# WebRPG: Automatic Web Rendering Parameters Generation for Visual Presentation

Zirui Shao<sup>1</sup>\*<sup>®</sup>, Feiyu Gao<sup>2</sup>\*<sup>®</sup>, Hangdi Xing<sup>1</sup><sup>®</sup>, Zepeng Zhu<sup>1</sup><sup>®</sup>, Zhi Yu<sup>1</sup>\*\*<sup>®</sup>, Jiajun Bu<sup>1</sup><sup>®</sup>, Qi Zheng<sup>2</sup><sup>®</sup>, and Cong Yao<sup>2</sup><sup>®</sup>

<sup>1</sup> Zhejiang Provincial Key Laboratory of Service Robot, Zhejiang University {shaozirui, xinghd, zapeng, yuzhirenzhe, bjj}@zju.edu.cn <sup>2</sup> Alibaba Group

feiyu.gfy@alibaba-inc.com, yongqi.zq@taobao.com, yaocong2010@gmail.com

Abstract. In the era of content creation revolution propelled by advancements in generative models, the field of web design remains unexplored despite its critical role in modern digital communication. The web design process is complex and often time-consuming, especially for those with limited expertise. In this paper, we introduce Web Rendering Parameters Generation (WebRPG), a new task that aims at automating the generation for visual presentation of web pages based on their HTML code. WebRPG would contribute to a faster web development workflow. Since there is no existing benchmark available, we develop a new dataset for WebRPG through an automated pipeline. Moreover, we present baseline models, utilizing VAE to manage numerous elements and rendering parameters, along with custom HTML embedding for capturing essential semantic and hierarchical information from HTML. Extensive experiments, including customized quantitative evaluations for this specific task, are conducted to evaluate the quality of the generated results. The dataset and code can be accessed at GitHub<sup>1</sup>.

**Keywords:** Generative model  $\cdot$  Visual Design Automation  $\cdot$  Web Rendering Parameters

# 1 Introduction

Recently, we are witnessing a revolution in content creation, driven by rapid advancements in generative models across domains such as image [21,48,55–57], text [3,42,62], and audio [6,7,32]. Numerous studies aim to leverage these advancements to enhance efficiency in graphic design, including advertisement [35, 40] and magazine [19,35,73] design. Nevertheless, the automation of web design, an essential part of graphic design [64], lacks exploration. Web design plays a significant role in the visual communication of web pages [61], impacting not

<sup>\*</sup> Equal contribution.

<sup>\*\*</sup> Corresponding author.

<sup>&</sup>lt;sup>1</sup> https://github.com/AlibabaResearch/AdvancedLiterateMachinery/tree/main/ DocumentUnderstanding/WebRPG



**Fig. 1:** Overview of the WebRPG task. The input consists of plain HTML code and the output comprises rendering parameters for each element. With browser rendering, plain HTML produces a disorganized visual presentation, while incorporating the generated rendering parameters significantly enhances the visual presentation.

only user satisfaction [9] but also user behavior [14]. Yet, it is a complex, timeconsuming task, especially challenging for those developers with limited design expertise, leading to substandard visual presentations [66]. Automating web design can simplify this process, enabling developers to create visually appealing web pages, and bridging the gap between technical development and aesthetic excellence.

Web pages are formed by  $\text{HTML}^2$  and  $\text{CSS}^3$  code, where HTML defines the content and structure, and CSS controls the visual presentation. With the advent of large language models (LLMs) [3,52,62], automating HTML code generation has become feasible. However, efforts in automatic visual presentation design, the core aspect of web design, currently center on specific subtasks such as layout generation [28, 49, 53], font recommendation [2, 71], and colorization [17, 27, 54], rather than designing a holistic web visual presentation from scratch.

Intuitively, leveraging generative models to learn design knowledge from existing web pages is a practical strategy for automated web visual design. However, the complexity of CSS coding practices poses challenges for its automatic generation [26]. To address this, we propose standardizing CSS using Rendering Parameters (RPs), which are defined by CSS properties that control the visual appearance of each web element [16]. Consequently, we introduce a novel task called **Web Rendering Parameters Generation (WebRPG** for short), which requires the automatic generation of rendering parameters for each web element based on the HTML code, as depicted in Fig. 1. With the help of a WebRPG system, HTML is the only prerequisite for obtaining an effective web visual presentation, which has the potential to achieve a faster web development workflow. With the integration of LLMs, a WebRPG system can even enable the realization of a fully automated web development workflow. Moreover, it can facilitate new applications, such as efficient exploration of various design options and dynamic personalization of web page styles.

Since there is no existing benchmark available for WebRPG, we develop automatic data processing steps to transform raw web pages into formalized We-

<sup>&</sup>lt;sup>2</sup> https://html.spec.whatwg.org/

<sup>&</sup>lt;sup>3</sup> https://www.w3.org/Style/CSS/specs.en.html

bRPG samples and construct a new dataset utilizing the Klarna dataset [22]. From a theoretical perspective, the WebRPG task presents two primary challenges: 1) Web pages comprise hundreds of elements, each with numerous RPs. 2) The visual presentation of web elements should be associated with the semantic and hierarchical information provided by HTML code. To address the challenges, variational autoencoder (VAE) [30] is employed to handle the large volume of rendering parameters for web elements, and specially designed HTML embedding is introduced to encode semantic and hierarchical information from HTML code. Using these modules, two WebRPG baselines are established, which are based on autoregressive and diffusion models, respectively. To verify the effectiveness of WebRPG baselines, metrics are designed to evaluate the overall appearance, layout, and style of the generated results. Both quantitative and qualitative experiments are conducted to assess the baselines.

Our main contributions are as follows:

- We introduce a novel task WebRPG for automatic web design from HTML code and create a new dataset.
- We explore the WebRPG task by establishing two baselines and propose solutions for its challenges.
- We design metrics to quantitatively evaluate the quality of generated results, and conduct qualitative experiments to analyze the strengths and weaknesses of the baselines.

# 2 Related Work

Generative models achieve notable success in image [4,50,55–57], text [3,42,62], and audio [6,7,32]. Image synthesis can create web visual presentations by generating screenshots but struggles with producing coherent text [55]. Moreover, image synthesis is limited to static images and cannot offer interactive, manipulable web pages.

