decoded back into sketches by the VAE. Additionally, the integration of LoRA [7] ensures smooth real-time interaction within the system.

To enhance system efficiency, we adopted caching mechanisms, including Noise Cache, Scheduler Cache, and Prompt Embed Cache. These mechanisms can improve system performance as follows:

- **Noise Cache:** Gaussian noise is pre-sampled and stored for each denoising step. This ensures that noise varies between steps but remains consistent within each step (i.e., $y_{t,\tau} \neq y_{t,\tau+1}$ and $y_{t+1,\tau} = y_{t,\tau}$). This consistency is crucial for image-to-image tasks that require stable noise across different time steps.

$$y_{t,\tau} = \sqrt{\gamma_\tau}x_0 + \sqrt{\delta_\tau}\epsilon \quad (3)$$

- **Scheduler Cache:** The noise strength coefficients γ_τ and δ_τ are pre-computed and stored. Caching these static values reduces recomputation overhead, thus improving processing speed, particularly at high frame rates (e.g., over 60 frames per second).

$$\hat{x}_{0,\tau-1,cf} = \frac{y_{\tau-1,cf} - \sqrt{\delta_{\tau-1}}\epsilon_{\tau-1,cf}}{\sqrt{\gamma_{\tau-1}}} \quad (4)$$

$$y_{\tau,cf} = \sqrt{\gamma_\tau}\hat{x}_{0,\tau-1,cf} + \sqrt{\delta_\tau}\epsilon_\tau \quad (5)$$

- **Prompt Embed Cache:** Prompt embeddings are pre-computed and stored in the cache. In interactive or streaming modes, these embeddings are utilized for Key and Value computations in the U-Net. This approach avoids the redundant computation of prompt embeddings during each inference, enhancing efficiency.

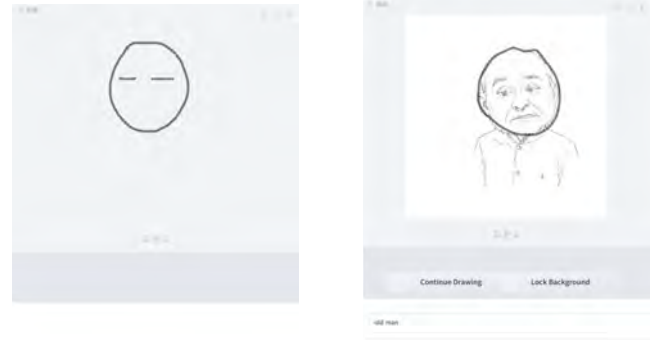
By incorporating these caching mechanisms, we effectively reduce redundant computations and optimize the performance of the StreamDiffusion pipeline, enabling high-efficiency, low-latency image generation in real-time applications.

3.4 Sketch Generation and Optimization

We use the informative drawings model [1] for converting images into sketches. The process begins by passing the input image through an encoder-decoder architecture with several ResNet [6] blocks to extract and process features. The initial line drawing is then subjected to depth prediction to compute geometric loss, ensuring accurate delineation of geometrically important areas. Simultaneously, the line drawing is processed with CLIP [15] to create embeddings, which are compared with the CLIP embeddings of the original image to compute semantic loss. This ensures that the generated sketch effectively conveys the semantic content of the original image. This approach allows for the efficient conversion of images into sketches.

To further refine the generated sketches, we have designed an optimizer that employs a recursive filter, an efficient image-smoothing technique. This filter achieves two-dimensional smoothing by repeatedly applying a one-dimensional filter. This approach effectively removes noise and unwanted details while

preserving the image's edge information. As a result, the final sketches presented to users are smoother and more refined, enhancing the overall drawing experience.



(a)

(b)

Fig. 5: Drawing interfaces used in our study. (a) Baseline interface: a canvas without guidance assistance. (b) Shadow guidance interface: interface with one guidance sketch fixed under canvas.

4. User Study

To evaluate the effectiveness of our proposed drawing assistance system and the impact of different guidance placements, we conducted a user study. We invited 8 participants (including 1 female and 7 males) to join in this user study. Among them, 2 participants had art training, and 6 participants had no formal art training.

We set up three experimental groups for comparison: the baseline interface (Figure 5a), the shadow guidance interface (Figure 5b), and our guidance interface. The shadow guidance interface fixes the guidance sketch at the bottom background layer of the canvas, while our interface places the guidance sketches on the right side of the canvas, allowing users to freely choose whether to display the guidance sketch in the background layer. Additionally, both the shadow guidance interface and ours are consistent in terms of image generation quality and speed. Since the study involved three different drawing interfaces, we first conducted the control experiment using the interface without any drawing assistance. To avoid any order bias, the sequence in which the remaining two drawing assistance interfaces were used was randomized. All participants used a graphics tablet (WACOM, drawing area 512×512) for portrait creation.

