

AutoFigure-Edit: Generating Editable Scientific Illustrations via Reference-Guided Styling

Zhen Lin^{1*}, Qiujie Xie^{2,1*}, Minjun Zhu^{2,1*},
Shichen Li^{1,3}, Qiyao Sun¹, Enhao Gu¹, Yiran Ding¹, Ke Sun¹,
Fang Guo¹, Panzhong Lu¹, Zhiyuan Ning¹, Yixuan Weng^{1†}, Yue Zhang^{1✉}

¹Engineering School, Westlake University, ²Zhejiang University, ³Soochow University

*: Equal contribution †: Project leader.

Correspondence author✉: zhangyue@westlake.edu.cn

Abstract

High-quality scientific illustrations are essential for communicating complex scientific and technical concepts, yet existing automated systems remain limited in editability, stylistic controllability, and efficiency. We present AUTOFIGURE-EDIT, an end-to-end system that generates fully editable scientific illustrations from long-form scientific text while enabling flexible style adaptation through user-provided reference images. By combining long-context understanding, reference-guided styling, and native SVG editing, it enables efficient creation and refinement of high-quality scientific illustrations. To facilitate further progress in this field, we release the video¹, full codebase² and provide a live demo for easy access and interactive use³.

1 Introduction

Creating high-quality scientific illustrations often takes researchers several days and demands both substantial domain expertise and professional-level design skills (Huang et al., 2026b). It requires a rigorous, logic-aware understanding of **long-form scientific texts** (>10k tokens), while the visual rendering must carefully balance **structural fidelity** and **image quality** to produce figures that are clear, accurate, and aesthetically appealing (Chang et al., 2025; Zhu et al., 2026b).

Research on **automatically generating scientific illustrations** from long-form scientific texts remains limited. Code-as-intermediate approaches (Belouadi et al., 2024a,b, 2025; Ellis et al., 2018) achieve strong geometric correctness, but often sacrifice visual aesthetics and readability (Zhu et al., 2026b). Meanwhile, end-to-end mainstream Text-to-Image (T2I) models can produce visually appealing illustrations yet frequently

fail to maintain **structural fidelity** on long scientific inputs (Liu et al., 2025; Huang et al., 2026a). Therefore, directly transforming long scientific texts into illustrations that are both **accurate** and **visually compelling** remains challenging.

In our previous work, we introduced AUTOFIGURE (Zhu et al., 2026b), an agentic framework grounded in the Reasoned Rendering paradigm that produces accurate and visually appealing illustrations through an iterative refinement process. Despite its ability to automatically generate high-quality illustrations, AUTOFIGURE has several limitations: (i) the generated visual elements are **fixed and non-editable**. Refinement can only be performed by modifying the user-provided textual prompt; (ii) generating illustrations in a desired style relies heavily on **prompt engineering**, which can be ambiguous and may lead to inaccurate or unintended stylistic outcomes; and (iii) its iterative sketch-to-render refinement **tightly couples layout planning with final raster rendering**, without exposing an explicit structural scaffold. As a result, fine-grained edits (e.g., layout adjustments) are difficult and text rendering is often unstable.

To address these limitations, we present a substantially enhanced system, named AUTOFIGURE-EDIT, that transforms long-form scientific text and a reference style image into a **fully editable** SVG illustration. It enables **reference-guided style control**, reducing reliance on ambiguous prompt engineering, and decouples layout planning from final rendering via an **explicit structural scaffold**, allowing layout edits directly on the vector scaffold without retrying the full sketch-to-render loop. Building on this design, AUTOFIGURE-EDIT provides the following features:

Scientific Illustration Generation. Directly transforms long-form scientific text into accurate, publication-quality illustrations.

Reference-Guided Style Control. Enables controllable visual adaptation via a user-provided ex-

¹<https://youtu.be/10IH8SyJjAQ>

²<https://github.com/ResearAI/AutoFigure-Edit>

³<https://deepscientist.cc/>

Method	Sci. Gen.	Editable	Style Control	GUI Support
StarVector & OmniSVG (Rodriguez et al., 2025; Yang et al., 2025)	Limited	✓	Prompt (Hard)	Web Only
AutomaTikZ (Belouadi et al., 2024a)	✓	✓	Prompt (Hard)	Web Only
DeTikZify (Belouadi et al., 2024b)	✓	✓	Sketch (Hard)	Web Only
GPT-Image (Hurst et al., 2024)	✓	✗	Prompt (Hard)	Web Only
Diagram Agent (Wei et al., 2025)	✓	✓	Prompt (Hard)	None
PaperBanana (Zhu et al., 2026a)	✓	✗	Prompt (Hard)	Web Only
EditBanana (BIT-DataLab, 2026)	✗	✓	✗	Web Only
SciFig & SciSketch (Huang et al., 2026b; Wang et al., 2025)	✓	✓	Prompt (Hard)	None
VisPainter (Sun et al., 2025)	✓	✓	Prompt (Hard)	None
AUTOFIGURE (Zhu et al., 2026b)	✓	✗	Prompt (Hard)	Web Only
AUTOFIGURE-EDIT (Ours)	✓	✓	Reference (Easy)	Web + Editor

Table 1: Comparison of AUTOFIGURE-EDIT with other relevant systems. “Sci. Gen.” denotes scientific illustration generation; “GUI Support” denotes the level of integrated graphical user interface and editing capabilities.

emplar while preserving semantic structure.

Editable SVG with Embedded Visual Editor.