Numerous efforts utilize generative models for graphic design, including advertising [35, 40], magazines [19, 23, 73], UI [8, 24, 25, 44], and posters [37, 74]. Yet, the designs restrict the element count to no more than 25. These methods primarily employ a one-dimensional sequence to represent designs, with each element defined by five tokens: four describe the bounding box, and one indicates the category (e.g., text, headline) [35]. However, the reliance on a simplistic flat input for the WebRPG task, which involves managing hundreds of elements and various RPs, leads to a substantial memory consumption increment, and performance degradation [12]. Moreover, the one-dimensional sequence neglects crucial hierarchical information in web pages.

Research focused on web pages has continuously emerged. In terms of understanding, efforts in web question answering [5, 72], web information extraction [34,69], and web pre-trained language models [10,41,60] have made notable progress in comprehending the essential semantic content and hierarchical structure of web pages. For instance, MarkupLM [41] stands out with its unique architecture and pre-training tasks, effectively encoding HTML content, which offers

insights for our research. Moreover, there are works aimed at web page design, such as optimizing the overall or specific block coloring of web pages [17,27,54], determining layouts based on given components like navigation bars [28,49,53], and recommending fonts for particular elements [2,71]. However, these studies focus only on specific subtasks of the web page design workflow, leaving the comprehensive design of web pages from scratch as an unexplored area.

## 3 Preliminary

#### 3.1 Task Definition

Web design is centered on visual presentation, i.e., the manipulation of CSS code. The complexity of CSS coding practices, including a wide range of selector options, makes the automatic generation of CSS code challenging [26]. To facilitate the model for learning web design, we standardize CSS by converting it into rendering parameters (RPs), which can be transformed back into CSS, with additional details in Sec. A.1. Consequently, the WebRPG task is defined as follows: given the HTML code, generate rendering parameters for each web element. Specifically, given a web page  $\mathcal{X}$ , whose HTML code is  $\mathcal{H}$ , it consists of a set of elements  $\mathcal{X} = \{X_1, X_2, \ldots, X_S\}$ , where S is the number of elements in  $\mathcal{X}$ . The visual appearance of element  $X_i$  is controlled by a set of RPs denoted as  $P_i = \{p_i^k \mid k \in \mathcal{W}\}$ , where  $\mathcal{W}$  indicates the indices for all RPs, and the complete set of RPs for  $\mathcal{X}$  is  $\mathcal{P} = \{P_1, P_2, \ldots, P_S\}$ . Therefore, the primary objective of the WebRPG task is to create a function f that generates RPs based on HTML code, that is,  $f : (\mathcal{H}) \mapsto \hat{\mathcal{P}}$ , where  $\hat{\mathcal{P}}$  represents the estimate of  $\mathcal{P}$ .

#### 3.2 Web Rendering Parameter Definition

The term "**Rendering Parameters (RPs)**" is employed to collectively describe the parameters controlling the visual appearance of each web element on the browser, as defined by CSS properties. Layout and visual style are crucial in the design of web pages [60,67], leading us to summarize 13 common CSS properties, divided into 3 categories as follows.

- Layout properties include *left*, top, width, and *height*.
- **Text properties** include *font-style*, *font-weight*, *font-size*, *line-height*, *text-align*, *text-decoration*, and *text-transform*.
- Color properties include *color* and *background-color*.

Various formats are available for web developers to define CSS properties. To standardize, we adopt the values computed by the browser [47] as the reference. Specifically, the values related to position and size are uniformly measured in integer pixels, and the values related to color correspond to 46 widely used colors. The vocabulary for all rendering parameters is available in Sec. A.1.



Fig. 2: Selected sub-page screenshots from our dataset. Notably, regions displayed are cropped due to space limitations.

## 4 Dataset Construction

#### 4.1 Data Pre-processing

Raw web pages cannot provide straightforward supervision for RPs. Thus, several pre-processings are conducted. Headless chrome<sup>4</sup> is used to render web pages and *selenium*<sup>5</sup> is employed to store HTML with only visible elements and record each element's selected CSS properties. Note that elements in this paper mean nodes in the  $DOM^6$  tree. The elements are stored following the DOM tree's preorder traversal. Since many web pages retain thousands of elements, we treat elements with a certain number of children as sub-pages with the semantic and hierarchical integrity preserved. The sub-pages are further cleaned while keeping the visual appearance, including removing uncommon HTML tags and intricate components like carousel images, as well as placing sub-pages at the top-left corner of the browser. Additionally, we only consider static components. Our models disregard the image on web pages, preserving only <img> tags. To guarantee data quality, a specific Visual Complexity (VC) metric is introduced to assist in filtering samples. The metric integrates three dimensions: color, size, and alignment, inspired by previous works [1, 15]. The definition of the VC metric is provided in Sec. A.2.

#### 4.2 Dataset Details

To accommodate the requirement for offline rendering, the Klarna dataset [22] is utilized to build our WebRPG dataset. The Klarna dataset, initially used for web information extraction, comprises 20K English product pages from 3K ecommerce sites, ensuring domain-specific diversity. The dataset stores all pages in MHTML<sup>7</sup> format, enabling offline rendering of the original pages in browsers with high fidelity.

The pre-processing in Sec. 4.1 is applied to the web pages with the browser canvas size setting to 1920\*1920 pixels, generating sub-pages containing between 32 and 128 child elements. The token length for each sample (sub-page) does not surpass 512. The size of RP vocabulary is 1993. The samples with a VC below

<sup>&</sup>lt;sup>4</sup> https://developer.chrome.com/blog/headless-chrome/

<sup>&</sup>lt;sup>5</sup> https://www.selenium.dev/

<sup>&</sup>lt;sup>6</sup> https://www.w3.org/DOM/DOMTR

<sup>&</sup>lt;sup>7</sup> https://en.wikipedia.org/wiki/MHTML





Fig. 3: Key components of WebRPG models. In the upper left, VAE compresses the RPs of each element into latent vectors shown in **blue**. In the top right, "Semantic" (Sem), "Hierarchical" (Hier), and "Character Count" (CharC) embeddings combine into the HTML embedding in **orange**. Below, two generative models are illustrated.

