

Attribute-Conditioned Layout GAN for Automatic Graphic Design

Jianan Li¹, Jimei Yang, Jianming Zhang¹, Chang Liu, Christina Wang¹, and Tingfa Xu

Abstract—Modeling layout is an important first step for graphic design. Recently, methods for generating graphic layouts have progressed, particularly with Generative Adversarial Networks (GANs). However, the problem of specifying the locations and sizes of design elements usually involves constraints with respect to element attributes, such as area, aspect ratio and reading-order. Automating attribute conditional graphic layouts remains a complex and unsolved problem. In this article, we introduce Attribute-conditioned Layout GAN to incorporate the attributes of design elements for graphic layout generation by forcing both the generator and the discriminator to meet attribute conditions. Due to the complexity of graphic designs, we further propose an element dropout method to make the discriminator look at partial lists of elements and learn their local patterns. In addition, we introduce various loss designs following different design principles for layout optimization. We demonstrate that the proposed method can synthesize graphic layouts conditioned on different element attributes. It can also adjust well-designed layouts to new sizes while retaining elements' original reading-orders. The effectiveness of our method is validated through a user study.

Index Terms—Generative adversarial networks, graphic design, attribute

1 INTRODUCTION

GRAPHIC design is very important in creating multimedia. It has a wide range of applications in areas including advertisement, books and websites, among many others. Good designs that are visually pleasing while clearly delivering information, highly rely on the expertise of experienced designers. Designers often need to create the same design in different sizes or retarget existing designs to new sizes to fit in various display resolutions, for example, from websites to mobile phone screens, which requires a great deal of human labor and intelligence. Automation tools for creating designs in desired sizes would assist designers and reduce design cycle time significantly.

The first step to render real designs is to generate graphic layouts by specifying locations and sizes of design elements. Conventional machine learning techniques, such as Generative Adversarial Networks (GANs) [1], [2] have been used to generate natural-looking digital images in pixel space [3], [4], [5], [6]. However, these methods are not suitable for synthesizing graphic designs which are not pixels but layouts of editable design elements. Recently, **LayoutGAN** [7] has been proposed to directly generate geometric parameters of graphic elements instead of image pixels from random variables. It demonstrates the feasibility of generating parameterized graphic

layouts with neural networks, but it is incapable of doing practical tasks such as automatic graphic design and retargeting when specific user data is provided. Such user data includes specific design elements and their corresponding content-based attributes. These attributes introduce additional spatial constraints to be obeyed when generating the layout. In this paper, we introduce three important attributes: 1) *Expected area* of elements with respect to the page size needs to be preserved. Distorted sizes will affect either the readability or the aesthetics. 2) *Aspect ratio* of specific elements are supposed to be retained, e.g., logos. 3) *Reading-order* or spatial distribution of elements should be preserved. For example, logos often appear on top, value propositions (headlines or images) in the middle and call-to-actions on the bottom for skyscraper banner ads. How to incorporate these attributes into automatic layout generation remains an open problem.

In this paper, we propose an attribute-conditioned layout GAN with user data, by incorporating elements' content-based attributes as conditions for automating graphic design. As shown in Fig. 1, we assume a set of graphic elements along with their content-based attributes are provided as input. We introduce attribute conditions to both the generator and the discriminator of our GAN model. Specifically, we do this by directly supplying the generator with attendant attributes of each element to generate attribute conditioned layouts. Instead of directly feeding attributes to the discriminator, we introduce an additional decoder branch to reconstruct the input attributes [8]. For training the proposed GAN model, we use the wireframe rendering layer similar to LayoutGAN [7] that rasterizes all parameterized elements into wireframe images. In addition to feeding all the elements to the discriminator to learn their global patterns, we also introduce a novel element dropout strategy which randomly removes some of the generated elements and feeds them to the discriminator. As there lacks global layout information delivered by a complete set of elements, the discriminator is forced to exploit local cues, such

• J. Li and T. Xu are with the Beijing Institute of Technology South Zhongguancun Street, Haidian District, Beijing 100081, China.
E-mail: {20090964, xutingfa}@bit.edu.cn.

• J. Yang, J. Zhang, C. Liu, and C. Wang are with the Adobe Inc., 345 Park Ave, San Jose, CA 95110 USA.
E-mail: {jimyng, jianmzha, cliu, cwang}@adobe.com.