Produces structurally organized, component-level editable SVGs and supports real-time refinement through an integrated interactive canvas.

Quantitative experiments and user studies (Section 4) demonstrate the effectiveness and practical value of AUTOFIGURE-EDIT in generating high-quality, editable scientific illustrations.

2 Related Work

Automation of scientific illustration has evolved from simple summarization to complex synthesis. However, achieving a balance between generation quality and post-generation editability remains challenging. Existing text-to-figure systems (Zhu et al., 2026a; Huang et al., 2026b; Zhu et al., 2026b) can automatically create high-quality illustrations from long textual descriptions, but they typically produce static outputs with limited support for iterative refinement, requiring full regeneration for even minor adjustments. To improve customizability, recent approaches convert rasterized renderings into vector representations. Nevertheless, their reliance on pre-rendered pixel inputs can lead to semantic information loss Sun et al. (2025). Meanwhile, editing tools such as EditBanana (BIT-DataLab, 2026) provide post-hoc modification capabilities, relying on externally provided images as inputs. In contrast, **AUTOFIGURE-EDIT offers a unified, end-to-end pipeline** that not only generates illustrations from scratch but also represents all components as fully editable objects, enabling precise control.

Programmatic Synthesis. Despite the proficiency of diffusion models (Saharia et al., 2022) in general visual synthesis, their limited structural transparency make them ill-suited to the strict compositional constraints of scientific figures. To en-

hance controllability, recent work has explored Text-to-Code-to-Image pipelines that use programmatic representations (e.g., TikZ or SVG) as intermediate forms (Belouadi et al., 2024a,b). However, purely programmatic generation is often brittle. Small syntax errors can trigger rendering failures, and the absence of an intuitive visual editing interface increases the effort required for iterative refinement. AUTOFIGURE-EDIT addresses these issues by combining long-form context understanding with robust structural reconstruction, striking better balances between stylistic flexibility and editability. The comparison with relevant systems is shown in Table 1.

3 AutoFigure-Edit

We introduce AUTOFIGURE-EDIT (Figure 1), an automated system that transforms long-form scientific text into structured, fully editable scientific illustrations and supports flexible style adaptation via user-provided reference images.

3.1 Framework Overview

The task of automated scientific illustration generation involves reconciling three competing goals: (i) **semantic faithfulness** to the text; (ii) **stylistic consistency** with the reference image; (iii) explicit structural decomposition to support **downstream editing**. Formally, given a long-form scientific text T and a reference style image I^{ref} , the objective is to learn a mapping:

$$S^* = \mathcal{F}(T, I^{\text{ref}}),$$

where S^* is an editable vector graphic that preserves the semantics of T while conforming to the visual style of I^{ref} .

Directly parameterizing this mapping is ill-posed due to the absence of explicit structural supervision and the entanglement of layout, instance, and

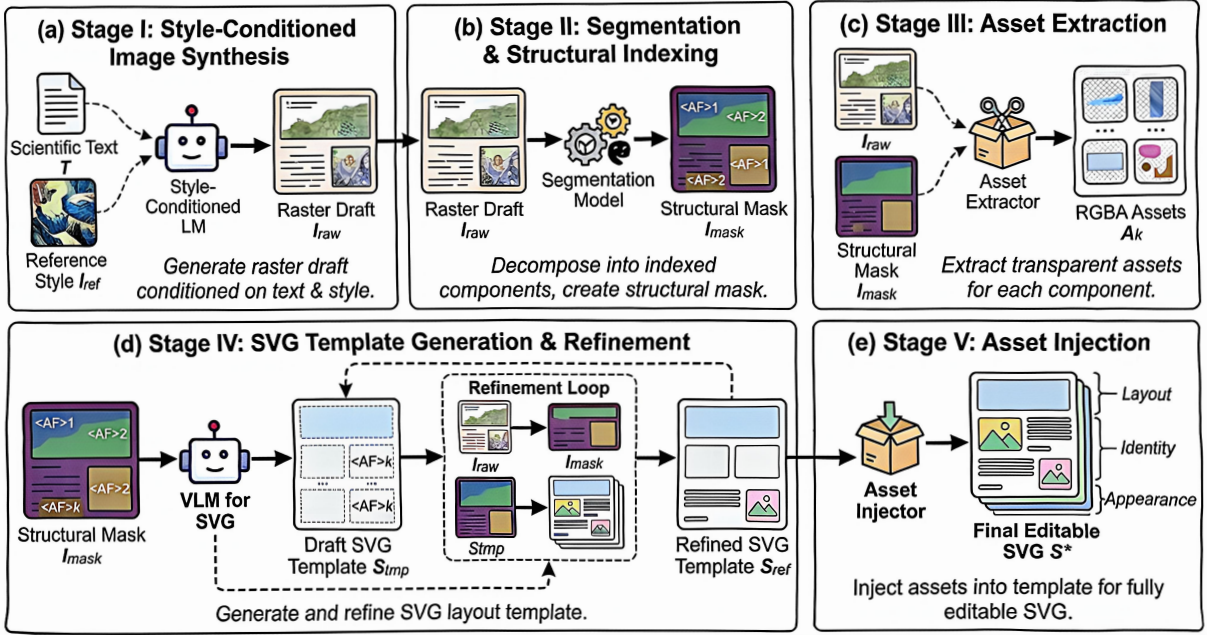


Figure 1: An overview of AUTOFIGURE-EDIT. This figure is also produced by AUTOFIGURE-EDIT and serves as a qualitative showcase of its generation quality.

visual appearance. We therefore **decompose the transformation into sequential stages** that progressively derive structure from an intermediate raster draft, disentangling layout planning, object identity, and visual rendering while preserving semantic and stylistic coherence.