0.1 are filtered out. After preprocessing, our dataset includes 88,418 samples, split into training and testing sets at an 8:2 ratio. Our dataset exceeds the size of established graphic design datasets such as CLAY [38] (50K samples) and RICO [44] (43K samples), ensuring it can meet our objectives. Screenshots of some samples are shown in Fig. 2. More details are provided in Sec. A.3.

### 5 Methodology

#### 5.1 Overview

As indicated in Sec. 3.1, the WebRPG task is formulated as a function that generates rendering parameters (RPs) for each web element based on the HTML code. Inspired by classical generation methods [50, 56, 57], we employ a latent generation approach. In the approach, VAE is leveraged to compress all RPs of an element into latent space representation (Sec. 5.2), and a generative model (Sec. 5.4) generates the latent vector based on the given HTML embeddings (Sec. 5.3), which is then decoded back into RPs by the decoder of VAE. The key components of our method are shown in Fig. 3.

#### 5.2 Rendering Parameters Compression

Assume a web page consists of S elements, with the appearance of each element  $X_i$  determined by  $\mathcal{W}$  rendering parameters  $P_i = \{p_i^k \mid k \in \mathcal{W}\}$ . The WebRPG model necessitates the processing of  $S \times \mathcal{W}$  values for both input and output. Expanding all  $p_i^k$  of  $X_i$  into a one-dimensional sequence, as per graphic design methods [24,35], leads to excessively long input and output lengths. To mitigate this challenge, we utilize VAE to compress the rendering parameters into a latent space. This ensures that the input length for the generative model correlates solely with S.

More precisely, given the RPs of an element  $P_i \in \mathbb{R}^{W*V}$ , where V is the size of RPs vocabulary (Sec. 3.2), and the corresponding latent vector is  $Z_i \in \mathbb{R}^d$ . We denote the generative distribution as  $p_\theta(P_i \mid Z_i)$  and the posterior as  $q_\phi(Z_i \mid P_i)$ , respectively. The learning objective of VAE is expressed as:

$$L_{VAE} = \frac{1}{S} \cdot \sum_{i=1}^{S} (-\mathbb{E}_{q_{\phi}(Z_{i}|P_{i})} \left[ \log p_{\theta}(P_{i} \mid Z_{i}) \right] + \lambda_{KL} \mathrm{KL} \left( q_{\phi}(Z_{i} \mid P_{i}) \| p(Z_{i}) \right) \right), (1)$$

where  $\theta$  and  $\phi$  are the encoder and decoder parameters,  $\mathbb{E}$  indicates the expectation, KL is the Kullback-Leibler divergence, and  $\lambda_{KL}$  is the hyperparameter to balance the two terms. The encoder and decoder of VAE both consist of a multilayer perceptron with five layers. To ensure that the latent space encompasses as many element appearances (i.e., combinations of RPs) as possible, the VAE is pre-trained using synthetic data.

#### 5.3 Encoding HTML

The visual presentation of a web page should be in harmony with the content and structure dictated by its HTML code. To this end, we design an **HTML embedding** that captures the essential information in the HTML code, establishing the input feature for the generative model (Sec. 5.4). HTML code essentially encompasses hierarchical information among elements and the textual content of each element [10]. The character count of each element is also crucial, as the size of an element generally exhibits a positive correlation with the length of characters. Therefore, our HTML embedding integrates three facets of information: semantics, hierarchy, and character count. Precisely, for an element  $X_i$ , its HTML embedding  $H_i \in \mathbb{R}^d$  is defined as:

$$H_i = \Lambda^{\text{Sem}}(H_i^{\text{Sem}}) + \Lambda^{\text{Hier}}(H_i^{\text{Hier}}) + \Lambda^{\text{CharC}}(H_i^{\text{CharC}}), \qquad (2)$$

where  $H_i^{\text{Sem}}$ ,  $H_i^{\text{Hier}}$  and  $H_i^{\text{CharC}}$  denote the semantic, hierarchical and character count embedding respectively, and  $\Lambda^{\circ}()$  is the linear projection layer.

Semantic embedding: The MarkupLM<sub>large</sub> model [41], a language model explicitly pre-trained for web understanding, is employed as the semantic extractor. Specifically, given an element  $X_i$  with HTML code tokens  $X_i = \{x_i^j \mid j \in \mathcal{L}\}$ , where  $\mathcal{L}$  denotes the token length, we calculate the semantic embedding of  $X_i$  as  $H_i^{\text{Sem}} = \text{Pool}(\text{MarkupLM}(x_i^1, x_i^2, \dots, x_i^{\mathcal{L}}))$ , where Pool( $\cdot$ ) denotes an average pooling operation.

**Hierarchical embedding:** The XPath embedding layer [41] is employed to model the hierarchical information of elements, taking their XPath expressions as input. XPath<sup>8</sup> is a query language for selecting elements from a web page, which is based on the DOM tree and can be used to easily locate an element. Specifically, for an element  $X_i$  with its corresponding XPath expression  $xp_i$ , we compute the hierarchical embedding directly as  $H_i^{\text{Hier}} = \text{XPathEmb}(xp_i)$ .

<sup>&</sup>lt;sup>8</sup> https://www.w3.org/TR/xpath-31/

**Character count embedding:** We establish a mapping mechanism that translates the raw count of characters into dense vector space. For an element  $X_i$  with the content of k characters, the character count embedding is calculated as  $H_i^{\text{CharC}} = \text{EmbCharC}(k)$ .

#### 5.4 Generative Models

Two generative models are implemented: autoregressive and diffusion model.

Autoregressive Model (AR): To enhance the model stability during training, a masked latent vector  $Z_{mask}$  of real RPs is introduced inspired by BART [36] and MaskGIT [4].  $Z_{mask}$  is constructed in two steps. Firstly, the real RPs are encoded into the latent vectors with the VAE encoder, i.e.,  $Z = \theta(\mathcal{P})$ . Then a special MASK vector and a binary mask  $M = \{m_i \mid i \in S\}$  are utilized to partially substitute the real latent vectors with the MASK as  $Z_{mask,i} = m_i \cdot MASK + (1 - m_i) \cdot \theta(P_i)$ .

