

Design-o-meter: Towards Evaluating and Refining Graphic Designs

Sahil Goyal^{*†}, Abhinav Mahajan^{*‡}, Swasti Mishra, Prateksha Udhayanan, Tripti Shukla,
K J Joseph, Balaji Vasan Srinivasan

[†]IIT Roorkee [‡]IIIT Bangalore Adobe Research

sahilg@ma.iitr.ac.in, abhinav.mahajan@iiitb.ac.in, {josephkj, balsrini}@adobe.com

Abstract

Graphic designs are an effective medium for visual communication. They range from greeting cards to corporate flyers and beyond. Off-late, machine learning techniques are able to generate such designs, which accelerates the rate of content production. An automated way of evaluating their quality becomes critical. Towards this end, we introduce Design-o-meter, a data-driven methodology to quantify the goodness of graphic designs. Further, our approach can suggest modifications to these designs to improve its visual appeal. To the best of our knowledge, Design-o-meter is the first approach that scores and refines designs in a unified framework despite the inherent subjectivity and ambiguity of the setting. Our exhaustive quantitative and qualitative analysis of our approach against baselines adapted for the task (including recent Multimodal LLM-based approaches) brings out the efficacy of our methodology. We hope our work will usher more interest in this important and pragmatic problem setting. Project Page: sahilg06.github.io/Design-o-meter.

1. Introduction

Graphic designs are becoming increasingly ubiquitous: advertisement content, menu cards at restaurants, campaign flyers, and so on. It is a composite of text, images, and shapes that harmoniously intermingle in an aesthetically pleasing way to convey the intended message effectively. A typical workflow of a graphic designer involves ideation, creation, and refinement stages. Each of these stages has its unique characteristics: *ideation* involves planning the design, *creation* involves aggregating the design elements and creating the first version, and *refinement* involves improving the design iteratively. The refinement stage is particularly prone to redundancy, as it involves fine-tuning details, adjusting layouts, and sometimes reworking significant portions of the design to meet the desired standards. Designers often undergo numerous feedback and revision cycles, which can be labor-intensive and time-consuming. Genera-

^{*}Equal contribution. Work done during internship at Adobe Research.

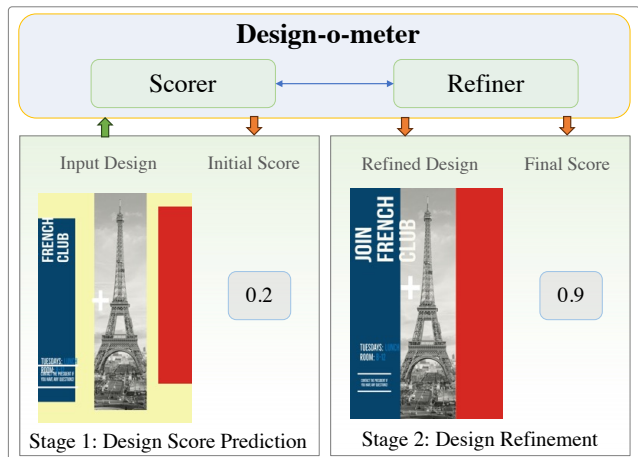


Figure 1. The figure presents an overview of Design-o-meter. It takes a design as input. The *scorer* evaluates the design and provides a unified design score. The *refiner* refines the design with the help of the design score to improve its aesthetic appeal.

tive AI technologies can work hand-in-hand with designers to supplement them in all phases of their creative workflow.

Off-late, such generative models [7, 8, 20, 22, 57] has been used for creating designs and layouts from user intents. Some of these methods can take in design assets from a user and generate designs by composing them [20, 21, 51], while some others can generate the entire graphic design from a text prompt [23, 27, 60]. These can generate large number of designs with very low latency. Coupled with the increasing online and offline consumption of designs, there is a strong demand for automatic tools to evaluate them.

The metrics currently used to evaluate graphic designs, like Fréchet Inception Distance (FID) [19], mean IOU and max IOU are often insufficient to capture the nuanced aspects of design quality. Further, these metrics strongly depend on having a high-quality reference design for comparison. Such evaluations might even penalize the creative freedom of the model (a new design might have a completely different location for the constituent elements but still might look good). Reference-free evaluation metrics like alignment and overlap [34] have also been proposed to

evaluate the generated layouts. These metrics fail to provide a holistic assessment of overall design quality. These limitations highlight the need for more comprehensive and robust evaluation methods that can assess the effectiveness, usability, and visual appeal of generated designs such that it corroborates with human perception and its use cases.

Towards this end, we propose Design-o-meter. As shown in Fig. 1, it contains a *scorer*, which quantitatively evaluates how good a design is, and a *refiner*, which takes in the score from the scorer and refines the design to improve the score. A design is a composite of components like text, images, shapes, and icons. Each component has its properties like color, location, size, content, opacity, shadow, and so on. A visually appealing design has an optimal value for each property and component. Hence, the search space for good designs is indeed combinatorially large. Further, the design trends also change with time. To effectively model such a complex design space, we propose a data-driven approach for the scorer and a novel meta-heuristic methodology for the refiner. Concretely, the scorer is modeled as a metric learning model that scores good-designs over bad-designs (Sec. 3.1), and the refiner is a genetic algorithm with a novel design-specific crossover operation called SWAN: Design Specific Crossover with Smart Snapping (Sec. 3.2).

We conduct thorough experimental analysis to test the mettle of both our scorer and refiner modules. Our experiments reveal that our scorer is able to capture small nuances in design documents that recent Multimodal LLMs fail to capture, while our newly proposed SWAN is able to efficiently navigate the complex design space to refine existing designs, comfortably outperforming recent state-of-the-art approaches. Further, we clearly ablate the different components of the method, and provide sensitivity analysis of the various design choices of our framework.

Our key contributions are summarized below:

- We develop a novel holistic framework, Design-o-meter, that can provide a comprehensive design score and further refine the designs to improve the score.
- We propose a metric learning model that learns to disambiguate between good designs and bad designs.
- Our novel smart snapping crossover methodology SWAN, refines designs to improve its aesthetic appeal.
- Our exhaustive quantitative and qualitative evaluation brings out the efficacy of our proposed Design-o-meter.