3.1.1 Stage I: Style-Conditioned Image Synthesis

We first generate a raster draft I^{raw} conditioned jointly on the input scientific text and a reference style image, using a style-conditioned text-to-image model (e.g., Gemini-3-Pro-Image-Preview). This stage translates textual descriptions into explicit visual entities while incorporating high-level stylistic cues from the reference image, thereby establishing semantic–stylistic alignment.

3.1.2 Stage II: Segmentation and Structural Indexing

To expose the compositional structure of the raster draft, we then apply instance segmentation (Carion et al., 2025) to **decompose I^{raw} into a set of visual components** $\{M_k, B_k\}_{k=1}^K$, where M_k and B_k denote the mask and bounding box of the k -th component, respectively. Instead of retaining the original appearance of each region, we construct a simplified structural rendering in which every instance is filled with a uniform tone and assigned a unique identifier token (e.g., $\langle \text{AF} \rangle k$). By suppressing texture and color while preserving spatial

configuration and instance identity, this transformation converts the raster image into an **indexed structural layout**, providing a coordinate-aware scaffold for subsequent vector generation.

3.1.3 Stage III: Asset Extraction

While Stage II enables explicit structural information, appearance cues must be preserved to ensure faithful reconstruction. Therefore, for each segmented instance (I^{raw}, B_k) , we extract the corresponding visual content and remove the background (Zheng et al., 2024) to obtain a transparent RGBA asset A_k . This process decouples geometric placement from visual texture, storing appearance as standalone icon-like assets while delegating spatial organization to the structural scaffold. As a result, subsequent layout modifications can be performed without altering stylistic details.

3.1.4 Stage IV: SVG Template Generation and Refinement

Given the indexed structural representation I^{mask} , we then prompt a vision-language model (e.g., Gemini-3.1-Pro-Preview) to generate an **SVG layout template** S^{tmp} containing placeholder elements aligned with the $\langle \text{AF} \rangle k$ identifiers.

To improve alignment with the target figure, we further perform a **lightweight refinement step** by re-prompting the vision-language model with the original raster draft, the structural mask, a rendered preview of the current SVG, and the correspond-

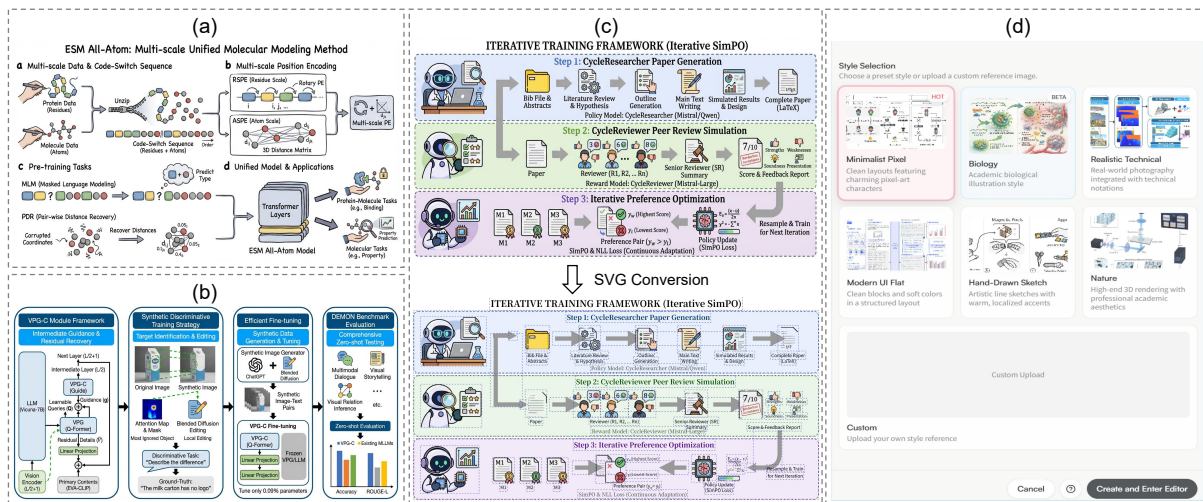


Figure 2: Representative outputs of AUTOFIGURE-EDIT. (a)-(b) are bitmap (PNG) figures generated from long-form scientific descriptions across various domains. (c) shows the PNG-to-SVG conversion case of AUTOFIGURE-EDIT, including the original bitmap (bottom) and its corresponding vectorized SVG result (down). (d) is the web interface of AUTOFIGURE-EDIT, allowing users to select predefined style templates or upload custom reference images.

ing SVG code. The model is instructed to correct discrepancies in two aspects: **positional consistency** (icon placement, text alignment, arrows, lines) and **stylistic consistency** (proportions, fonts, stroke widths, and colors). Identifier mappings and placeholder group structures are preserved to ensure compatibility with subsequent asset injection. In practice, the refinement process requires only 0–2 iterations to obtain a satisfactory template S^{ref} .

3.1.5 Stage V: Asset Injection

Finally, the extracted appearance assets $\{A_k\}_{k=1}^K$ are injected into the refined SVG template S^{ref} by replacing each placeholder with its corresponding asset. This process produces a **fully editable SVG** S^* where layout, object identity, and visual appearance **remain independently manipulable**. Users can subsequently modify geometry, adjust style, or update individual components without disrupting the overall composition.