Here M is generated using a mask scheduling function  $\gamma(r) \in (0, 1]$  following MaskGIT [4], and the MASK vector is a learnable parameter with the same shape as  $Z_i$ . Additionally, it is important to highlight that during inference, all  $Z_i$  are masked, i.e.,  $M = \{m_i = 1 | 1 \le i \le S\}$ .

As depicted in Fig. 3, the model inputs the sum of  $\mathcal{Z}_{mask}$  and  $\mathcal{H}$  to generate  $\hat{\mathcal{Z}}$ , which is then decoded by the VAE decoder as  $\hat{\mathcal{P}} = \phi(\hat{\mathcal{Z}})$ . The VAE and generative models are trained jointly, thus the training loss is as follows:

$$L = \log p_{\psi}(\mathcal{P}|\mathcal{H}, \mathcal{Z}_{mask}) + L_{VAE}, \tag{3}$$

where  $\psi$  is the parameters of the generative model.

**Diffusion Model:** Diffusion models [21, 48, 75] have recently emerged as a new class of generative models with high performance. These models are characterized by forward and reverse Markov processes of length T. In our rendering parameters compression (VAE) model, rendering parameters  $\mathcal{P}$  are encoded into a latent space, i.e.,  $\mathcal{Z} = \theta(\mathcal{P})$ . These latent vectors  $\mathcal{Z}$ , which align more closely with a Gaussian distribution, improve compatibility with the noise distribution in diffusion models. Following successful models [11, 57, 59], our diffusion model can be interpreted as an equally weighted sequence of denoising autoencoders  $\mathcal{E}(\mathcal{Z}_t, t, \mathcal{H}); t = 1 \dots T$ , which are trained to predict the noise  $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  in  $\mathcal{Z}_t$ . The  $\mathcal{Z}_t$  is obtained from a forward process starting from  $\mathcal{Z}_0$  (where  $\mathcal{Z}_0 = \mathcal{Z}$ ), defined as  $\mathcal{Z}_t = \sqrt{\alpha_t}\mathcal{Z}_{t-1} + \sqrt{1 - \alpha_t}\boldsymbol{\epsilon}$ , with  $\alpha_t$  being a predefined set of coefficients. As illustrated in Fig. 3,  $\mathcal{Z}_t$  and  $\mathcal{H}$  are added and input into the model. Our diffusion model employs the standard variational lower bound objective as its training loss, and we jointly optimize the VAE, leading to the overall loss function:

$$L = \mathbb{E}_{\mathcal{Z}, \epsilon \sim \mathcal{N}(0,1), t} \left[ \| \boldsymbol{\epsilon} - \boldsymbol{\epsilon}_{\psi}(\mathcal{Z}_t, t, \mathcal{H}) \|_2^2 \right] + L_{VAE}.$$
(4)

During inference, the predicted  $\hat{Z}$  is progressively obtained through a reverse process, expressed as  $Z_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( Z_t - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}} \epsilon_{\psi}(Z_t, t, \mathcal{H}) \right)$ . Subsequently,  $\hat{Z}$  is decoded to  $\hat{\mathcal{P}}$  via a single pass through the VAE decoder  $\phi$ . Additionally,  $Z_T$  is random Gaussian noise.

### 6 Experiment

#### 6.1 Evaluation Metrics

Three metrics are utilized to assess the quality of the generated rendering parameters. Fréchet Inception Distance (FID), Element Intersection over Union (Ele. IoU), and newly introduced Style Consistency Score (SC Score) enable the evaluation of the overall appearance, layout, and style of generated web pages respectively. As indicated in Sec. 3.2, "layout" refers to layout properties, while "style" encompasses text properties and color properties.

**Fréchet Inception Distance** FID [20], a metric initially proposed in the domain of image generation, measures the similarity of generated data to real ones in feature space. Inspired by Lee et al. [35], a binary classifier is trained to distinguish between real and noise-added RPs. This classifier is employed to generate representative features of RPs for calculating FID. We also introduce layout-specific and style-specific FID models. Further details are in Sec. A.4.

**Elements Intersection over Union** Ele. IoU is a metric for evaluating the similarity between generated layouts and real ones, based on adaptation to the Maximum IoU [29]. As the elements of real and generated web pages correspond one-to-one, IoU is computed between the corresponding pairs. Denote the real layouts as  $B = \{b_i\}_{i=1}^N$  and the generated ones as  $\hat{B} = \{\hat{b}_i\}_{i=1}^N$ , with N being the element count, and  $\hat{b}_i$  and  $\hat{b}_i$  as corresponding elements. The Ele. IoU can be calculated as follows:

$$EleIoU(B, \hat{B}) = \frac{1}{N} \sum_{i=1}^{N} IoU(b_i, \hat{b}_i).$$
(5)

**Style Consistency Score** The "Principle of Similarity" of Gestalt theory suggests that people tend to perceive elements with similar style as a whole [31,65], highlighting the importance of style consistency among elements. Hence, the SC Score assesses whether elements with the same style on a real web page retain that consistency on the generated page, beyond merely visual similarity. An example explanation is provided in Sec. C. Elements are deemed to have the same style only if all their style properties are identical [60]. Specifically, for a web page  $W = \{e_i \mid i \in N\}$  with N being the number of elements, the style consistency subset of the page is defined as  $S \subseteq W, \forall e_i, e_j \in S, style(e_i) = style(e_j)$ .

Thus the real web page W and its generated page  $\hat{W}$  are divided into style consistency subsets  $W = \{S_j \mid j \in M\}$  and  $\hat{W} = \{S_k \mid k \in N\}$ , respectively. Given N and M can differ, we apply a max operation for optimal matching. The SC Score is then calculated as:

$$SCScore(W, \hat{W}) = \sum_{j=1}^{M} w_j \cdot \max_k J(S_j, \hat{S}_k),$$
(6)

where J(A, B) is the Jaccard similarity coefficient. Additionally, under the assumption that style consistency subsets with more elements are more semantically valuable, we utilize a weight  $w_j = \frac{|S_j|}{\sum_{l=1}^M |S_l|}$ .

