

This CVPR Workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

OpenCOLE: Towards Reproducible Automatic Graphic Design Generation

Naoto Inoue^{*} Kento Masui^{*} Wataru Shimoda^{*} Kota Yamaguchi CyberAgent, Japan



Figure 1. Generated graphic designs from OpenCOLE, given short user intentions.

Abstract

Automatic generation of graphic designs has recently received considerable attention. However, the state-of-the-art approaches are complex and rely on proprietary datasets, which creates reproducibility barriers. In this paper, we propose an open framework for automatic graphic design called OpenCOLE, where we build a modified version of the pioneering COLE and train our model exclusively on publicly available datasets. Based on GPT4V evaluations, our model shows promising performance comparable to the original COLE. We release the pipeline and training results to encourage open development.¹

1. Introduction

Graphic design serves as an essential medium for visual communication and is ubiquitously present in our daily lives. It encompasses a variety of multi-modal elements such as text or images, and the composition of these elements necessitates an in-depth comprehension of many aesthetic facets, including layout and readability. Automatic graphic design has long been a goal within the research community, with a rich history of exploration and experimentation [6, 20, 28]. Over the past decade, learning-based formulations, particularly those employing neural networks, have been utilized in various approaches aimed at resolving specific design sub-tasks, such as layout generation [8, 9, 13–15, 17, 30], font recommendation [12, 25, 31], and colorization [16, 22–24, 29]. With the substantial advancements in modeling texts and images, some ambitious papers have begun to tackle the generation or editing of complete graphic designs [10, 18, 27].

Most recently, COLE [11] has exhibited outstanding proficiency in generating graphics based on brief intention prompts. COLE deconstructs the intricate generation procedure into several distinct stages. The outputs comprise a background layer, an object image layer, and text layers with a broad spectrum of typographic properties such as font type and line spacing. An off-the-shelf graphic renderer rasterizes these layers to obtain the preview image. Owing to this methodological breakdown, texts are presumed to maintain legibility and editability, which is a crucial requirement in graphic design.

The primary shortcoming of COLE is its lack of publicly accessible datasets and codes. Because of its cascaded process and reliance on a graphical renderer, its proprietary nature poses significant challenges to reproduction.

^{*:} equal contribution

https://github.com/CyberAgentAILab/OpenCOLE

We propose OpenCOLE, an open-source implementation of COLE that utilizes only publicly available datasets and models. This aims to democratize the development of generative models for graphic design and foster a more inclusive community. Following COLE, we present a framework composed of three stages with minor modifications. For the dataset, we utilize Crello [27], which is a publicly available graphic design dataset, and synthesize all the necessary data for training each stage of OpenCOLE.

We experimentally show that our OpenCOLE matches COLE in terms of GPT4V-based evaluation on five aspects specific to graphic design on DESIGNERINTEN-TION benchmark [11]. We also compare OpenCOLE with LLM-augmented text-to-image models and show it performs better than DeepFloydIF [3] but it still lags behind SDXL1.0 [21] and DALL-E3 [2]. We further discuss potential challenges to promote the research for automatic graphic design generation.

Our contributions can be summarized as follows:

- We propose to construct an open-sourced automatic graphic design generation pipeline by leveraging open-sourced data and models to facilitate the democratization of development.
- Our empirical results suggest that our OpenCOLE performs almost on par with COLE.

2. Method

Following COLE [11], the aim of our OpenCOLE architecture is to convert a user's intention to create a graphic design into a complete design composed of visual elements. Our OpenCOLE implementation consists of a design plan generation module, an image generation module, and a typography generation module, as shown in Fig. 2.

First, a design plan generation module takes a user's intention and translates it into a detailed explanation and information for design composition in JSON format, which we refer to as a design plan. Next, an image generation module generates an image that covers the entire canvas following the design plan. Then, a typography generation module synthesizes typographic attributes, such as font or color, from the generated image and the design plan. Finally, a graphic renderer composes the image and the typography attributes to place the texts to form a final design image.

This design choice mimics the architecture of COLE, where the original architecture consists of several corresponding components, namely a Design-LLM, Image Generation Modules, and a Typography-LMM. We introduce slight adjustments for each component to complement the ambiguous part in the original COLE paper for our Open-COLE implementation.

2.1. Design Plan Generation

In the design plan generation module, we use an LLM to convert a rather vague user intention into a design plan that works as a schematic of the design composition. This design plan instructs the other modules to generate design elements for the final composition.