3.2 Applications

We demonstrate the utility of AUTOFIGURE-EDIT through three representative application scenarios, showcasing its generation capability, stylistic adaptability, and editability. Beyond being a technical system, AUTOFIGURE-EDIT serves as a **productivity tool for researchers** across fields, lowering the barrier to high-quality scientific illustrations.

High-quality Scientific Illustration Generation. The primary use case of AUTOFIGURE-EDIT is the generation of publication-level illustrations directly from long-form scientific texts. Given a

method section or system overview spanning thousands of tokens, the system automatically extracts the key entities, relations, and procedural stages, and transforms them into illustrations that are both **accurate** and **visually appealing**. This capability substantially **reduces the time and expertise** traditionally required to translate dense technical content into clear visual form. In Figure 2 (a)-(c), we present representative examples of the generated results, demonstrating both semantic fidelity to the source text and high visual quality.

Style Adaptation. Given a user-provided reference image, AUTOFIGURE-EDIT can adapt to a **wide range of visual styles**, changing color palettes, typography, icon aesthetics, spacing density, and visual hierarchy while **preserving semantic structure**. Rather than relying on prompt-level style descriptions, the system explicitly conditions on a visual exemplar and transfers high-level stylistic attributes in a controlled manner. This enables users to experiment with multiple visual appearances of the same scientific content, facilitating alignment with venue- or lab-specific visual standards and reducing reliance on manual graphic design expertise. In Figure 2 (d), we illustrate the input interface of AUTOFIGURE-EDIT, where users provide both the source text and a reference style image to guide generation.

Interactive Editing with SVG Output. The generation result of AUTOFIGURE-EDIT is a structurally organized SVG file in which semantic elements (e.g., modules, connectors, annotations) are explicitly represented, enabling **fine-grained**

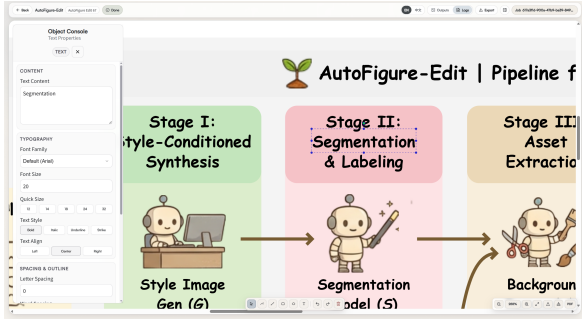


Figure 3: The embedded interactive canvas enables users to freely manipulate individual components within the generated SVG.

manipulation at the component level. This eliminates the common limitation of raster-based generation, where even minor revisions require re-generating the entire image. More importantly, AUTOFIGURE-EDIT further provides **an embedded Visual Editor** that supports real-time updates. Users can reposition objects, modify text, and adjust stylistic attributes while preserving the overall layout. In Figure 3, we present the interface of the embedded interactive canvas, which allows users to directly manipulate individual elements within the generated figure.

In summary, AUTOFIGURE-EDIT transforms scientific illustration generation into an editable and style-controllable process. For individual researchers, AUTOFIGURE-EDIT provides substantial time savings, improved visual clarity, and seamless integration into writing workflows. For the broader community, AUTOFIGURE-EDIT promotes more standardized, accessible, and reproducible scientific communication, enabling clearer dissemination of complex ideas. Additional qualitative results are provided in Figures 5 and 6.

4 Evaluation

To comprehensively evaluate the usability of AUTOFIGURE-EDIT, we conduct (i) automated evaluations on FigureBench (Zhu et al., 2026b) and (ii) a user study involving 217 participants.

4.1 Quantitative Analysis

Experimental Setup. We adopt the research-paper subset of FigureBench (Zhu et al., 2026b) as our evaluation dataset, which contains long-form method sections paired with publication-quality illustrations and provides a realistic testbed for scientific figure generation (Appendix A).

We evaluate on 200 method descriptions sampled from FigureBench: 100 samples are gen-

erated without reference style conditioning, and the remaining 100 are generated with reference style images. The style-conditioned subset is further divided into five groups, where each group shares the same reference image, enabling evaluation of style consistency and robustness under fixed stylistic constraints. Example reference styles are provided in Figure 6. We compare against three categories of baselines, including end-to-end text-to-image generation (GPT-Image (Hurst et al., 2024)), text-to-code generation (HTML-Code and SVG-Code) (Rodriguez et al., 2025; Malashenko et al., 2025; Yang et al., 2024), and multi-agent frameworks (Diagram Agent (Wei et al., 2025) and AUTOFIGURE (Zhu et al., 2026b)). We use NanoBanana-Pro as the text-to-image model and Gemini-3-Pro as the vision-language model for scaffold synthesis and iterative refinement. The quantitative results are summarized in Table 2.

Overall Performance. AUTOFIGURE-EDIT consistently outperforms prior approaches across Visual Design, Communication Effectiveness, and Content Fidelity, demonstrating its ability to **generate publication-quality scientific illustrations and achieving a strong balance** among visual quality and scientific fidelity.

Effect of Reference Conditioning. Reference conditioning reveals a clear trade-off between Visual Design and Content Fidelity. When reference images are provided, Content Fidelity improves consistently across all three sub-dimensions: Accuracy (8.83), Completeness (8.26), and Appropriateness (8.37), surpassing both the original AUTOFIGURE and the non-conditioned variant and suggesting better semantic grounding for long procedural inputs. In contrast, Visual Design slightly drops (e.g., Aesthetic **7.37 vs. 8.32**), likely because fixed references can constrain stylistic expressiveness. Despite this, the Win-Rate increases from **76.0%** to **83.0%**, indicating that reference conditioning yields figures that are more reliably preferred overall. This also suggests that blind pairwise preference is a more holistic and robust indicator than scalar ratings, as it better matches practical selection scenarios.