#### 6.2 Implementation

Two baselines are implemented: autoregressive (WebRPG-AR) and diffusion model (WebRPG-DM). The VAE, hierarchical embedding, and character count embedding are jointly trained with the backbone, and the semantic embedding is produced by frozen pre-trained MarkupLM<sub>large</sub> [41]. The XPath embedding layer is initialized following Li et al. [41]. All baselines are based on Transformer [63] and have approximately 50M of parameters to ensure fair comparison, whose hidden dimensions are 128. The dimensions of latent vector and HTML embedding d is 128. For optimization, AdamW [45] is used with a learning rate of 1.2e-4. All models are trained for 1M steps with a batch size of 300.

Additionally, LLMs have been gaining adoption in different domains. We assess GPT-4 [51,70], StarCoder2-7b [46], DeepSeek-Coder-6.7b [18], CodeLlama-13B [58] on the WebRPG task. GPT-4 is one of the state-of-the-art LLMs, while the others are open-source models known for code generation. Due to limited resources, we randomly select 10% of test samples. The prompt template employs in-context learning [3], incorporating a task description, three demonstrations, and a test instance. Further details are available in Sec. A.5.

#### 6.3 Quantitative and Qualitative Evaluation

We present quantitative results in Tab. 1 and qualitative results in Fig. 4. Regarding the results of real data, in addition to the normally rendered web page (Tab. 1, "Real Web Pgae"), we also report the web page rendered using only HTML code (Tab. 1, "Plain HTML"). Since the browser would apply default CSS when custom CSS is absent, some models perform worser than the plain HTML due to unreasonable generated RPs. The FIDs for real data are calculated between the test set and other real web pages.

The experimental results show that WebRPG-AR consistently surpasses other baselines. Its sequential decoding mechanism allows for more refined control based on previously generated results [13,68]. As shown in Fig. 4 a, b, e, WebRPG-AR demonstrates impressive visual quality in detail.

The performance of WebRPG-DM is suboptimal across all metrics. It only tends to produce standard web visual presentations in simpler cases, as illustrated in Fig. 4 e, such as bolding prices, adding background color to buttons, and aligning a few elements. This implies that diffusion models may be inappropriate for this task. There are two plausible explanations: First, unlike images and videos in Euclidean space, web elements are non-Euclidean due to their hierarchical arrangement, while diffusion models are confined to Euclidean space [33]. Second, the WebRPG task demands meticulous adjustments and detailed control for realism, a limitation of diffusion models [43].

	Overall	Overall Layout		Style	
Model	$\mathrm{FID}\downarrow$	$\mathrm{FID}_{\mathrm{layout}}\downarrow$	Ele. Io U $\uparrow$	$\mathrm{FID}_{\mathrm{style}}\downarrow$	SC Score $\uparrow$
WebRPG-AR WebBPG-DM	<b>0.1281</b> 62.021	<b>0.1520</b> 60.942	<b>0.7082</b> 0.0357	<b>0.2124</b>	<b>0.9474</b> 0.3671
WebRPG-AR*	0.1324	0.2877	0.7069	0.1359	0.9485
$WebRPG-DM^*$	61.135	60.870	0.0356	105.86	0.3649
GPT4*	4.2141	47.732	0.0347	8.8898	0.5515
$StarCoder2-7b^*$	11.899	51.432	0.0309	18.186	0.3639
$DeepSeek-Coder-6.7b^*$	5.8219	55.744	0.0330	7.4542	0.3949
CodeLlama-13b*	9.2826	55.427	0.0278	11.625	0.3864
Real Web Page Plain HTML	$0.0027 \\ 8.5342$	$0.0015 \\ 52.438$	$1.0000 \\ 0.0354$	$0.0074 \\ 8.4951$	$1.0000 \\ 0.3668$

**Table 1:** WebRPG baselines quantitative comparison with bold figures for best results.

 "\*" stands for the result in the randomly selected test set.

**Table 2:** Ablation study based on WebRPG-AR. Best results in bold. " $\mathcal{Z}_{mask}$ " is detailed in Sec. 5.4. "H.E." stands for HTML embedding. "S.", "H.", and "C." stand for semantic, hierarchical, and character count embeddings.

#	VAE	$\mathcal{Z}_{mask}$	s.	H. E H.	с.	$\begin{array}{c} \text{Overall} \\ \text{FID} \downarrow \end{array}$	$\begin{array}{c} Layout\\ FID_{layout} \downarrow \end{array}$	Ele. IoU $\uparrow$	$\begin{array}{c} {\rm Style} \\ {\rm FID}_{\rm style} \downarrow \end{array}$	SC Score $\uparrow$
1		$\checkmark$	√ √	√ √	<b>√</b>	0.9702 0.1487	5.4668 0.2055	0.5954 0.6462	15.923 0.2944	0.8053 0.9332
3	 ✓	√	•	• √	▼ √	0.1407	0.2005	0.6620	0.3055	0.9323
$\frac{4}{5}$	$\checkmark$	$\checkmark$	√ √	$\checkmark$	$\checkmark$	$0.3003 \\ 0.1575$	$0.3770 \\ 0.3152$	$0.6345 \\ 0.6769$	$1.9048 \\ 0.3065$	$0.8982 \\ 0.9434$
6	$\checkmark$	√	√	√	$\checkmark$	0.1281	0.1520	0.7082	0.2124	0.9474

GPT-4's performance on the WebRPG task surpasses that of WebRPG-DM and falls short of WebRPG-AR. Open-source LLMs underperform compared to GPT-4. As illustrated in the Fig. 4 a,b,e, GPT-4 can effectively handle element styles, such as adding background colors to buttons and applying distinct colors for prices. However, the performance of GPT-4 in layout is limited. As demonstrated in Fig. 4 a-c, GPT-4 tends to generate simplistic vertical arrangements when faced with complex HTML structures. With regular HTML, as depicted in Fig. 4 e, GPT-4 achieves a layout that is similar to the real page. Therefore, we conclude that GPT-4 demonstrates basic capability in WebRPG tasks with regular HTML, but its performance with complex HTML is less effective. Additionally, we notice that LLMs do not generate RPs for all elements, causing many to use the browser's default CSS, resulting in performance similar to plain HTML.