The design plan consists of 4 items: description, keywords, captions, and headings. Detailed explanations are as follows:

- Description: A detailed caption of the final design composition. The caption explains what the image is about, the depicted object, and the context.
- Keywords: A list of keywords to express the image in various aspects, such as color, object, and genre.
- Captions: A background caption and an object caption. The background caption mainly explains the color and objects behind the main subject. The objects caption explains individual objects that should appear in the design with their coordination.
- Headings: A heading and a subheading for the final composed design. Since a design is often made up of multiple texts, we explicitly specify the text here for the typography generation module.

We have used a pre-trained LLM with in-context learning prompt [7] to generate a design plan. The in-context learning prompt contains examples of 5 pairs for a user intention and a design plan to facilitate the generation of a design plan. This approach is a simplification of the original COLE that fine-tunes Llama [26] for JSON generation.

2.2. Image Generation

In the image generation module, we transform a design plan into a graphic design image without text layers. Since an off-the-shelf text-to-image model tends to yield illegible texts when instructed to generate graphic design images, we fine-tune the generator on a dataset of text-image pairs. We concatenate the background and object caption in the design plan to form the input text. We render the original graphic document data, excluding text layers, to create the target image.

This single-stage generation process deviates from the original COLE's two-stage generation. COLE first employs a text-to-background model to generate an embellished background image. Then, COLE uses a text-to-object model to take the background image as an additional conditioning input. We adopted the single-stage approach due to the challenges of obtaining reliable input-target pairs for both stages.

2.3. Typography Generation

In the typography generation module, we fine-tune large multi-modal models (LMMs) to generate typographic attributes given the outputs from the former modules: the de-



Figure 2. An architecture of OpenCOLE. We mimic the architecture of COLE with adjustments. The user intention is first converted to a design plan with GPT3.5 and in-context learning. Then, the image generation module and the typography generation module synthesize design elements following a design plan. Finally, the graphic renderer composes the final image.

sign plan and the generated image. This module follows the original TypographyLMM of COLE. The typographic attributes consist of fine-grained text styling parameters, such as font, font size, color, and letter spacing. The typography LMM handles fine-grained attributes as language via JSON format, which contains all typographic attributes in a design for per-text elements with hierarchical structures.

3. Experiments

3.1. Implementation Details

For the base dataset, we use the Crello dataset [27]. It provides around 22k design templates for many domains, such as social media posts, banner ads, blog headers, and printed posters. Each sample is stored in a vector format and contains images, texts, layouts, and typography information.

Dataset for In-Context Learning For the design plan generation with an LLM and in-context learning, we prepared an intention-to-design-plan dataset following the original COLE with Crello dataset.

For intention creation, we used GPT3.5 with adjusted prompt from COLE. This adjustment includes adding intention examples that were referenced but not in the original COLE's prompt, and providing information such as title, format, keywords, and texts.

For design plan creation, we followed COLE to use LLaVA-1.5-13B model [19], but with a slightly simplified JSON structure. Since the prompts used for creating a design plan with LLaVA in the original COLE are not available, we implemented our design plan generation using the divide-and-conquer approach. For each item of a design plan (description, keywords, captions, and headings), we designed a specific prompt to extract only the corresponding information from an image using LLaVA. This is to reduce the complexity of the task and to improve the stability of the LLM in successfully extracting meaningful information as JSON text. After each piece of information for an image is extracted, we merge them to form a single design plan.

Design Plan Generation We used GPT3.5 as a base LLM for design plan generation with in-context learning. We randomly selected 5 examples from the intention-to-design-plan dataset synthesized from the Crello dataset for in-context examples.

Image Generation We fine-tune SDXL1.0 [21] for 10,000 iterations with a constant learning rate of 8×10^{-7} with 1,000 warmup steps using Simpletuner [5]. Images with less than 0.5M pixels in the area are filtered to avoid upscaling artifacts. SDXL1.0 has a maximum token length limit of 77. During training, we randomly omit some sentences in the input to meet the limit. During sampling, we employ Compel [1] to feed all the information in the text input and generate images with size of 1024×1024 .

Typography Generation Unlike COLE, we merge headings and subheadings into a single list because the Crello dataset does not have such attributes. Similarly to COLE, we employ LLaVA1.5-7B for the base LMM. We fine-tune the model for 6 epochs with a batch size of 32, learning rate of 2e-4, max token length of 4096, LoRA rank 128, LoRA alpha 256, and input resolution of 336×336 .

3.2. Benchmark, Metrics, and Baselines

Following COLE [11], we test all the models on DE-SIGNERINTENTION benchmark, which contains 200 graphic design intention prompts. These prompts span six salient categories within graphic design: advertising, events, marketing, posts, covers & headers, and creative. The quality of the generated design images is evaluated on a scale from 1 (poor) to 10 (excellent) using GPT4V [4] to assess five crucial aspects.