4.2 User Study

Experimental Setup. To assess the real-world usability of AUTOFIGURE-EDIT, we conduct a **deployment-based user study** via our public web-

Method	Visual Design			Communication		Content Fidelity			Overall	Win-Rate
	Aesthetic	Express.	Polish	Clarity	Flow	Accuracy	Complete.	Appropriate.		
HTML-Code	5.90	5.04	5.84	7.17	7.38	6.99	6.37	6.15	6.35	11.0%
SVG-Code	5.00	4.19	4.89	6.34	6.48	6.15	5.53	5.37	5.49	31.0%
GPT-Image	4.24	3.47	4.00	5.63	5.63	4.77	4.08	4.25	4.51	7.0%
Diagram Agent	2.25	1.73	2.04	2.67	2.49	2.11	1.72	1.94	2.12	0.0%
AutoFigure	7.28	6.99	6.92	7.34	7.87	6.96	6.51	6.40	7.03	53.0%
AUTOFIGURE-EDIT(w/o Ref.)	8.32	8.66	8.16	8.10	8.51	8.59	8.07	7.95	8.29	76.0%
AUTOFIGURE-EDIT(with Ref.)	7.37	7.14	7.37	8.15	8.43	8.83	8.26	8.37	7.99	83.0%

Table 2: Quantitative comparison of illustration generation methods on FigureBench. Scores are averaged across Visual Design, Communication Effectiveness, and Content Fidelity dimensions. Overall denotes the mean of all sub-metrics, and Win-Rate reflects blind pairwise human preference.




Figure 4: Results of the human user study. The numbers indicate the mean scores. AUTOFIGURE-EDIT achieves consistently high satisfaction in most metrics.

site⁴. Users can freely generate editable scientific illustrations and refine the resulting SVGs using the embedded visual editor. Feedback is collected through an integrated interface: once generation completes, a rating dialog is automatically shown, where users evaluate both the rendered figure and its corresponding SVG. All scalar metrics are rated on a 5-point Likert scale (1 = lowest, 5 = highest). We additionally collect a binary *Usability* metric, indicating whether the figure is directly usable in an academic paper without major modifications. Additional details are provided in Appendix B.

Result Analysis. We collect 262 evaluation samples from 217 unique participants. The aggregated results and rating distributions are illustrated in Figure 4. For the generated PNG figures, **AUTOFIGURE-EDIT achieves strong performance across all evaluation dimensions**, with mean scores of 4.04 (Scientific Semantic Correctness), 4.11 (Information Completeness), 3.95 (Visual Presentation Quality), and 4.09 (Style Consistency). Notably, **ratings are heavily concentrated at the highest level**: 48% and 51% of evaluations assign Score 5 to semantic correctness and completeness, respectively, and 50% to style consistency. Low ratings (Score 1–2) are rare for semantic dimensions (generally below 12%), demon-

strating that the system reliably preserves scientific meaning and structural integrity across diverse user inputs.

Practical usability further confirms the system’s effectiveness. **126** out of 262 users consider the generated figure directly suitable for inclusion in an academic paper without additional modification. Given that direct usability requires conceptual correctness, satisfactory layout and stylistic quality, this result **highlights the system’s readiness for real-world research workflows** rather than merely benchmark-level adequacy.

For PNG-to-SVG reconstruction, the average Conversion Correctness score reaches 3.60, with the majority of evaluations concentrated in the upper-middle to high range (Scores 3–5) and 36% achieving Score 5. Very low scores remain limited, indicating that catastrophic structural failures are uncommon. While minor geometric deviations may occasionally arise during reconstruction, the fully editable SVG output ensures that such issues can be corrected with minimal manual effort, thereby preserving downstream usability.

Overall, the empirical results show that AUTOFIGURE-EDIT achieves strong semantic fidelity, high informational completeness, and substantial real-world usability in deployment, suggesting it can be effectively integrated into academic figure production workflows.

5 Conclusion

In this paper, we presented AUTOFIGURE-EDIT, an end-to-end system that generates fully editable scientific illustrations from long-form text with reference-guided style adaptation and native SVG editing. Quantitative evaluations and a deployment-based user study showed that AUTOFIGURE-EDIT consistently outperformed prior methods and produced outputs that users frequently judged ready for academic publication.

⁴<https://deepschientist.cc/>

Acknowledgments

This work is supported by a project from Westlake Institute for Advanced Study under Grant Number 210230006022314, the National Natural Science Foundation of China Key Project under Grant Number 625B2152 (Research on Evaluation Paradigms and Closed-Loop Optimization Mechanisms for the Automated Discovery of AI Algorithms), and the 2025 Dean’s Special "PhD Student Project" (WU2025D007) of the School of Engineering, Westlake University.

Limitations

While our deployment and user study highlight the practical utility of AUTOFIGURE-EDIT in scientific illustration generation, some limitations remain:

Dependence on foundation models. Our pipeline currently relies on closed-source vision and vision-language models (e.g., Gemini-3.1-Pro-Preview and NanoBanana-Pro) for style-conditioned raster synthesis and SVG template refinement. This reliance may incur usage costs, raise data privacy concerns, and limit the reproducibility of our pipeline. While current open-source alternatives struggle with the complex spatial reasoning required for this specific task, future work will explore integrating powerful open-weight models to mitigate these accessibility limitations once their capabilities sufficiently mature.