Fig. 4: Qualitative comparison of WebRPG baselines.



Fig. 5: Case visualization from the ablation study.

It is worth noting that WebRPG-AR exhibits the ability to render diverse web pages. For example, Fig. 4 d shows WebRPG-AR's creation of a page with a vertical layout (originally horizontal), preserving the pattern and order consistency across four groups. This finding suggests that the model successfully learns web design knowledge and applies it effectively to render web pages from HTML code. Further cases are available in Sec. B.1.

Furthermore, we calculate the FID on screenshots of rendered web pages, following conventional image generation practices [55,57]. The results, shown in Sec. B.2, are consistent with Tab. 1. Additionally, we conduct a human evaluation, detailed in Sec. B.4, with results that also align with Tab. 1.



**Fig. 6:** *Left*: Trends in WebRPG-AR performance relative to the number of elements and average depth of elements within the DOM tree. *Right*: WebRPG-AR failure cases with real web pages on the left, generated results on the right, and highlights in green.

#### 6.4 Ablation Study

We conduct a series of ablation experiments based on WebRPG-AR, as shown in Tab. 2. #1 uses a one-dimensional flat input instead of VAE. #2 removes  $\mathcal{Z}_{mask}$  (in Sec. 5.4). #3 and #5 respectively remove the corresponding embedding layers, while #4 substitutes hierarchical embedding with one-dimensional positional embedding. Additionally, we visualize some cases from #3, #4, and #5 in Fig. 5. All models are trained to convergence following the settings in Sec. 6.2.

The results of #1 demonstrate the effectiveness of using VAE for rendering parameters compression. Although #2 is comparable to #6, the incorporation of  $Z_{mask}$  enhances the model stability during training. The results of #3, #4, and #5 reveal that all three embeddings play critical roles in web design. Hierarchical embedding helps layout arrangement significantly. The simplification to 1D positional embedding leads to a disorganized layout, as illustrated in Fig. 5 c. Semantic embedding enhances the model with the capacity to perceive semantic relationship. For example, as Fig. 5 b shows, the model struggles to horizontally align elements like "select a color" and "sunshine," suggesting challenges in identifying key-value pairs without semantic information. Character count embedding helps to predict appropriate element sizes for full content display, as in Fig. 5 d, where a narrow "price" and "sunshine" width leads to incomplete text display.

#### 6.5 Discussion on Failure Cases

To investigate the boundaries of the model's capabilities, we analyze several failure cases generated by WebRPG-AR. The left side of Fig. 6 reveals that both layout (Ele. IoU) and style (SC Score) metrics decrease with an increase in the number of elements or the average depth of elements within the DOM tree. This trend may be attributed to two factors: the inherent complexity of a page increases with more elements or greater depth, and the training set lacks web pages with a large number of elements or significant depths (details in Sec. A.3). Regarding error types, layout issues mainly include misalignments and overlaps, as shown in Fig. 6 a and b. For style, the model struggles to recognize web page elements with identical semantic functions, such as the "Add to Cart" buttons



**Fig. 7:** The HTML code generated by GPT-4 and the corresponding web page visual results generated by WebRPG-AR. Screenshots use **green** <img> placeholders.

illustrated in Fig. 6 c, which should appear identical. Moreover, we observe two primary error scenarios: elements positioned at the end of the HTML code tend to be more error-prone, as seen with the element in the bottom right corner of Fig. 6 b, likely due to the characteristics of the autoregressive model [12]; additionally, pages with large-scale images pose challenges, as shown in Fig. 6 a, since the model does not take the original images as input. The discussion above highlights the need for further research.

#### 6.6 Discussion on the Integration of LLM and WebRPG Model

Recently, LLMs have enabled the possibility of automatically generating HTML code [39]. Consequently, we hypothesize that integrating LLM into a WebRPG system could facilitate a fully automated web development workflow. We employ GPT-4 [51,70] to validate this hypothesis. As Fig. 7 illustrates, WebRPG-AR effectively creates visual presentations of web pages based on generated HTML, demonstrating the potential of a fully automated web development workflow through the integration of LLM and WebRPG. Additional cases and the prompt for automatically generating HTML are provided in Sec. B.3.

# 7 Conclusion and Limitations

This paper presents WebRPG, a task that automates web design by generating rendering parameters for web elements from HTML. We introduce a new dataset, two baseline models, and evaluation metrics. Results show the autoregressive baseline most effectively generates web visual presentations.

Nevertheless, this study has limitations that warrant further investigation in future research. The proposed model can undergo fine-tuning to support design tasks such as partial web page design by masking specific elements. Additionally, it can be adapted to analyze raster images by replacing <img> tokens with image embeddings. The employment of established CSS frameworks like Tailwind<sup>9</sup> could standardize CSS, thereby potentially simplifying the WebRPG task. However, sourcing web pages based on these frameworks presents challenges. Furthermore, design options and control mechanisms of the results are worth exploring. Future research will address these aspects.

<sup>&</sup>lt;sup>9</sup> https://tailwindcss.com/

15

### Acknowledgements

This work is supported by the National Natural Science Foundation of China (Grant No. 62372408) and the National Key R&D Program of China (No. 2021YFB2701100).