We obtain COLE's results from their project page. Following COLE, we manually test text-to-image models as additional baselines. These baselines take GPT4augmented prompts as inputs and use DeepFloydIF [3], SDXL1.0 [21], or DALL-E3 [2] to generate the final image.

Table 1. Comparison of different models on full DESIGNER-INTENTION. GPT4V evaluation considers the following aspects: (i) *design and layout*, (ii) *content relevance*, (iii) *typography and color*, (vi) *graphics and images*, and (v) *innovation*. [†] indicates that GPT-4 converts user intentions into detailed prompts fed to text-to-image models. Higher is better.

	(i)	(ii)	(iii)	(iv)	(v)	Avg.
Text fixed						
DeepFloyd/IF [3] [†]	5.6	6.3	4.8	6.5	5.0	5.6
SDXL [21] [†]	7.2	7.6	7.1	7.9	6.3	7.2
DALL-E3 [2] [†]	7.8	8.3	7.6	8.7	7.2	7.9
Text editable						
COLE [11]	6.0	6.9	5.7	6.2	5.1	6.0
OpenCOLE	6.3	7.0	5.6	7.1	5.3	6.3



Figure 3. Failure cases of OpenCOLE.

3.3. Quantitative Evaluation

In Tab. 1, we see that OpenCOLE performs competitively with the original COLE. For reference, we also show the results of simple text-to-image baselines DESIGNERINTEN-TION. There is still room for improvement for OpenCOLE since it largely lags behind DALL-E3 and SDXL1.0 in all the aspects, though the texts in the baselines are not editable and sometimes illegible.

3.4. Qualitative Evaluation

Fig. 1 demonstrates the generated designs by OpenCOLE. We observe that OpenCOLE can generate plausible graphic designs. The generated texts do not always accurately reflect intentions, but most of them are at least partially related. Notably, the generated images tend to contain empty spaces for text placement, and the typography LMM appears to position texts in these spaces by recognizing them.

Fig. 3 shows failure cases of OpenCOLE. We observe that OpenCOLE tends to generate some low-legibility texts. Some factors, we believe, include generating excessively long sentences without line breaks, generating images with small empty spaces for text placement, overlooking color contrasts, and choosing very thin fonts.

Fig. 4 presents a comparison between COLE and Open-COLE. Both methods can generate plausible designs, but they exhibit different tendencies in designs from them. While COLE generates simple designs with distinct texts, OpenCOLE generates more informative images with complex texts. We attribute these differences to variations in the datasets used and modifications made to certain modules.

4. Discussion

While we believe that OpenCOLE constitutes a pivotal initial stride towards the open-source progression of auto"Design an invitation for a Halloween party with a spooky and playful design. The party date is October 31st and guests are encouraged to drop in for the celebration."

"Create a logo for a honey brand called Royal Honey with a playful and fun design that showcases the sweetness and naturalness of the product."

"Design a poster to advertise a virtual reality event. The poster should emphasize the ability to create one's own experience. The website for the event is www.yoursite.com."

"Create an Instagram story to advertise an antique vase available at a store with a brown color scheme. The promotion offers up to a 50% discount and is only available from October 22-25, 2023."

Intention

OpenCOLE

Figure 4. The comparisons between COLE and OpenCOLE. The middle and right images are generated designs from the left intentions by COLE and OpenCOLE, respectively.

COLE

mated graphic design generation, many discussions persist. The most notable point is the dependency on black box GPT4V assessment. We observe that some generated images without textual content are highly ranked by GPT4V. This often happens in GPT4-augmented text-to-image baselines. Evaluation considering both the input intention and output-generated images would be necessary.