Error propagation. As AUTOFIGURE-EDIT derives vector structures from an intermediate raster draft, upstream segmentation errors (e.g., incorrectly merged split visual components) can cascade through the pipeline, requiring manual adjustments via the embedded editor.

Scope and evaluation constraints. The embedded visual editor is designed for localized, component-level refinements and is not intended to replace comprehensive graphic design software. Furthermore, our current user study was primarily intended as a usability evaluation to assess real-world workflow efficiency. Validating the system across a wider variety of highly specialized scientific domains and conducting rigorous expert-only correctness checks remains an area for future work.

Failure-mode breakdown. Our inspection of representative SVG conversion failures suggests that remaining errors mainly arise from three sources. First, segmentation and asset extraction may fail when SAM-3 selects regions inconsistent with the text prompt, or when the RMBG-

2.0 background remover mistakenly removes non-background icon details. Second, structural template generation remains challenging for complex non-axis-aligned figures, such as molecular structures in chemistry, and may also mis-handle relational cues such as arrow directions or connector endpoints. Third, structural indexing and asset injection can be affected by small or dense asset labels, leading to incorrect <AF>k associations and asset-template mismatch. Future work will investigate more robust segmentation, background removal, and vision-language models to gradually mitigate these failure modes.

We hope that our system will set a new standard for automated figure generation workflows, bridging the gap between complex scientific concepts and accessible, high-quality visual communication.

Ethics and Broader Impact Statement

We acknowledge the significant ethical considerations associated with powerful generative technologies like AUTOFIGURE-EDIT. The primary risk involves the potential for misuse, where the system could be used to generate scientifically plausible but factually incorrect or misleading schematics to support false claims. To mitigate this risk, we are committed to transparency and responsible deployment. Our mitigation strategy is twofold. First, we explicitly declare that AUTOFIGURE-EDIT is an assistive tool with limitations. This disclaimer, stating that the system is not a substitute for expert verification and may not produce perfectly reliable outputs, will be prominently placed in this paper and in the README file of the public code repository. Second, the open-source license governing AUTOFIGURE-EDIT will include a mandatory attribution clause. This clause requires any academic publication using a figure generated by our tool to (a) include a specific section that discusses the role AI played in the work, and (b) explicitly caption the figure as having been generated by AUTOFIGURE-EDIT. These requirements are designed to ensure transparency and accountability in downstream use of our technology, fostering a research environment where AI tools augment, rather than compromise, scientific integrity.

References

Jonas Belouadi, Eddy Ilg, Margret Keuper, Hideki Tanaka, Masao Utiyama, Raj Dabre, Steffen Eger,

- and Simone Ponzetto. 2025. Tikzero: Zero-shot text-guided graphics program synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17793–17806.
- Jonas Belouadi, Anne Lauscher, and Steffen Eger. 2024a. [Automatikz: Text-guided synthesis of scientific vector graphics with tikz](#). In *The Twelfth International Conference on Learning Representations*.
- Jonas Belouadi, Simone Paolo Ponzetto, and Steffen Eger. 2024b. [DeTikZify: Synthesizing graphics programs for scientific figures and sketches with TikZ](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, volume 37, pages 85074–85108.
- BIT-DataLab. 2026. Edit-banana. <https://github.com/BIT-DataLab/Edit-Banana>. GitHub repository.
- Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, and 1 others. 2025. Sam 3: Segment anything with concepts. *arXiv preprint arXiv:2511.16719*.
- Yifan Chang, Yukang Feng, Jianwen Sun, Jiaxin Ai, Chuanhao Li, S. Kevin Zhou, and Kaipeng Zhang. 2025. [Sridbench: Benchmark of scientific research illustration drawing of image generation model](#). *Preprint*, arXiv:2505.22126.
- Kevin Ellis, Daniel Ritchie, Armando Solar-Lezama, and Josh Tenenbaum. 2018. Learning to infer graphics programs from hand-drawn images. *Advances in neural information processing systems*, 31.
- Jen-Yuan Huang, Tong Lin, and Yilun Du. 2026a. [Long-text-to-image generation via compositional prompt decomposition](#). In *The Fourteenth International Conference on Learning Representations*.
- Siyuan Huang, Yutong Gao, Juyang Bai, Yifan Zhou, Zi Yin, Xinxin Liu, Rama Chellappa, Chun Pong Lau, Sayan Nag, Cheng Peng, and Shraman Pramanick. 2026b. [Scifig: Towards automating scientific figure generation](#). *Preprint*, arXiv:2601.04390.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Luping Liu, Chao Du, Tianyu Pang, Zehan Wang, Chongxuan Li, and Dong Xu. 2025. [Improving long-text alignment for text-to-image diffusion models](#). In *The Thirteenth International Conference on Learning Representations*.
- Boris Malashenko, Ivan Jarsky, and Valeria Efimova. 2025. Leveraging large language models for scalable vector graphics processing: A review. *arXiv preprint arXiv:2503.04983*.
- Juan A Rodriguez, Abhay Puri, Shubham Agarwal, Issam H Laradji, Pau Rodriguez, Sai Rajeswar, David Vazquez, Christopher Pal, and Marco Pedersoli. 2025. Starvector: Generating scalable vector graphics code from images and text. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16175–16186.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, and 1 others. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494.
- Jianwen Sun, Fanrui Zhang, Yukang Feng, Chuanhao Li, Zizhen Li, Jiaxin Ai, Yifan Chang, Yu Dai, and Kaipeng Zhang. 2025. [From pixels to paths: A multi-agent framework for editable scientific illustration](#). *Preprint*, arXiv:2510.27452.
- Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan. 2024. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*.
- Zihang Wang, Yilun Zhao, Kaiyan Zhang, Chen Zhao, Manasi Patwardhan, and Arman Cohan. 2025. [SciSketch: An open-source framework for automated schematic diagram generation in scientific papers](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 403–417, Suzhou, China. Association for Computational Linguistics.
- Jingxuan Wei, Cheng Tan, Qi Chen, Gaowei Wu, Siyuan Li, Zhangyang Gao, Linzhuang Sun, Bihui Yu, and Ruifeng Guo. 2025. From words to structured visuals: A benchmark and framework for text-to-diagram generation and editing. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 13315–13325.
- Yiying Yang, Wei Cheng, Sijin Chen, Xianfang Zeng, Fukun Yin, Jiaxu Zhang, Liao Wang, Gang YU, Xingjun Ma, and Yu-Gang Jiang. 2025. [OmniSVG: A unified scalable vector graphics generation model](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Zhiyu Yang, Zihan Zhou, Shuo Wang, Xin Cong, Xu Han, Yukun Yan, Zhenghao Liu, Zhixing Tan, Pengyuan Liu, Dong Yu, and 1 others. 2024. Matplotlib: Method and evaluation for llm-based agentic scientific data visualization. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 11789–11804.
- Peng Zheng, Dehong Gao, Deng-Ping Fan, Li Liu, Jorma Laaksonen, Wanli Ouyang, and Nicu Sebe. 2024. Bilateral reference for high-resolution dichotomous image segmentation. *CAAI Artificial Intelligence Research*.