Manuscript received 9 Aug. 2019; revised 24 May 2020; accepted 27 May 2020.

Date of publication 2 June 2020; date of current version 2 Sept. 2021.

(Corresponding author: Jianan Li.)

Recommended for acceptance by T. Schreck.

Digital Object Identifier no. 10.1109/TVCG.2020.2999335

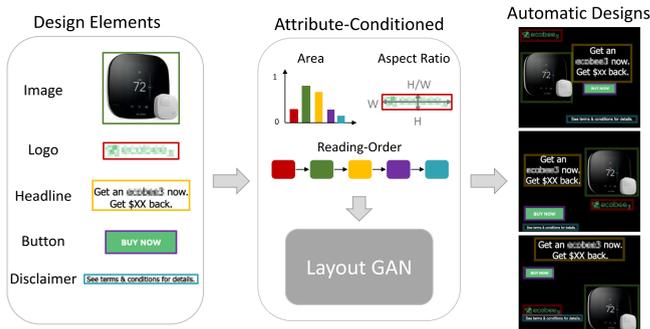


Fig. 1. We present a model for automating attribute-conditioned graphic designs by incorporating different attributes such as area, aspect ratio, and reading-order of given design elements.

as alignment and non-overlapping of the remaining elements to better learn local patterns. We further propose several hand-crafted loss functions following different design principles, such as non-overlapping, alignment and preservation of reading-orders to facilitate model optimization.

We apply our model to synthesize attribute conditional graphic layouts of different aspect ratios and retarget existing layouts to new sizes while preserving elements' reading-orders. For evaluation, we automatically generate a number of graphic designs, and compare the generated results with those produced by template-based methods, novice designers and professional designers. Experimental comparisons show that our automatic designs are better than those produced by template-based methods and novice designers, and sometimes comparable to professional designs.

To sum up, this paper comprises the following contributions: 1) A conditional generator that synthesizes attribute conditional structured data for layout design. 2) A conditional discriminator with element dropout that exploits and learns both global and local layout patterns. 3) Several hand-crafted loss functions based on various design principles for layout optimization.

2 RELATED WORK

Automatic layout generation is a major point of interest in the field of graphic design [9], [10], [11], [12], [13]. Harrington *et al.* [14] suggest an energy function to measure the aesthetics of a layout. Morcillo *et al.* [15] create graphic designs by maximizing design quality through cost functions to meet specific design constraints. Similarly, O'Donovan *et al.* [16] propose an energy-based model derived from design principles to create layouts, and further develop an interactive tool [17]. Pang *et al.* [18] present a web design interaction allowing designers to direct user attention using visual flow on web designs. Swearngin *et al.* [19] provide an interactive system that automatically infers a vector representation of screenshots and helps designers reuse or redraw design elements. Li *et al.* [7] apply generative adversarial networks to generate a set of relational graphic elements for layout design. Previous methods have used attribute information for design tasks with other graphic elements, such as fonts [20]. There are also works for 3D layout synthesis [21], [22], [23], [24], [25], [26].

Designs are now displayed on an increasingly large variety of devices and platforms, and hence, adjusting existing graphic layouts to new sizes has attracted much attention

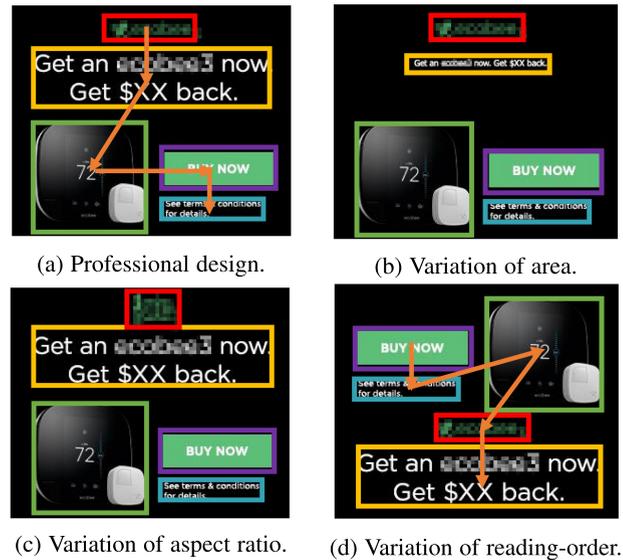


Fig. 2. Attributes to be preserved for design elements.

recently. Kumar *et al.* [27] transfer design and content between web pages by capturing structural relationships among elements and balancing local and global concerns through optimization. Baluja *et al.* [28] propose to retarget by segmenting a web page into coherent regions. Cho *et al.* [29] learn an attention map which leads to a pixel-wise mapping from source to target grid for image retargeting [30]. No previous work has applied generative adversarial networks to adjusting layouts to new sizes.