2. Related Work

In the following subsections, we first discuss related methodologies that evaluate designs and then talk about approaches that refine them. Finally, we provide a brief summary of generic algorithms, which is the basis of our proposed refiner module.

2.1. Design Evaluation

Design evaluations strategies can be grouped into heuristic-based approaches and data-driven approaches:

Design Heuristics based Approaches: Design heuristics are cognitive tools that designers and engineers use to measure different design aesthetics. Prior works [4, 18, 36, 37, 43, 50, 62] use heuristics as a fixed set of formulas, each formula measuring a particular aspect of the design. Such methods do not consider the overlap between design elements, which is common in graphic designs. Despite the valuable insights provided by the heuristic rules, they don't offer a comprehensive view of the overall design. Also, heuristics are often subjective, and their concrete meaning can deviate, which might lead to inconsistencies.

Data-driven Approaches: Most of the data-driven approaches [14, 49, 56] formulate the problem as a regression or classification task, i.e., mapping images to aesthetic ratings given by human annotators. Dou *et al.* [14] train a convolution network to predict the aesthetic ratings of webpages. They use a dataset [48] of webpage screenshots with human-annotated aesthetic ratings. Wan *et al.* [56] obtain global, local, and aesthetic features from the webpage layout to predict the aesthetics. The above-mentioned techniques rely on human-annotated aesthetic ratings, which have several drawbacks. Creating a high-quality human-annotated dataset is time-consuming and expensive. It may also introduce subjectivity and bias.

An interesting approach that requires minimal or no human interference is the use of Siamese Networks [5]. Ever since they emerged, Siamese Networks have found application in various diverse fields such as self-supervised representation learning [1, 3, 6, 61], audio-visual synchronization [10, 16, 45], and measuring aesthetics [29, 30, 53, 63], etc. Their popularity is due to several reasons, such as the sharing of parameters among its twin networks, the ability to navigate the search space, and the extraction of distinctive features for downstream tasks. They allow learning in unsupervised settings. They compare the data instances in pairs instead of directly using labels and are robust to data imbalance. Zhao *et al.* [63] propose a Siamese-based deep learning framework to estimate personality scores of a graphic design. Aesthetics++ [30] train a Siamese-based network to estimate the aesthetics of a design. Kong *et al.* [29] and Aesthetics++ [30] use comparatively less human annotation to create training pairs but still introduce bias and may not be accurate. Tabata *et al.* [53] arbitrarily moves the layout elements without any guidance to generate negative examples for training. However, designs generated randomly are not necessarily bad and can sometimes exhibit creative and unique layouts. We follow the successful practice of existing works [30, 63] utilizing Siamese networks. Instead of human annotation and random per-

turbations without guidance, we use intelligent transformations to create design pairs for training. Moreover, unlike the methods mentioned above, our approach incorporates layout-information-rich color encodings alongside the design renderings, providing a more comprehensive representation of the design elements and additional guidance to the scoring model.

Recently, there are some efforts [9, 26] that tries to evaluate a graphic design using large multi-modal LLMs like GPT-4o [40] and LLaVA [35]. As these models are trained on huge amounts of data which includes design data too, they are able to analyse graphic designs well. We compare with LLaVA-NeXT and GPT-4o in Sec. 4.1.

2.2. Refining Layouts

Prior works like [38, 43] minimize energy functions for typical design principles like white space, symmetry, and alignment using simulated annealing. This approach is not scalable and requires significant computational time (up to 40 minutes to generate one optimized layout), making it impractical for real-time or large-scale applications. Pang *et al.* [44] generate a set of candidate designs by randomly perturbing the existing design and then selecting the best out of the perturbed designs using heuristic rules. However, such an approach is highly inefficient as the design space is vast and complex. Also, relying on a fixed set of heuristic formulas is not a good approach, as discussed in the section Sec. 2.1. Aesthetics++ [30] generates the candidate designs from the input design by traversing a segmentation tree (created using hierarchical segmentation of the graphic design) and leveraging design principles. The candidate designs with the highest aesthetic score is taken as the refinement suggestion. They use a human-annotated dataset for a data-driven approach. However, the dataset used may not represent all user preferences, leading to biased or skewed results. Also, generating multiple candidate designs through tree traversal is computationally intensive and time-consuming.

RUIE [47] models the refinement as a denoising task and trains a transformer [55] model to denoise the layouts. RUIE only focuses on aligning UI elements; it neglects content and other critical aspects of design, like color schemes and interactive elements. Moreover, reliance on transformer architectures introduces significant latency in the training and inference phases. FlexDM [22] employs multitask learning in a single transformer-based model to solve various design tasks by predicting masked fields in incomplete vector-graphic documents. While this approach allows for flexible design refinement, it requires significant computational resources, making implementation and scaling challenging.

In contrast to the above approaches, our method is computationally efficient, providing a more practical solution

for design refinement tasks. Our genetic algorithm-based refinement module takes approximately 30 sec to generate a refined design. We also avoid human intervention in our approach to get unbiased results.

2.3. Genetic Algorithms

Metaheuristic algorithms are applied to address complex real-world problems across domains like engineering, economics, management, etc. A genetic algorithm [28] (GA) is a metaheuristic algorithm inspired by natural selection. It is a population-based search algorithm used for optimization. Classical GA has an objective function (fitness function), chromosome representation of the population, and operators inspired by biology (selection, crossover, and mutation). Population is improved iteratively using the genetic operators and selecting the fittest.

Multiobjective Genetic Algorithms (MOGAs) differ primarily from standard GAs in how they assign fitness functions, while the remaining steps follow the same procedure as in GAs. They focus mainly on convergence and diversity. NSGA [52] (Non-dominated sorting genetic algorithm) is a multiobjective genetic algorithm that finds multiple pareto-optimal solutions in a single run. It lacks elitism, needs to specify the sharing parameter (σ_{share}), and has high computation complexity. To solve these, Deb *et al.* [12] introduce a fast, elitist, non-dominated sorting genetic algorithm, NSGA-II. It has been applied in various real-world applications, demonstrating its versatility and effectiveness.

We remove the dependency on multiple objectives as we represent our model score as our unified objective function. We introduce a novel crossover method for designs SWAN: Design Specific Crossover with Smart Snapping, and refine the mutation functions for our task. With a unified objective function, we bypass the need to filter the pareto-front, which is essential for real user-based scenarios.