Before the formal experiment began, participants were allowed to freely use all three drawing interfaces to familiarize themselves with their features and functions. After the user became familiar with the interfaces, they selected three themes from the provided options and planned the content for each theme. Participants were instructed to follow this sequence: first, using the baseline interface, and then, in a randomized order, using the shadow guidance interface and our guidance interface, drawing the same theme content in each case. After completing one theme, they proceeded to draw the next one. We timed their drawing from the moment they began, allowing them to fully utilize the interface features and include as much detail as possible. The timer was stopped when participants considered their drawings complete.

Table 1: The result of the System Usability Scale (SUS)

#	Question	Mean	SD
1	I think that I would like to use this system frequently.	4.61	0.64
2	I found the system unnecessarily complex.	3.20	1.40
3	I thought the system was easy to use.	4.61	0.64
4	I think that I would need the support of a technical person to be able to use this system.	2.88	1.36
5	I found the various functions in this system were well integrated.	4.68	0.47
6	I thought there was too much inconsistency in this system.	2.75	1.61
7	I would imagine that most people would learn to use this system very quickly.	4.56	0.50
8	I found the system very cumbersome to use.	2.87	1.73
9	I felt very confident using the system.	4.56	0.50
10	I need to learn a lot of things before I could get going with this system.	3.13	1.75




Fig. 6: Some results of the user study. Users initially drew on the baseline interface, with results shown in the "Baseline" row. Subsequently, users randomly chose between our interface and the shadow guidance interface. The results are displayed in the "Shadow Guidance" and "Ours" rows respectively.

After completing all the drawing tasks, participants filled out a questionnaire. The questionnaire focused on several aspects: the ease of use of the interface, the effectiveness of the drawing assistance, and which type of drawing assistance was most helpful in improving the users' drawings. Our questionnaire consists of two parts: the System Usability Scale (SUS) and a user preference experiment.

5. Results

We analyzed the collected online questionnaire results and users' drawings to explore the effectiveness of our designed drawing assistance system and to assess which type of assistance is more helpful in improving users' drawing outcomes.

5.1 User Experience Study

Figure 6 shows some examples of drawings created by users using a baseline interface, shadow guidance interface, and our interface. Both the Shadow guidance interface and our interface can infer the user's intentions from incomplete hand-drawn sketches and prompts, generating guidance sketches to assist further draw-

ing. The key difference is that the shadow guidance interface provides only one guidance sketch, which is fixed at the background layer of the canvas, whereas our interface offers four guidance sketches placed on the right side of the canvas. Users can choose whether to fix these reference sketches in the background of the canvas according to their needs. Compared to the baseline interface, our interface significantly enhances the quality of the user's drawings. Visually, users utilizing our interface can create more refined images with greater detail.

We adopted the standard of the System Usability Scale (SUS) to evaluate our interface in the questionnaire (1 for strongly disagree and 5 for strongly agree). For odd-numbered questions, the score was subtracted from 1, and for even-numbered questions, the score was subtracted from 5; the converted scores were then summed and multiplied by 2.5. The calculated SUS score for our method was 70.48, indicating that our method performs well in terms of usability and provides effective drawing assistance.

5.2 User Preference Experiment

The way in which the guidance sketch is provided can have a

significant impact on users' creative processes. To explore this, we conducted a survey in the questionnaire to determine which method of providing the guidance sketch is most effective. In the survey, we asked users which of the three interfaces helped them best achieve their initial drawing intentions. The results showed that 75% of users chose our interface, 25% chose the shadow guidance interface, and none chose the baseline interface. For the question "Which interface do you think is better for you to draw?", 87.5% of respondents selected our interface, 12.5% chose the shadow guidance interface, and 0% chose the baseline interface. For question "Which is your desired outcome in the three groups?", 62.5% of users selected our interface, while 37.5% chose the shadow guidance interface. The experimental results demonstrate that our drawing assistance method is more positively received by users. When users draw freely, the absence of a guidance sketch in the background layer allows them to focus more on expressing their ideas without being influenced by the guidance sketch.

6. Conclusion

This work proposed a drawing guidance system that can help users draw using the diffusion model, which generates the corresponding RGB image by using the simple user input. The resulting sketch lines are extracted to assist the user in drawing. According to the use study, our system can assist users with no drawing foundation at all, and it is better to draw sketches in line with their ideas. At the same time, our system can assist users who already have a certain painting foundation in carrying out more creative painting work.

Our work still has some limitations. First, some participants pointed out that the reference images generated each time are independent, sometimes the subsequent generation can rely on the current generation's results, which will offer the user more reasonable tips. Also, it would become convenient to adjust the transparency of the guidance sketch image. Finally, the current guidance system relied heavily on pre-trained models, there may be a certain difference from the user's expected drawing intentions.

Acknowledgement

We greatly thank the participant for joining our user study. This research has been supported by the Kayamori Foundation of Informational Science Advancement.