Dawei Zhu, Rui Meng, Yale Song, Xiyu Wei, Sujian Li, Tomas Pfister, and Jinsung Yoon. 2026a. [Paperbanana: Automating academic illustration for ai scientists](#). *Preprint*, arXiv:2601.23265.

Minjun Zhu, Zhen Lin, Yixuan Weng, Panzhong Lu, Qiujie Xie, Yifan Wei, Sifan Liu, QiYao Sun, and Yue Zhang. 2026b. [Autofigure: Generating and refining publication-ready scientific illustrations](#). In *The Fourteenth International Conference on Learning Representations*.

A Quantitative Evaluation

We compare against three categories of baselines: (1) **end-to-end text-to-image** methods (Sun et al., 2024), where we use GPT-Image (Hurst et al., 2024) to directly generate a scientific schematic from the paper text following standardized instructions; (2) **text-to-code** methods, where an LLM produces HTML and SVG code (Rodriguez et al., 2025; Malashenko et al., 2025; Yang et al., 2024) that is automatically rendered into images; and (3) **multi-agent** frameworks, represented by Diagram Agent (Wei et al., 2025) and AUTOFIGURE (Zhu et al., 2026b). In our experiments, we employ FigureBench as evaluation benchmark, which adopts a VLM-as-a-judge paradigm designed for structural reasoning and long-context scientific illustration assessment, and reports referenced, multi-dimensional scores. Evaluation follows a multi-dimensional rubric covering **Visual Design** (aesthetic quality, visual expressiveness, professional polish), **Communication Effectiveness** (clarity, logical flow), and **Content Fidelity** (accuracy, completeness, appropriateness). We additionally report **Win-Rate**, computed via blind pairwise comparisons against the reference illustrations, measuring how often a method is preferred as producing the more suitable figure for a given description.

B User Study

The user study consists of two complementary parts: (i) Figure Evaluation, which assesses the quality of the originally generated figure, and (ii) SVG Conversion Evaluation, which evaluates the fidelity of the converted SVG.

All scalar metrics are rated on a 5-point Likert scale (1 = lowest, 5 = highest). The process of evaluation is conducted through an integrated feedback interface on our website. After AUTOFIGURE-EDIT generates the final output, a rating dialog is automatically presented to the user. If the user dismisses the dialog, it will reappear when the user attempts to download the generated figure; submission of the evaluation is required prior to download. This design ensures that all collected ratings correspond to actual usage scenarios and are provided by users who have interacted with the generated result. To minimize ambiguity and promote consistent interpretation of the evaluation criteria, each rating dimension in the feedback dialog is accompanied by a small “?” icon. Clicking this icon opens a detailed scoring guideline that standardizes

the evaluation process. The guideline is structured into three components: a **Definition** that formalizes the metric, a **Guiding Question** that anchors user judgment, and a **Scoring Rubric** that specifies the semantic meaning of each score level. This design helps align user understanding with the intended evaluation protocol and reduces variance caused by subjective interpretation. The detailed scoring guidelines are provided as follows:

Part I: Figure Evaluation (PNG)

Scientific Semantic Correctness (1–5). (i) **Definition:** Measures whether the figure accurately represents the scientific concepts, processes, and relationships described in the input method text. (ii) **Question:** Does the figure correctly reflect the scientific content described in the method section? A score of 5 indicates fully correct semantic representation, while 1 indicates a misleading or incorrect depiction.

Information Completeness (1–5). (i) **Definition:** Measures whether all key components and steps described in the method text are present in the figure. (ii) **Question:** Are all essential elements from the method text included in the figure? Higher scores indicate more comprehensive coverage of essential elements.