3. Methodology

Our proposed approach Design-o-meter, takes an input a design, and proposes a quantitative score from a *scorer* module, and in-turn uses this score to refine the design using a *refiner* module. A design is a composite of a set of elements, which are put together in an aesthetic way. Hence, a design D , can be represented as its metadata D_{meta} . For instance, a design with a background image and foreground text saying ‘Hi!’, can be represented as $D_{meta} = [\{x_1, y_1, w_1, h_1, image_location\}, \{x_2, y_2, w_2, h_2, 'Hi!'\}]$, where x_i, y_i, w_i, h_i refers to the location and dimension information, in its most simplistic form. The image and text can have more attributes like opacity, overlay, shadow, font details, emphasis and so on. Using D_{meta} , we can render its rendition image $\mathcal{I}(D_{meta}) \in \mathbb{R}^{H \times W \times 3}$ and a color encoded layout $\mathcal{L}(D_{meta}) \in \mathbb{R}^{H \times W \times 3}$, as illustrated in Fig. 3. Our scorer $\mathcal{S}(\mathcal{I}(D_{meta}), \mathcal{L}(D_{meta}))$

learns to measure the goodness and global aesthetics of the input design. The refinement module $R(\mathbf{d}_{meta})$, takes as input $\mathbf{d}_{meta} \subset \mathbf{D}_{meta}$, which is a set of actionable layout attributes that we are interested in optimizing. It refines these attributes to generate \mathbf{d}_{meta}^* , by maximizing the corresponding design score from the scorer: $\mathbf{d}_{meta}^* \mid \mathcal{S}(I(\mathbf{d}_{meta}^*), L(\mathbf{d}_{meta}^*)) > \mathcal{S}(I(\mathbf{d}_{meta}), L(\mathbf{d}_{meta})) \forall \mathbf{d}_{meta}$. \mathbf{d}_{meta}^* can be rendered to obtain the refined design. We detail about our scorer and refiner in the next few sub-sections.

3.1. Design Scorer

We propose to use a data-driven, self-supervised approach towards learning the scorer module. Different from other works which tries to regress a scalar score from the input design [14, 49, 56], we learn a Siamese model in a contrastive setup, to differentiate between good designs and bad designs.

Model Architecture: The scorer function \mathcal{S} is a composition of a feature extractor network \mathcal{F} , and a final scoring block S_{block} . We use a four layer convolutional network as the feature extractor, and a three layer fully connected network for the scoring block. Given the input design \mathbf{D}_{meta} , its design score is computed as:

$$\mathcal{S}(\mathcal{I}(\mathbf{D}_{meta}), \mathcal{L}(\mathbf{D}_{meta})) = S_{block}(\mathcal{F}(\mathcal{I}(\mathbf{D}_{meta}), \mathcal{L}(\mathbf{D}_{meta}))) \quad (1)$$

We train our entire model from scratch. Experimentally, we find that using pretrained feature extractors like DINO-V2 [42], CLIP [46], BLIP [33] and ViT [55] gives inferior results. This is because these models are trained on natural images, which are distinctly different from the design data, and thus struggle in gauging and extracting important features specific to quantify design aesthetics.

Our light-weight convolutional network with group-normalization accepts the rendition image $I(\mathbf{D}_{meta})$ and the colour-coded layout map $\mathcal{L}(\mathbf{D}_{meta})$ concatenated across the channel dimension as follows:

$$\{N, C, H, W\} + \{N, C, H, W\} \equiv \{N, 2 * C, H, W\}$$

where N , C , H , and W are batch size, number of channels, height, and width, respectively. The layout map $\mathcal{L}(\mathbf{D}_{meta})$ effectively warps the multi-layer information of a design into the scorer. This information helps the model to give attention to the constituent components of the design, and its relative

Table 1. Color codes for layout encoding.

Element type	Color (R, G, B)
Image	(0, 100, 0)
Text	(0, 0, 100)
Text and Text overlap	(0, 0, 0)
Text and Image Overlap	(100, 0, 0)
Image and Text Overlap	(100, 100, 0)
Image and Image Overlap	(0, 100, 100)

positions while proposing a design score. Tab. 1 summarize

the color coding and Fig. 3 shows layout map and its corresponding designs rendition.

Training Details: Similar to the traditional Siamese training setup, the weights of \mathcal{S} is shared. We curate good designs \mathbf{D}_{meta}^{good} and bad designs \mathbf{D}_{meta}^{bad} (explained in the next sub-section) to train the model using the following loss function:

$$L_{scorer} = \alpha L_{rank} + \beta L_{sim}; \quad (2)$$

where L_{rank} is hinge loss:

$$L_{rank} = \max(0, m - (\mathcal{S}(\mathbf{D}_{meta}^{good}) - \mathcal{S}(\mathbf{D}_{meta}^{bad}))) \quad (3)$$

and L_{sim} is a similarity loss, formulated as below:

$$L_{sim} = \ln(e^{2 * P_{sim}(\mathcal{S}(\mathbf{D}_{meta}^{good}), \mathcal{S}(\mathbf{D}_{meta}^{bad}))} + 1); \quad (4)$$

P_{sim} is an embedding similarity computed as the dot product between the tanh-activations of the ‘‘good’’ and ‘‘bad’’ design pairs as follows:

$$P_{sim} = \frac{\mathcal{F}(\mathbf{D}_{meta}^{good}) \cdot \mathcal{F}(\mathbf{D}_{meta}^{bad})}{\max(\|\mathcal{F}(\mathbf{D}_{meta}^{good})\|_2, \|\mathcal{F}(\mathbf{D}_{meta}^{bad})\|_2, \epsilon)}; \epsilon > 0 \quad (5)$$

Dataset Creation: Given a set of designs from any design dataset [59], we first filter them on two criteria: 1) total number of elements should be at most 10, 2) text should not overlap with images; to create a list of good designs \mathbf{D}_{meta}^{good} . Next, we surgically modify these designs to make them bad by altering the location and scale of its constituent elements. We employ 22 such types of transformations to create \mathbf{D}_{meta}^{bad} . The location based perturbations are:

Noise addition: In order to simulate imperfect designs, noise from standard normal distribution with mean 0 is added to the center coordinates of the design elements. By varying the standard deviation across 0.05, 0.1, 0.2, and .5, we control the degree of perturbation, and thus the degree of badness in the graphic design.