3 ATTRIBUTE-CONDITIONED LAYOUT GAN

3.1 Design Representation

Instead of generating pixel-level images as a layout representation, the model synthesizes a graphic layout comprised of a set of N design elements represented as $\{(p_1, \theta_1), \dots, (p_N, \theta_N)\}$, where p is a vector of the class label and $\theta = [\theta^1, \dots, \theta^m]$ denotes a set of m geometric parameters that can be in various forms for different graphic layouts [7]. We assume an element is a rectangle, or bounding box represented by its normalized center coordinates, width and height, i.e., $\theta = [x^C, y^C, w, h]$, which is very common in various graphic designs.

3.2 Attribute Description

A good professional design, as exemplified in Fig. 2a, is built on a layout in which all design elements are arranged in proper sizes and locations according to their content-based attributes such as area, aspect ratio and reading-order.

3.2.1 Expected Area

The size of a design element in the layout usually highly depends on its content. For example, for a text element, the area of the generated bounding box should match with the text length, so that the target text can be rendered into the generated region with proper font size to ensure readability and aesthetic soundness. Otherwise, as shown in Fig. 2b, generating a small bounding box (in yellow) for a large amount of text would make the rendered text too small, thus ruining a design.

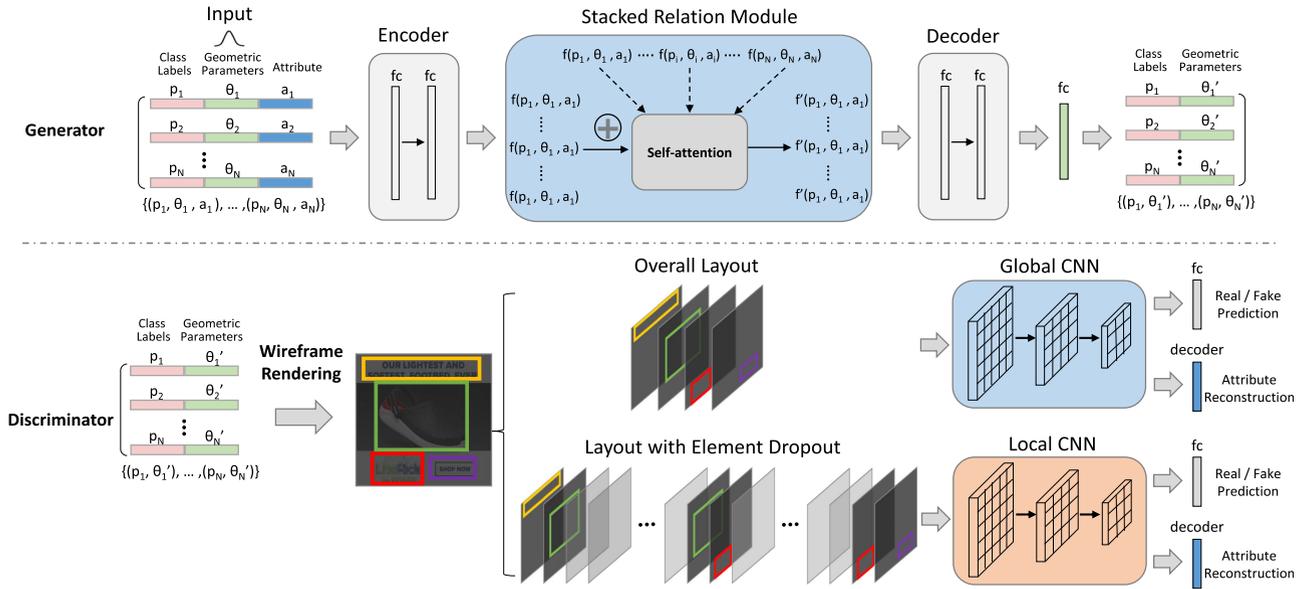


Fig. 3. Overall architecture of attribute-conditioned layout GAN. The generator takes graphic elements with class labels, random geometric parameters and attributes as input. An encoder followed by a stacked relation module embeds the input and contextually refines the embedded features for each element. The refined features are then decoded back to geometric parameters to assemble layouts. The discriminator takes generated structured data as input and renders all elements and a random partial list of them into two wireframe images, upon which two CNNs learn to capture both global and local behaviours of elements respectively. The discriminator is further tasked for attribute reconstruction to enhance discrimination.

The expected area, determined by elements' content, is thus an essential attribute to be considered for layout design.

3.2.2 Aspect Ratio

For an image element, the aspect ratio of the generated bounding box is supposed to be fixed. Given a generated bounding box of an arbitrary aspect ratio, current methods mostly fit the image into the generated region either via cropping or via warping. However, cropping the content could result in information loss and the warped content may result in unwanted geometric distortion, shown as the red bounding box in Fig. 2c. The aspect ratio of the image element is thus a necessary attribute to be preserved in the output layout.

3.2.3 Reading-Order