Visual Presentation Quality (1–5). (i) **Definition:** Evaluates the visual clarity, readability, and overall professionalism of the figure. (ii) **Question:** Is the figure visually clear, well-aligned, and suitable for academic publication? Users should consider whether the figure would be suitable for academic publication.

Style Consistency (1–5). (i) **Definition:** Measures how well the generated figure matches the style of the provided reference image. (ii) **Question:** Does the generated figure follow the visual style of the reference figure?

Usability (Binary: 0/1). (i) **Definition:** Determines whether the figure is directly usable in an academic paper without major modifications. (ii) **Question:** Can this figure be directly used in a paper?

Part II: SVG Conversion Evaluation

Conversion Correctness (1–5). (i) **Definition:** Measures whether the structural and semantic elements are correctly preserved during the png-to-SVG conversion. This includes: (a) Correct placement of components; (b) Correct correspondence between original objects and SVG elements; (c)

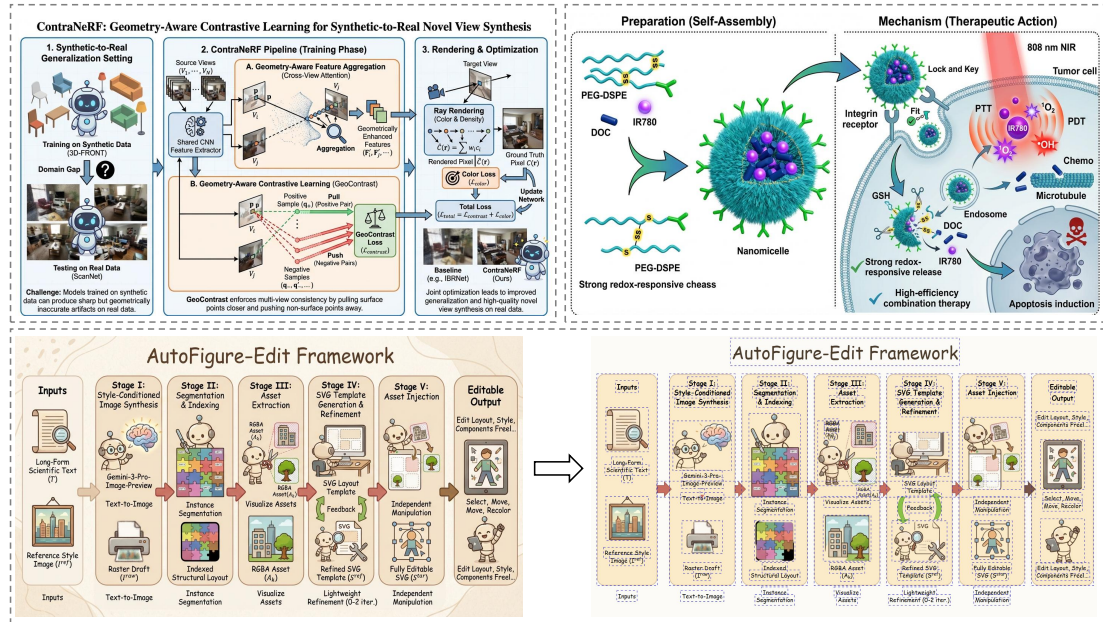


Figure 5: Qualitative results of AUTOFIGURE-EDIT.

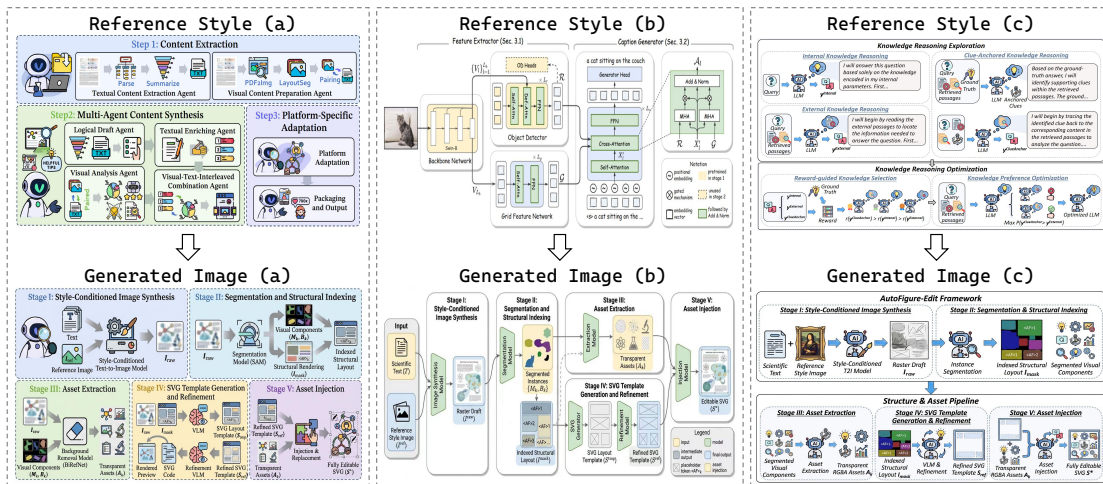


Figure 6: Comparison between Reference Images and Generated Images.

No missing or duplicated elements (ii) **Question:** Are all elements correctly preserved and positioned after conversion to SVG?

C License

AUTOFIGURE-EDIT is released under the MIT License. Users are solely responsible for any content generated using our demo (including the resulting figures and SVGs), as well as for ensuring compliance with applicable laws, third-party rights, and publication policies. We do not assume liability for any direct or indirect damages arising from the use of the software or from the generated outputs.