Moving specific type and group of elements differently: Not all elements of a bad-design will have equal amount of layout shift. In-order to accommodate this aspect, we arbitrarily move specific elements in the graphic design such as the largest element, smallest element, two largest elements, and two smallest elements.

Clutter: Another aspect of designs that makes them naturally bad is clutter, and hence, we clutter all the elements of a graphic design at different positions such as center, top-left, top-right, bottom-left, and bottom-right.

Scale-based transformations are similar to the position-based transformations. The height and width of the design elements are modified instead of the coordinates of the center. Further, the position-based and scale-based transformations are combined to cover more cases. A critical aspect of

these augmentations is that they are agnostic to the dataset and are modeled on how humans consider a design good and bad. We can improve these augmentations further to incorporate any newer cognitive constraints, which we leave for future explorations. Though the model is trained only on location and scale perturbations, we see in our experiments that the model is able to learn design principles beyond just layout principles. We attribute this to the contrastive learning objective in which the model is trained.

3.2. Design Refiner

As a design is a composite of multiple components, search space that constitutes all the designs is vast. Traversing this space of designs in a meaningful way would enable us to find better aesthetically pleasing designs, which improves over the initial design. Here, we propose an efficient approach based on a meta-heuristic algorithm, equipped with our novel SWAN: Design Specific Crossover with Smart Snapping. Our scorer (Sec. 3.1), guides SWAN, to refine its design aesthetics by acting as the objective function being optimized. Our approach expedites convergence and makes the complex algorithm more deterministic.

Revisiting Genetic Algorithms: Genetic algorithms are a type of meta-heuristic algorithms used to solve non-linear optimization problems. In our case, we want to optimise the vector \mathbf{d}_{meta}^* on the non-linear objective \mathcal{S} , such that the score $\mathcal{S}(\mathbf{d}_{meta}^*)$ is maximised. All genetic algorithms follow a three-step framework: 1) *Initialisation*: where random vectors in the search space (\mathbf{d}_{meta}) are sampled, 2) *Exploitation*: where the “fitness” of the samples is carried out using the scorer and only the “fittest” samples are retained, and 3) *Exploration*: where new samples are created with the knowledge of the fittest samples to explore a diverse search space, through crossover and mutation steps.

These steps are carried on until a specified number of iterations or stopping criteria is met. NSGA-II [13] introduces a faster filtering operation for selecting the “fittest” samples. Deb *et al.* [11] was an augmentation of the previous work and it tailor makes an algorithm specifically for multi-objective optimization. However, many works [25, 32] argue that this is at the cost of appropriate search space exploration, which is detrimental to overall performance. Since we are explicitly working on single-objective optimization, we use its predecessor. We adapt NSGA-II to our setup by replacing its crossover mechanism with a novel design-specific crossover approach, explained next.

SWAN: Design Specific Crossover with Smart Snapping

An overview of the approach is presented in Algorithm 1 and Fig. 2. In line 1, two parents (vectors in the search space, $\mathbf{d}_{meta}^{P1}, \mathbf{d}_{meta}^{P2} \in \mathbf{d}_{meta}$) are randomly chosen from the set of “fit” samples (samples remaining after the exploitation step). The objective of crossover function is to

Algorithm 1 Design Specific Crossover

Input: Design attributes of Parent 1: \mathbf{d}_{meta}^{P1} ; Design attributes of Parent 2: \mathbf{d}_{meta}^{P2} ;
Output: Design attributes of Child: $\mathbf{d}_{meta}^{child}$

- 1: $\mathbf{V} \leftarrow \text{Sample}(\mathbf{d}_{meta}^{P1}, \mathbf{d}_{meta}^{P2})$ \triangleright *Initialise Parent mapping*
- 2: $\mathbf{L1} \leftarrow \mathbf{d}_{meta}^{P1}(V == 0)$ \triangleright *Parent1 elements mapping*
- 3: $\mathbf{L2} \leftarrow \mathbf{d}_{meta}^{P2}(V == 1)$ \triangleright *Parent2 elements mapping*
- 4: $\mathbf{C} = \mathbf{L1}$ \triangleright *Initialise Current Canvas*
- 5: **for** $\mathbf{Element} \in \mathbf{L2}$ **do**
- 6: **if** $\text{Type}(\mathbf{Element}) == \text{Image or SVG}$ **then**
- 7: $\mathbf{G} \leftarrow \text{GridGeneration}(\mathbf{C})$ \triangleright *Grid Line Generation*
- 8: $\mathbf{Slot} \leftarrow \text{FindSlot}(\mathbf{G})$ \triangleright *Find Appropriate Slot*
- 9: $\mathbf{Elem}^* \leftarrow \text{Snap}(\mathbf{Element}, \mathbf{Slot})$ \triangleright *Refinement*
- 10: **else if** $\text{Type}(\mathbf{Element}) == \text{Text}$ **then**
- 11: $\mathbf{E} \leftarrow \text{EulerDistances}(\mathbf{Element}, \mathbf{C})$
- 12: $\mathbf{Ele}_{min} \leftarrow \text{MapElement}(\mathbf{C}, \min(\mathbf{E}))$
- 13: $\mathbf{Elem}^* \leftarrow \text{MinAlign}(\mathbf{Element}, \mathbf{Ele}_{min})$
- 14: $\mathbf{C} \leftarrow \mathbf{C.append}(\mathbf{Elem}^*)$ \triangleright *Update Canvas*
- 15: $\mathbf{d}_{meta}^{child} = \mathbf{C}$

fuse the knowledge of both the parents to create a “smarter” offspring, which becomes a new sample in the population. In regular crossover, there is blind copy-pasting of parent data, and we find that to be unfit for our task at hand, which motivates us to propose SWAN: Design Specific Crossover with Smart Snapping. The first step is to initialize a random mask, \mathbf{V} , which assigns what parent is responsible for the percolation of which element in the new Child vector ($\mathbf{d}_{meta}^{child}$). Next, we copy element attributes from the first parent $\mathbf{L1}$ using the mask \mathbf{V} and build our current canvas \mathbf{C} in Line 4 and the third column in Fig. 2. For the remaining image elements from the second parent, we build grid lines from the edges of the design elements inside the current canvas and find all the boxes thus made after the intersection of the lines. We then find the box with the most similar size and proximity to the element from Parent2, and then ‘Snap’ it inside it and update the Canvas as shown in Fig. 2. For text elements, we find the Euler distances from the center of the text and all the other centers of elements from the current canvas. We find the element with least distance and find out what the cost (distance moved) is of either x-aligning or y-aligning with the element. We choose the minimum of it and update the current canvas to continue.