### References

- 1. Alemerien, K., Magel, K.: Guievaluator: A metric-tool for evaluating the complexity of graphical user interfaces. In: SEKE. pp. 13–18 (2014)
- Azadi, S., Fisher, M., Kim, V.G., Wang, Z., Shechtman, E., Darrell, T.: Multicontent gan for few-shot font style transfer. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7564–7573 (2018)
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. Advances in neural information processing systems 33, 1877–1901 (2020)
- Chang, H., Zhang, H., Jiang, L., Liu, C., Freeman, W.T.: Maskgit: Masked generative image transformer. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 11305–11315 (2022)
- Chen, L., Chen, X., Zhao, Z., Zhang, D., Ji, J., Luo, A., Xiong, Y., Yu, K.: Websrc: A dataset for web-based structural reading comprehension. In: Conference on Empirical Methods in Natural Language Processing (2021)
- Chen, N., Zhang, Y., Zen, H., Weiss, R.J., Norouzi, M., Chan, W.: Wavegrad: Estimating gradients for waveform generation. In: International Conference on Learning Representations (2021)
- Chen, N., Zhang, Y., Zen, H., Weiss, R.J., Norouzi, M., Dehak, N., Chan, W.: Wavegrad 2: Iterative refinement for text-to-speech synthesis. arXiv preprint arXiv:2106.09660 (2021)
- Cheng, C.Y., Huang, F., Li, G., Li, Y.: Play: parametrically conditioned layout generation using latent diffusion. In: Proceedings of the 40th International Conference on Machine Learning. ICML'23, JMLR.org (2023)
- Cyr, D., Head, M., Larios, H.: Colour appeal in website design within and across cultures: A multi-method evaluation. International journal of human-computer studies 68(1-2), 1–21 (2010)
- Deng, X., Shiralkar, P., Lockard, C., Huang, B., Sun, H.: Dom-lm: Learning generalizable representations for html documents. arXiv preprint arXiv:2201.10608 (2022)
- Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in neural information processing systems 34, 8780–8794 (2021)
- 12. Dong, Z., Tang, T., Li, L., Zhao, W.X.: A survey on long text modeling with transformers. arXiv preprint arXiv:2302.14502 (2023)
- Du, Y., Chen, Z., Jia, C., Yin, X., Li, C., Du, Y., Jiang, Y.G.: Context perception parallel decoder for scene text recognition. arXiv preprint arXiv:2307.12270 (2023)
- Flavian, C., Gurrea, R., Orus, C.: Web design: a key factor for the website success. Journal of Systems and Information Technology 11(2), 168–184 (2009)
- Fu, F., Chiu, S.Y., Su, C.H.: Measuring the screen complexity of web pages. In: Human Interface and the Management of Information. Interacting in Information Environments: Symposium on Human Interface 2007, Held as Part of HCI International 2007, Beijing, China, July 22-27, 2007, Proceedings, Part II. pp. 720–729. Springer (2007)

- 16 Z. Shao et al.
- 16. Furht, B. (ed.): Cascading Style Sheets, pp. 58–58. Springer US, Boston, MA (2008)
- Gu, Z., Lou, J.: Data driven webpage color design. Computer-Aided Design 77, 46–59 (2016)
- Guo, D., Zhu, Q., Yang, D., Xie, Z., Dong, K., Zhang, W., Chen, G., Bi, X., Wu, Y., Li, Y., et al.: Deepseek-coder: When the large language model meets programming– the rise of code intelligence. arXiv preprint arXiv:2401.14196 (2024)
- Herzig, R., Bar, A., Xu, H., Chechik, G., Darrell, T., Globerson, A.: Learning canonical representations for scene graph to image generation. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) Computer Vision – ECCV 2020. pp. 210–227. Springer International Publishing, Cham (2020)
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems 33, 6840–6851 (2020)
- 22. Hotti, A., Risuleo, R.S., Magureanu, S., Moradi, A., Lagergren, J.: The klarna product page dataset: A realistic benchmark for web representation learning. arXiv preprint arXiv:2111.02168 (2021)
- Hui, M., Zhang, Z., Zhang, X., Xie, W., Wang, Y., Lu, Y.: Unifying layout generation with a decoupled diffusion model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1942–1951 (2023)
- Inoue, N., Kikuchi, K., Simo-Serra, E., Otani, M., Yamaguchi, K.: Layoutdm: Discrete diffusion model for controllable layout generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10167– 10176 (2023)
- Jyothi, A.A., Durand, T., He, J., Sigal, L., Mori, G.: Layoutvae: Stochastic scene layout generation from a label set. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9895–9904 (2019)
- Kaluarachchi, T., Wickramasinghe, M.: A systematic literature review on automatic website generation. Journal of Computer Languages p. 101202 (2023)
- Kikuchi, K., Inoue, N., Otani, M., Simo-Serra, E., Yamaguchi, K.: Generative colorization of structured mobile web pages. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3650–3659 (2023)
- Kikuchi, K., Otani, M., Yamaguchi, K., Simo-Serra, E.: Modeling visual containment for web page layout optimization. In: Computer Graphics Forum. vol. 40, pp. 33–44. Wiley Online Library (2021)
- Kikuchi, K., Simo-Serra, E., Otani, M., Yamaguchi, K.: Constrained graphic layout generation via latent optimization. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 88–96 (2021)
- Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
- 31. Koffka, K.: Principles of gestalt psychology (1955)
- Kong, Z., Ping, W., Huang, J., Zhao, K., Catanzaro, B.: Diffwave: A versatile diffusion model for audio synthesis. arXiv preprint arXiv:2009.09761 (2020)
- Koo, H.: A survey on generative diffusion models for structured data. arXiv preprint arXiv:2306.04139 (2023)
- Kumar, V., Dhar, M., Khattar, D., Lal, Y.K., Mishra, A., Shrivastava, M., Varma, V.: Swde: A sub-word and document embedding based engine for clickbait detection. arXiv preprint arXiv:1808.00957 (2018)