References

- [1] Compel. https://github.com/damian0815/ compel. Accessed: 2024-03-15. 3
- [2] DALL-E3. https://openai.com/dall-e-3. Accessed: 2024-03-15. 2, 3, 4
- [3] Deepfloyd IF. https://github.com/deep-floyd/ IF. Accessed: 2024-03-15. 2, 3, 4
- [4] GPT-4(V). https://cdn.openai.com/ contributions/gpt-4v.pdf. Accessed: 2024-03-15.3
- [5] Simpletuner. https://github.com/bghira/ SimpleTuner. Accessed: 2024-03-15. 3
- [6] Maneesh Agrawala, Wilmot Li, and Floraine Berthouzoz. Design principles for visual communication. *Communications of the ACM*, 54(4), 2011.
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. 2020. 2
- [8] Kamal Gupta, Alessandro Achille, Justin Lazarow, Larry Davis, Vijay Mahadevan, and Abhinav Shrivastava. Layout-Transformer: Layout generation and completion with selfattention. In *ICCV*, 2021. 1
- [9] Naoto Inoue, Kotaro Kikuchi, Edgar Simo-Serra, Mayu Otani, and Kota Yamaguchi. LayoutDM: Discrete Diffusion Model for Controllable Layout Generation. In *CVPR*, 2023.
- [10] Naoto Inoue, Kotaro Kikuchi, Edgar Simo-Serra, Mayu Otani, and Kota Yamaguchi. Towards flexible multi-modal document models. In *CVPR*, 2023. 1
- [11] Peidong Jia, Chenxuan Li, Zeyu Liu, Yichao Shen, Xingru Chen, Yuhui Yuan, Yinglin Zheng, Dong Chen, Ji Li, Xiaodong Xie, et al. Cole: A hierarchical generation framework for graphic design. arXiv preprint arXiv:2311.16974, 2023. 1, 2, 3, 4
- [12] Shuhui Jiang, Zhaowen Wang, Aaron Hertzmann, Hailin Jin, and Yun Fu. Visual font pairing. In *IEEE Transaction on Multimedia*, 2019. 1
- [13] Z. Jiang, J. Guo, S. Sun, H. Deng, Z. Wu, V. Mijovic, Z. Yang, J. Lou, and D. Zhang. LayoutFormer++: Conditional Graphic Layout Generation via Constraint Serialization and Decoding Space Restriction. In *CVPR*, 2023. 1
- [14] Akash Abdu Jyothi, Thibaut Durand, Jiawei He, Leonid Sigal, and Greg Mori. LayoutVAE: Stochastic scene layout generation from a label set. In *CVPR*, 2019.
- [15] Kotaro Kikuchi, Edgar Simo-Serra, Mayu Otani, and Kota Yamaguchi. Constrained graphic layout generation via latent optimization. In ACM MM, 2021. 1
- [16] Kotaro Kikuchi, Naoto Inoue, Mayu Otani, Edgar Simo-Serra, and Kota Yamaguchi. Generative colorization of structured mobile web pages. 2023. 1
- [17] Jianan Li, Jimei Yang, Aaron Hertzmann, Jianming Zhang, and Tingfa Xu. LayoutGAN: Generating graphic layouts with wireframe discriminators. In *ICLR*, 2019. 1
- [18] Jinpeng Lin, Min Zhou, Ye Ma, Yifan Gao, Chenxi Fei, Yangjian Chen, Zhang Yu, and Tiezheng Ge. Autoposter:

A highly automatic and content-aware design system for advertising poster generation. In *ACM MM*, 2023. 1

- [19] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 3
- [20] Simon Lok and Steven Feiner. A survey of automated layout techniques for information presentations. In *SmartGraphics*, 2001. 1
- [21] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952, 2023. 2, 3, 4
- [22] Qianru Qiu, Xueting Wang, and Mayu Otani. Multimodal color recommendation in vector graphic documents. In ACM MM, 2023. 1
- [23] Qianru Qiu, Xueting Wang, Mayu Otani, and Yuki Iwazaki. Color recommendation for vector graphic documents based on multi-palette representation. 2023.
- [24] Xinyu Shi, Ziqi Zhou, Jing Wen Zhang, Ali Neshati, Anjul Kumar Tyagi, Ryan Rossi, Shunan Guo, Fan Du, and Jian Zhao. De-stijl: Facilitating graphics design with interactive 2d color palette recommendation. In *CHI*, 2023. 1
- [25] Wataru Shimoda, Daichi Haraguchi, Seiichi Uchida, and Kota Yamaguchi. Towards diverse and consistent typography generation. 2024. 1
- [26] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288, 2023. 2
- [27] Kota Yamaguchi. Canvasvae: Learning to generate vector graphic documents. In *ICCV*, 2021. 1, 2, 3
- [28] Xuyong Yang, Tao Mei, Ying-Qing Xu, Yong Rui, and Shipeng Li. Automatic generation of visual-textual presentation layout. *TOMM*, 12(2), 2016. 1
- [29] Lin-Ping Yuan, Ziqi Zhou, Jian Zhao, Yiqiu Guo, Fan Du, and Huamin Qu. Infocolorizer: Interactive recommendation of color palettes for infographics. *IEEE TVCG*, 28(12), 2021.
 1
- [30] Junyi Zhang, Jiaqi Guo, Shizhao Sun, Jian-Guang Lou, and Dongmei Zhang. Layoutdiffusion: Improving graphic layout generation by discrete diffusion probabilistic models. In *ICCV*, 2023. 1
- [31] Nanxuan Zhao, Ying Cao, and Rynson WH Lau. Modeling fonts in context: Font prediction on web designs. *Comput. Graph. Forum*, 37(7), 2018. 1