The mutation step in NSGA-II can also be improved. Standard mutation operation, which adds Gaussian noise to elements, might result in partial overflow in the design elements of their canvas size. We identify such cases and add noise only if it is within the canvas area. This allows our method to skip over degenerate cases.




Figure 2. A visual illustration of SWAN: Design Specific Crossover with Smart Snapping. Given two parents, SWAN first randomly decides which element to pick from either of the parents to generate the child. Then, it copies over the content from the first parent to the current canvas. Next, it identifies potential area within the canvas to host elements from the second parent, guided by grid-lines. Finally, the content from the second parent is ‘snapped’ into the identified areas by changing its attributes. This allows SWAN to generate better results when compared with regular crossover.



Figure 3. Samples from Crello dataset and their layout encodings.

4. Experiments and Results

Datasets: We utilize the widely-used Crello [59] dataset for our evaluation. It contains 23182 design templates. The data is divided into 18768 / 2315 / 2278 examples for train, validation, and test splits. We use multiple guided perturbations techniques to make the design pairs as discussed in Sec. 3.1. If the evaluation dataset includes the same or similar transformations used in training, we call it a *biased* setting; otherwise, it is called *unbiased*. We employ two *unbiased* setups: color and cross-match. In the color dataset, the transformation used makes the design’s color scheme bad by recoloring certain elements of it. We traverse the design, find overlapping elements (say, SVG preceding a Text), we extract the colors: C_{SVG}, C_{text} and randomly choose one of the elements (say, the SVG) and recolor it such that the $CIELAB.Distance(C_{text}, C_{SVG}^*) \in (2, 3)$. This ensures that there are distinct colours but it makes the design unaesthetic. In the cross-match dataset, we randomly pair a *good* version of a design x with the *bad* version of a design y , where $x \neq y$.

Baselines: We compare our *scorer* with recent Multimodal LLM-based design scorers. Specifically, we com-

pare against GPT-4o [40] and LLaVA-NeXT [35]. We compare the *refiner* against leading design refinement approaches. Specifically, we compare against SmartText++ [31], FlexDM [22], and COLE [27] for text box placement (refine-text setting); and against CanvasVAE [59], FLeXDM [22], DocLap [64], GPT-4 [41], and GPT-4V [39] for full design refinement (refine-all setting).

Implementation Details: We employ a learning rate of $1e^{-4}$, the Adam optimizer [15] with gradient coefficients of [0.5, 0.99], L2 regularization with weight decay as 0.005 and every 5 epochs, we schedule the learning rate to divide by half. We set a hard margin, $m = 0.2$, and loss parameters $\alpha = 0.8$ and $\beta = 0.2$ (see Eq. (2)). We add sensitivity analysis on these hyper-parameters in Sec. 5.4. All the convolution layers in the scorer model are as follows (representing the number of filters, kernel size, and stride, respectively): (64, 3, 1). We use group-normalization [58] (ngroup=2) in our scoring module. For SWAN, we find that $population_size = 100$, $n_trials = 1500$, $p = 0.3$ gives good performance.

Evaluation Metrics. We introduce Rank Accuracy (RAcc) to evaluate the design scoring methods. The higher the RAcc, the better the scorer. We define the Rank Accuracy as follows:

$$RAcc = \frac{1}{N} \sum_{i=1}^N \mathbb{I}\{\mathcal{S}(g_i) > \mathcal{S}(b_i)\}; \quad (6)$$

If the score given by a scorer \mathcal{S} to a *good* design (g_i) is more than that of a *bad* design (b_i) in a (g_i, b_i) pair, we consider it as correctly classified.

Following FlexDM [22], we use mean Intersection over Union (mIOU) and mean boundary displacement error [31]

(mBDE) to evaluate our refinement module for refine-text experiment. Following CanvasVAE [59], we use Type-wise mIoU (T-mIoU) to evaluate our refinement module for refine-all settings. In T-mIoU, first we calculate the overall IoU for each type of design element (text, image, SVG). Then we take the average of IoU of all the types.

$$\text{T-mIoU} = \frac{1}{N} \sum_{i=1}^N \frac{\sum_{j=1}^{n_i} x_{ij} \cap y_{ij}}{\sum_{j=1}^{n_i} x_{ij} \cup y_{ij}}; \quad (7)$$

where element x_{ij} belongs to type i and is j th element of the ground truth design. y_{ij} belongs to type i and is the j th element of refined design. n_i is total number of elements of type i and N is total number of types.

4.1. Evaluating the Scorer

We use three settings (biased, unbiased color, and unbiased cross-match) for evaluation. Each data instance in these datasets is a pair of designs, where the task of the model is to provide a better score to the *good* design.

We use the following prompt for the LLM-based evaluators: *"I will show you two designs, and you should give each design a design score between 1-100, which follows design principles and reason, and justify your score briefly and succinctly and then output which design has the higher score"*. See Tab. 2 for results. Despite having only $\sim 410k$ parameters compared to the billions in multi-modal LLMs, our model achieves the highest RAcc across all three datasets. This shows that multi-modal LLM-based evaluation of graphic design [9, 26] is not good at gauging micro-aesthetic differences between designs [17].

4.2. Evaluating the Refiner

Following earlier works [22, 27], we consider two settings: refining a single text box and refining whole design.

Table 4. When compared to state-of-the-art approached that refines placement of text on a design, Design-o-meter outperforms them in SingleText and MultipleText settings. Further, we ablate SWAN to find that it significantly helps to improve performance.