Automatic Web Rendering Parameters Generation for Visual Presentation

17

- Lee, H.Y., Jiang, L., Essa, I., Le, P.B., Gong, H., Yang, M.H., Yang, W.: Neural design network: Graphic layout generation with constraints. In: Computer Vision– ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16. pp. 491–506. Springer (2020)
- 36. Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., rahman Mohamed, A., Levy, O., Stoyanov, V., Zettlemoyer, L.: Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In: Annual Meeting of the Association for Computational Linguistics (2019)
- Li, C., Zhang, P., Wang, C.: Harmonious textual layout generation over natural images via deep aesthetics learning. IEEE Transactions on Multimedia 24, 3416– 3428 (2022)
- Li, G.e.a.: Learning to denoise raw mobile ui layouts for improving datasets at scale. pp. 1–13 (2022)
- Li, J., Li, G., Li, Y., Jin, Z.: Enabling programming thinking in large language models toward code generation. arXiv preprint arXiv:2305.06599 (2023)
- Li, J., Yang, J., Zhang, J., Liu, C., Wang, C., Xu, T.: Attribute-conditioned layout gan for automatic graphic design. IEEE Transactions on Visualization and Computer Graphics 27(10), 4039–4048 (2020)
- 41. Li, J., Xu, Y., Cui, L., Wei, F.: Markuplm: Pre-training of text and markup language for visually rich document understanding. In: Annual Meeting of the Association for Computational Linguistics (2021)
- Li, X., Thickstun, J., Gulrajani, I., Liang, P.S., Hashimoto, T.B.: Diffusion-lm improves controllable text generation. Advances in Neural Information Processing Systems 35, 4328–4343 (2022)
- Li, X., Thickstun, J., Gulrajani, I., Liang, P.S., Hashimoto, T.B.: Diffusion-lm improves controllable text generation. Advances in Neural Information Processing Systems 35, 4328–4343 (2022)
- Liu, T.F.e.a.: Learning design semantics for mobile apps. In: UIST. pp. 569–579 (2018)
- 45. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (2017)
- 46. Lozhkov, A., Li, R., Allal, L.B., Cassano, F., Lamy-Poirier, J., Tazi, N., Tang, A., Pykhtar, D., Liu, J., Wei, Y., et al.: Starcoder 2 and the stack v2: The next generation. arXiv preprint arXiv:2402.19173 (2024)
- 47. Network, M.D.: Computed value css: Cascading style sheets (2023)
- Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: International Conference on Machine Learning. pp. 8162–8171. PMLR (2021)
- O'Donovan, P., Agarwala, A., Hertzmann, A.: Designscape: Design with interactive layout suggestions. In: Proceedings of the 33rd annual ACM conference on human factors in computing systems. pp. 1221–1224 (2015)
- van den Oord, A., Vinyals, O., kavukcuoglu, k.: Neural discrete representation learning. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 30. Curran Associates, Inc. (2017)
- 51. OpenAI: Gpt-4 technical report (2023)
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. Advances in Neural Information Processing Systems 35, 27730–27744 (2022)

- 18 Z. Shao et al.
- O'Donovan, P., Agarwala, A., Hertzmann, A.: Learning layouts for singlepagegraphic designs. IEEE transactions on visualization and computer graphics 20(8), 1200–1213 (2014)
- Qiu, Q., Otani, M., Iwazaki, Y.: An intelligent color recommendation tool for landing page design. In: 27th International Conference on Intelligent User Interfaces. pp. 26–29 (2022)
- 55. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical textconditional image generation with clip latents. arXiv preprint arXiv:2204.06125 1(2), 3 (2022)
- Razavi, A., van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with vq-vae-2. In: Neural Information Processing Systems (2019)
- 57. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
- Roziere, B., Gehring, J., Gloeckle, F., Sootla, S., Gat, I., Tan, X.E., Adi, Y., Liu, J., Remez, T., Rapin, J., et al.: Code llama: Open foundation models for code. arXiv preprint arXiv:2308.12950 (2023)
- Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image superresolution via iterative refinement. IEEE Transactions on Pattern Analysis and Machine Intelligence 45(4), 4713–4726 (2022)
- 60. Shao, Z., Gao, F., Qi, Z., Xing, H., Bu, J., Yu, Z., Zheng, Q., Liu, X.: Gem: Gestalt enhanced markup language model for web understanding via render tree. In: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. pp. 6132–6145 (2023)
- 61. Thorlacius, L.: The role of aesthetics in web design. Nordicom Review 28 (05 2007)
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 (2023)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
- Wang, P.: The influence of artificial intelligence on visual elements of web page design under machine vision. Computational Intelligence and Neuroscience 2022 (2022)
- 65. Wertheimer, M.: Gestalt theory. (1938)
- 66. Williams, R.: The non-designer's design book: Design and typographic principles for the visual novice. Pearson Education (2015)
- Xiang, P., Yang, X., Shi, Y.: Web page segmentation based on gestalt theory. In: 2007 IEEE International Conference on Multimedia and Expo. pp. 2253–2256 (2007)
- Xiao, Y., Wu, L., Guo, J., Li, J., Zhang, M., Qin, T., Liu, T.Y.: A survey on non-autoregressive generation for neural machine translation and beyond. IEEE Transactions on Pattern Analysis and Machine Intelligence 45, 11407–11427 (2022)
- Xie, C., Huang, W., Liang, J., Huang, C., Xiao, Y.: Webke: Knowledge extraction from semi-structured web with pre-trained markup language model. Proceedings of the 30th ACM International Conference on Information & Knowledge Management (2021)
- Yang, Z., Li, L., Lin, K., Wang, J., Lin, C.C., Liu, Z., Wang, L.: The dawn of lmms: Preliminary explorations with gpt-4v (ision). arXiv preprint arXiv:2309.17421 9 (2023)

Automatic Web Rendering Parameters Generation for Visual Presentation

- Zhao, N., Cao, Y., Lau, R.W.: Modeling fonts in context: Font prediction on web designs. In: Computer Graphics Forum. vol. 37, pp. 385–395. Wiley Online Library (2018)
- 72. Zhao, Z., Chen, L., Cao, R., Xu, H., Chen, X., Yu, K.: Tie: Topological information enhanced structural reading comprehension on web pages. In: Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. pp. 1808–1821 (2022)
- Zheng, X., Qiao, X., Cao, Y., Lau, R.W.: Content-aware generative modeling of graphic design layouts. ACM Transactions on Graphics (TOG) 38(4), 1–15 (2019)
- Zhou, M., Xu, C., Ma, Y., Ge, T., Jiang, Y., Xu, W.: Composition-aware graphic layout GAN for visual-textual presentation designs. In: Raedt, L.D. (ed.) Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022. pp. 4995–5001. ijcai.org (2022)
- Zhu, Y., Wu, Y., Olszewski, K., Ren, J., Tulyakov, S., Yan, Y.: Discrete contrastive diffusion for cross-modal and conditional generation. arXiv preprint arXiv:2206.07771 (2022)