References

- [1] Chan, C., Durand, F. and Isola, P.: Learning to generate line drawings that convey geometry and semantics, *CVPR* (2022).
- [2] Choi, J., Cho, H., Song, J. and Yoon, S. M.: SketchHelper: Real-Time Stroke Guidance for Freehand Sketch Retrieval, *IEEE Transactions on Multimedia*, Vol. 21, No. 8, pp. 2083–2092 (online), DOI: 10.1109/TMM.2019.2892301 (2019).
- [3] Eitz, M., Richter, R., Hildebrand, K., Boubekeur, T. and Alexa, M.: Photosketcher: Interactive Sketch-Based Image Synthesis, *IEEE Computer Graphics and Applications*, Vol. 31, No. 6, pp. 56–66 (online), DOI: 10.1109/MCG.2011.67 (2011).
- [4] Ghosh, A., Zhang, R., Dokania, P. K., Wang, O., Efros, A. A., Torr, P. H. S. and Shechtman, E.: Interactive Sketch Fill: Multiclass Sketch-to-Image Translation, *Proceedings of the IEEE international conference on computer vision* (2019).
- [5] Ha, D. and Eck, D.: A Neural Representation of Sketch Drawings (2017).
- [6] He, K., Zhang, X., Ren, S. and Sun, J.: Deep Residual Learning for Image Recognition, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016).
- [7] Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L. and Chen, W.: Lora: Low-rank adaptation of large language models, *arXiv preprint arXiv:2106.09685* (2021).
- [8] Huang, Z., Peng, Y., Hibino, T., Zhao, C., Xie, H., Fukusato, T. and Miyata, K.: dualface: Two-stage drawing guidance for freehand portrait sketching, *Computational Visual Media*, Vol. 8, pp. 63–77 (2022).
- [9] Huang, Z., Xie, H., Fukusato, T. and Miyata, K.: AniFaceDrawing: Anime Portrait Exploration during Your Sketching, *ACM SIGGRAPH 2023 Conference Proceedings, SIGGRAPH '23* (2023).
- [10] Isola, P., Zhu, J.-Y., Zhou, T. and Efros, A. A.: Image-to-Image Translation with Conditional Adversarial Networks (2018).
- [11] Kanayama, H., Xie, H. and Miyata, K.: Illustration Drawing Interface with Image Retrieval and Adjustable Grid Guidance, *2023 Nicograph International (NicoInt)*, pp. 54–61 (online), DOI: 10.1109/NICOINT59725.2023.00018 (2023).
- [12] Kingma, D. P. and Welling, M.: Auto-Encoding Variational Bayes (2022).
- [13] Kodaira, A., Xu, C., Hazama, T., Yoshimoto, T., Ohno, K., Mitsuhori, S., Sugano, S., Cho, H., Liu, Z. and Keutzer, K.: StreamDiffusion: A Pipeline-level Solution for Real-time Interactive Generation (2023).
- [14] Lee, Y. J., Zitnick, C. L. and Cohen, M. F.: ShadowDraw: real-time user guidance for freehand drawing, *ACM Trans. Graph.*, Vol. 30, No. 4 (online), DOI: 10.1145/2010324.1964922 (2011).
- [15] Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J. et al.: Learning transferable visual models from natural language supervision, *International conference on machine learning*, PMLR, pp. 8748–8763 (2021).
- [16] Rombach, R., Blattmann, A., Lorenz, D., Esser, P. and Ommer, B.: High-Resolution Image Synthesis with Latent Diffusion Models (2022).
- [17] Ronneberger, O., Fischer, P. and Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation (2015).
- [18] Su, Q., Li, W. H. A., Wang, J. and Fu, H.: EZ-sketching: three-level optimization for error-tolerant image tracing, *ACM Trans. Graph.*, Vol. 33, No. 4 (online), DOI: 10.1145/2601097.2601202 (2014).
- [19] Wang, J., Nakano, K., Chen, D., Huang, Z., Fukusato, T., Miyata, K. and Xie, H.: A Study on Cognitive Effects of Canvas Size for Augmenting Drawing Skill, *2024 Nicograph International (NicoInt)*, pp. 49–53 (online), DOI: 10.1109/NICOInt62634.2024.00018 (2024).
- [20] Xie, J., Hertzmann, A., Li, W. and Winnemöller, H.: PortraitSketch: face sketching assistance for novices, *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology, UIST '14*, New York, NY, USA, Association for Computing Machinery, p. 407–417 (online), DOI: 10.1145/2642918.2647399 (2014).
- [21] Zhang, L., Rao, A. and Agrawala, M.: Adding Conditional Control to Text-to-Image Diffusion Models (2023).
- [22] Zhang, S.-H., Guo, Y.-C. and Gu, Q.-W.: Sketch2model: View-aware 3d modeling from single free-hand sketches, *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6012–6021 (2021).