	SingleText		MultipleText	
	mIoU (\uparrow)	mBDE (\downarrow)	mIoU (\uparrow)	mBDE (\downarrow)
SmartText++	0.047	0.262	0.023	0.300
FlexDM	0.357	<u>0.098</u>	0.110	0.141
COLE	<u>0.402</u>	-	0.172	-
Ours w/o SWAN	0.376 ± 0.01	0.116 ± 0.01	0.311 ± 0.05	0.128 ± 0.01
Design-o-meter	0.42 ± 0.01	0.08 ± 0.01	0.38 ± 0.05	0.06 ± 0.01

Table 2. Comparison with LMM evaluators in biased, unbiased color (UC) and unbiased cross-match (UCM) setting.

Setting	Model	RAcc (\uparrow)	Params
Biased	LLaVA-NeXT	34.17	7B
	GPT-4o	<u>68.84</u>	-
	Design-o-meter	94.97	$\sim 410k$
UC	LLaVA-NeXT	44.22	7B
	GPT-4o	<u>72.86</u>	-
	Design-o-meter	90.45	$\sim 410k$
UCM	LLaVA-NeXT	31.00	7B
	GPT-4o	<u>63.50</u>	-
	Design-o-meter	87.50	$\sim 410k$

Text Box Refinement (Refine-Text): We divide the experiment into two settings: SingleText and MultiText. SingleText refers to the evaluation of designs containing only one text element, whereas MultiText refers to the evaluation of designs containing multiple text elements. The task is to predict the correct position and size of the target text element, keeping the aspect ratio of the element fixed. Our settings for the experiment are as follows: We choose a random text element as the target element. We randomly initialize the values of the center coordinates and scale of the target. We optimize these values using our refinement module. We summarize the quantitative results in table Tab. 4. In the second-last row, we selectively turn-off SWAN from Design-o-meter. The experiment reveals that our approach outperforms state-of-the-art approaches, and showcases the efficacy of SWAN. We add qualitative results in Fig. 4 (b).

Full Refinement (Refine-All): In this experiment we randomly initialize the values of the center coordinates and scale of all elements except the background element. Tab. 3 showcases comparison with 7 baseline approaches, and Design-o-meter comfortably outperforms all of them. We add qualitative results in Fig. 4 (b).

5. Discussion and Analysis

5.1 Visualizing the scorer: We conduct a simple occlusion based sensitivity analysis to measure the contribution of each region of a design towards the design score. We slide a $60 \text{ px} \times 60 \text{ px}$ occluding window across the design, replacing the pixels within the window with the design’s mean pixel value at each location. This modified design is then evaluated to generate a new score. The difference between the new score and the original score is used as the pixel value for the window’s center in the sensitivity map. This process creates a pixel-wise sensitivity map highlighting areas of the design that most positively or negatively impact the design score. The visualization in Fig. 5 shows that the scorer is indeed paying attention to the key areas of the design while making its decision.

5.2 On the Efficacy of Layout Encoding: Including layout color encoding in the input to the scorer model makes it easy for the model to capture the basic design principles such as overlap and alignment. We experiment with three configurations to verify this hypothesis: only rendering, only layout, rendering + layout; and train the model on Crello dataset. We notice that using rendering + layout as input performs the best with RAcc of 94.97 while the other two configurations give 86.65 and 86.99 RAcc.

5.3 Effect of Normalization Layers: Every graphic design is unique and complex. Assuming that designs belong to a common distribution and using normalization methods like BatchNorm [24] degrades the performance of the model.

Table 3. Comparison with state-of-the-art approaches when all elements of a design are refined (Refine-All setting).

	CanvasVAE	FlexDM	DocLap	GPT-4 0-shot	GPT-4 1-shot	GPT-4V 0-shot	GPT-4V 1-shot	Design-o-meter
T-mIOU (\uparrow)	42.39	<u>50.08</u>	43.75	30.75	29.97	28.81	35.17	54.44

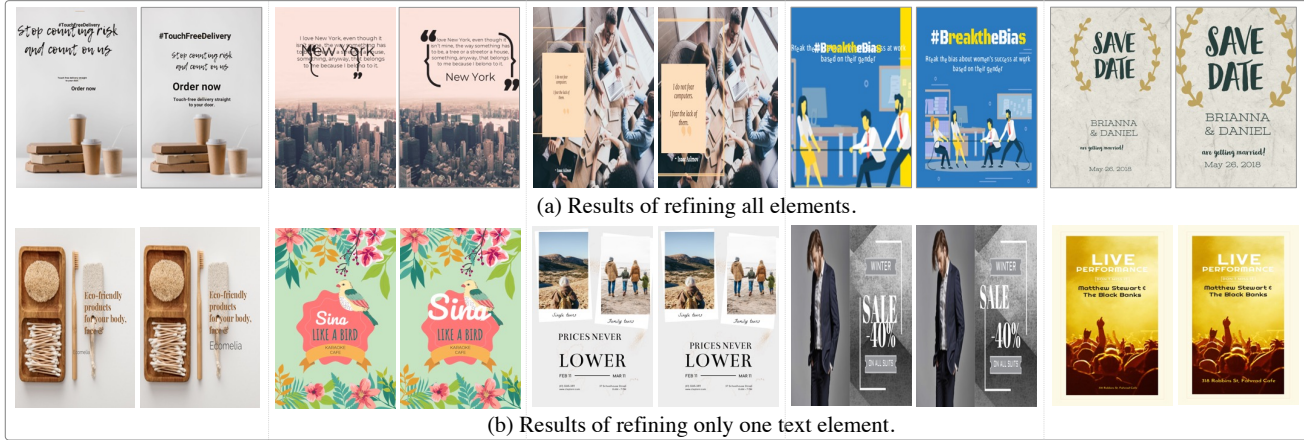


Figure 4. Qualitative results of Design-o-meter refining graphic designs. The top sub-figure refines all elements of a design (Refine-All setting), while the bottom one refines a single text box (Refine-Text setting). We see that our approach is able to automatically improve the position and scale parameters of design elements, making them more visually appealing.




Figure 5. These occlusion-based sensitivity maps show the locations of the design with positive (red) or negative (blue) impacts on a design score prediction. The numbers correspond to scores.

We experiment with BatchNorm [24], InstanceNorm [54], LayerNorm [2], and GroupNorm [58]. Table Tab. 5 shows the results. We ob-

Table 5. We find that GroupNorm gives the best results with the scorer.

Method	RAcc (\uparrow)
CNN + BatchNorm	68.10
CNN + LayerNorm	82.36
CNN + InstanceNorm	85.18
CNN + GroupNorm	94.97

serve that GroupNorm performs best along with the rendering+layout as input.

5.4 Sensitivity Analysis: We experiment with different types of margins in Eq. (3). Hard margin (H-Margin), Table 6. Varying α and β .

α	β	RAcc \uparrow
1.0	0.0	91.50
0.9	0.1	93.90
0.8	0.2	94.97

transformation-based margin (TB-Margin), and adaptive margin (Ada-Margin) gave 94.97, 94.10 and 93.65 RAcc values. We use H-Margin for the results. Next, we analyze the contribution of α and β parameters in Eq. (2). Decreasing α and increasing β , seems to have a positive effect, and hence we use $\alpha = 0.8$ and $\beta = 0.2$ for all experiments.

6. Conclusion

In this paper, we propose Design-o-meter, a novel framework which can evaluate and refine graphic designs. It takes an existing user design as input, and provide a unified design score. Further, it refines the input design to improve the design score. Our exhaustive experimental analysis brings out the efficacy of Design-o-meter. It would be interesting to see how our scorer can aid other layout and design generation frameworks as an off-the-shelf discriminator. We leave this for future exploration. We hope Design-o-meter will kindle interest in this practical and relevant setting of evaluating and refining graphic designs.

References

- [1] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Mike Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *European Conference on Computer Vision*, pages 456–473. Springer, 2022. 2
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization, 2016. 8
- [3] Adrien Bardes, Jean Ponce, and Yann LeCun. Vireg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021. 2
- [4] Michael Bauerly and Yili Liu. Computational modeling and experimental investigation of effects of compositional elements on interface and design aesthetics. *International journal of human-computer studies*, 64(8):670–682, 2006. 2
- [5] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a” siamese” time delay neural network. *Advances in neural information processing systems*, 6, 1993. 2
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 2
- [7] Shang Chai, Liansheng Zhuang, and Fengying Yan. Layoutdm: Transformer-based diffusion model for layout generation, 2023. 1
- [8] Jian Chen, Ruiyi Zhang, Yufan Zhou, Rajiv Jain, Zhiqiang Xu, Ryan Rossi, and Changyou Chen. Towards aligned layout generation via diffusion model with aesthetic constraints, 2024. 1
- [9] Yutao Cheng, Zhao Zhang, Maokke Yang, Hui Nie, Chunyuan Li, Xinglong Wu, and Jie Shao. Graphic design with large multimodal model. *arXiv preprint arXiv:2404.14368*, 2024. 3, 7
- [10] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *Asian conference on computer vision*, pages 251–263. Springer, 2016. 2
- [11] Kalyanmoy Deb and Himanshu Jain. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: solving problems with box constraints. *IEEE transactions on evolutionary computation*, 18(4):577–601, 2013. 5
- [12] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. 3
- [13] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197, 2002. 5
- [14] Qi Dou, Xianjun Sam Zheng, Tongfang Sun, and Pheng-Ann Heng. Webthetics: quantifying webpage aesthetics with deep learning. *International Journal of Human-Computer Studies*, 124:56–66, 2019. 2, 4
- [15] John Duchi, Elad Hazan, and Yoram Singer. Adaptive sub-gradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011. 6
- [16] Sahil Goyal, Shagun Uppal, Sarthak Bhagat, Yi Yu, Yifang Yin, and Rajiv Ratn Shah. Emotionally enhanced talking face generation. *arXiv preprint arXiv:2303.11548*, 2023. 2
- [17] Daichi Haraguchi, Naoto Inoue, Wataru Shimoda, Hayato Mitani, Seiichi Uchida, and Kota Yamaguchi. Can gpts evaluate graphic design based on design principles?, 2024. 7
- [18] Steven J Harrington, J Fernando Naveda, Rhys Price Jones, Paul Roetling, and Nishant Thakkar. Aesthetic measures for automated document layout. In *Proceedings of the 2004 ACM symposium on Document engineering*, pages 109–111, 2004. 2
- [19] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 1
- [20] Daichi Horita, Naoto Inoue, Kotaro Kikuchi, Kota Yamaguchi, and Kiyoharu Aizawa. Retrieval-augmented layout transformer for content-aware layout generation, 2024. 1
- [21] Hsiao Yuan Hsu, Xiangteng He, Yuxin Peng, Hao Kong, and Qing Zhang. Posterlayout: A new benchmark and approach for content-aware visual-textual presentation layout. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6018–6026, 2023. 1
- [22] Naoto Inoue, Kotaro Kikuchi, Edgar Simo-Serra, Mayu Otani, and Kota Yamaguchi. Towards flexible multi-modal document models, 2023. 1, 3, 6, 7
- [23] Naoto Inoue, Kento Masui, Wataru Shimoda, and Kota Yamaguchi. Opencole: Towards reproducible automatic graphic design generation, 2024. 1
- [24] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift, 2015. 7, 8
- [25] Hisao Ishibuchi, Ryo Imada, Yu Setoguchi, and Yusuke Nojima. Performance comparison of nsga-ii and nsga-iii on various many-objective test problems. In *2016 IEEE Congress on Evolutionary Computation (CEC)*, pages 3045–3052. IEEE, 2016. 5
- [26] 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. 3, 7
- [27] Peidong Jia, Chenxuan Li, Yuhui Yuan, Zeyu Liu, Yichao Shen, Bohan Chen, Xingru Chen, Yinglin Zheng, Dong Chen, Ji Li, Xiaodong Xie, Shanghang Zhang, and Baining Guo. Cole: A hierarchical generation framework for multi-layered and editable graphic design, 2024. 1, 6, 7
- [28] Sourabh Katoch, Sumit Singh Chauhan, and Vijay Kumar. A review on genetic algorithm: past, present, and future. *Multimedia tools and applications*, 80:8091–8126, 2021. 3
- [29] Shu Kong, Xiaohui Shen, Zhe Lin, Radomir Mech, and Charless Fowlkes. Photo aesthetics ranking network with attributes and content adaptation. In *Computer Vision—ECCV*

- 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, *Proceedings, Part I 14*, pages 662–679. Springer, 2016. 2
- [30] Wenyuan Kong, Zhaoyun Jiang, Shizhao Sun, Zhuoning Guo, Weiwei Cui, Ting Liu, Jian-Guang Lou, and Dongmei Zhang. Aesthetics++: Refining graphic designs by exploring design principles and human preference. *IEEE Transactions on Visualization and Computer Graphics*, 2022. 2, 3
- [31] Chenhui Li, Peiying Zhang, and Changbo Wang. Harmonious textual layout generation over natural images via deep aesthetics learning. *IEEE Transactions on Multimedia*, 24:3416–3428, 2021. 6
- [32] Hui Li, Kalyanmoy Deb, Qingfu Zhang, Ponnuthurai N Suganthan, and Lei Chen. Comparison between moea/d and nsga-iii on a set of novel many and multi-objective benchmark problems with challenging difficulties. *Swarm and Evolutionary Computation*, 46:104–117, 2019. 5
- [33] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR, 2022. 4
- [34] Jianan Li, Jimei Yang, Jianming Zhang, Chang Liu, Christina Wang, and Tingfa Xu. Attribute-conditioned layout gan for automatic graphic design. *IEEE Transactions on Visualization and Computer Graphics*, 27(10):4039–4048, 2020. 1
- [35] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, January 2024. 3, 6
- [36] David Chek Ling Ngo and John G Byrne. Another look at a model for evaluating interface aesthetics. *International Journal of Applied Mathematics and Computer Science*, 11(2):515–535, 2001. 2
- [37] David Chek Ling Ngo, Lian Seng Teo, and John G Byrne. A mathematical theory of interface aesthetics. In *Visual mathematics*, volume 2. Mathematical Institute SASA, 2000. 2
- [38] Peter O’Donovan, Aseem Agarwala, and Aaron Hertzmann. Designscape: Design with interactive layout suggestions. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pages 1221–1224, 2015. 3
- [39] OpenAI. Chatgpt can now see, hear, and speak. <https://openai.com/index/chatgpt-can-now-see-hear-and-speak/>, 2023. 6
- [40] OpenAI. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>, 2024. 3, 6
- [41] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Felipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolaus Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Valone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward,

- Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024. [6](#)
- [42] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. [4](#)
- [43] Peter O’Donovan, Aseem Agarwala, and Aaron Hertzmann. Learning layouts for single-pagegraphic designs. *IEEE transactions on visualization and computer graphics*, 20(8):1200–1213, 2014. [2, 3](#)
- [44] Xufang Pang, Ying Cao, Rynson WH Lau, and Antoni B Chan. Directing user attention via visual flow on web designs. *ACM Transactions on Graphics (TOG)*, 35(6):1–11, 2016. [3](#)
- [45] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. A lip sync expert is all you need for speech to lip generation in the wild. In *Proceedings of the 28th ACM international conference on multimedia*, pages 484–492, 2020. [2](#)
- [46] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [4](#)
- [47] Soliha Rahman, Vinoth Pandian Sermuga Pandian, and Matthias Jarke. Ruite: Refining ui layout aesthetics using transformer encoder. In *26th International Conference on Intelligent User Interfaces-Companion*, pages 81–83, 2021. [3](#)
- [48] Katharina Reinecke and Krzysztof Z Gajos. Quantifying visual preferences around the world. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 11–20, 2014. [2](#)
- [49] Shoko Sawada, Tomoyuki Suzuki, Kota Yamaguchi, and Masashi Toyoda. Visual explanation for advertising creative workflow. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pages 1–8, 2024. [2, 4](#)
- [50] David Schölgens, Sven Müller, Christine Bauer, Roman Tilly, and Detlef Schoder. Aesthetic measures for document layouts: Operationalization and analysis in the context of marketing brochures. In *Proceedings of the 2016 ACM Symposium on Document Engineering*, pages 21–30, 2016. [2](#)
- [51] Mohammad Amin Shabani, Zhaowen Wang, Difan Liu, Nanxuan Zhao, Jimei Yang, and Yasutaka Furukawa. Visual layout composer: Image-vector dual diffusion model for design layout generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9222–9231, 2024. [1](#)
- [52] Nidamarthi Srinivas and Kalyanmoy Deb. Multiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary computation*, 2(3):221–248, 1994. [3](#)
- [53] Sou Tabata, Hiroki Yoshihara, Haruka Maeda, and Kei Yokoyama. Automatic layout generation for graphical design magazines. In *ACM SIGGRAPH 2019 Posters*, pages 1–2. 2019. [2](#)
- [54] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization, 2017. [8](#)
- [55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [3, 4](#)
- [56] Hongyan Wan, Wanting Ji, Guoqing Wu, Xiaoyun Jia, Xue Zhan, Mengting Yuan, and Ruili Wang. A novel webpage layout aesthetic evaluation model for quantifying webpage layout design. *Information Sciences*, 576:589–608, 2021. [2, 4](#)
- [57] Haohan Weng, Danqing Huang, Yu Qiao, Zheng Hu, Chinyew Lin, Tong Zhang, and C. L. Philip Chen. Desigen: A pipeline for controllable design template generation, 2024. [1](#)
- [58] Yuxin Wu and Kaiming He. Group normalization, 2018. [6, 8](#)
- [59] Kota Yamaguchi. Canvasvae: Learning to generate vector graphic documents. *ICCV*, 2021. [4, 6, 7](#)
- [60] Zhenhua Yang, Dezhi Peng, Yuxin Kong, Yuyi Zhang, Cong Yao, and Lianwen Jin. Fontdiffuser: One-shot font generation via denoising diffusion with multi-scale content aggregation and style contrastive learning, 2023. [1](#)
- [61] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021. [2](#)
- [62] Mathieu Zen and Jean Vanderdonckt. Towards an evaluation of graphical user interfaces aesthetics based on metrics. In *2014 IEEE Eighth International Conference on Research Challenges in Information Science (RCIS)*, pages 1–12. IEEE, 2014. [2](#)
- [63] Nanxuan Zhao, Ying Cao, and Rynson WH Lau. What characterizes personalities of graphic designs? *ACM Transactions on Graphics (TOG)*, 37(4):1–15, 2018. [2](#)
- [64] Wanrong Zhu, Jennifer Healey, Ruiyi Zhang, William Yang Wang, and Tong Sun. Automatic layout planning for visually-rich documents with instruction-following models. *arXiv preprint arXiv:2404.15271*, 2024. [6](#)