For a good design layout, information is presented smoothly based on the reading-order of design elements, as shown in Fig. 2a. In contrast, a layout with disordered elements would not convey information as clearly, as depicted in Fig. 2d. We use simple heuristics by following the left-to-right and top-to-down observation order commonly used in banner ads. In particular, we first compute the distance to the layout origin, i.e., the upper-left corner of the layout canvas, for element i as $d_i = \sqrt{x_i^L{}^2 + y_i^T{}^2}$, where (x_i^L, y_i^T) denotes the upper-left coordinates of the predicted bounding box. Then we sort the computed distance for all elements in ascending order, and take the sorted order of element i as its reading-order o_i . Such reading-order is supposed to be retained when adjusting an existing layout to new sizes. It should be taken as an additional attribute when needed.

3.3 Attribute Conditioning

We introduce element attributes in both the generator and the discriminator built upon LayoutGAN [7] for layout design.

3.3.1 Generator Architecture

As illustrated in Fig. 3, the generator takes as input a layout $z = \{(p_1, \theta_1, a_1), \dots, (p_N, \theta_N, a_N)\}$ consisting of N initial graphic elements, each of which is represented by a pre-defined class label p (one hot vector), randomly sampled geometric parameters θ , and expected element attributes a . The generator learns to capture and adjust the relations among all elements jointly through the relation module (implemented as self-attention in Wang *et al.* [31]). It does so by embedding the feature of each graphic element as a function of its relations with all the other elements in the design, and outputs p and contextually refined θ' to assemble graphic layouts $G(z) = \{(p_1, \theta'_1), \dots, (p_N, \theta'_N)\}$ similar as real ones x . Differently from LayoutGAN, we further incorporate an attribute vector a as input conditions to generate attribute conditional layouts. In our task, $a_i = [s_i, r_i, d_i]$ represents the expected area, aspect ratio and reading-order of element i , which determines the sizes and relative locations of the generated bounding box. Here, r_i equals the expected aspect ratio (bounding box height divided by width) if element i is supposed to be ratio-fixed, and 0 otherwise. Denote h'_i and w'_i as the height and width of the bounding box originally predicted by the generator, respectively. To strictly obey the aspect ratio constraint introduced by r_i , the final output bounding box height \hat{h}'_i can be computed with:

$$\hat{h}'_i = 1[r_i = 0]h'_i + r_i w'_i, \quad (1)$$

where the Iverson bracket indicator function $1[r_i = 0]$ equals 1 if $r_i = 0$, and 0 otherwise.

3.3.2 Discriminator Architecture

The discriminator renders synthesized structured data into wireframe images using a differentiable wireframe rendering layer [7], and learns to differentiate between synthesized and real layouts by capturing the relations among elements using

CNNs from the visual domain. To discriminate attribute conditional layouts, we further task the discriminator for attribute reconstruction by introducing an additional decoder (implemented as fully connected layers) as inspired by Odena *et al.* [8]. As element-wise information has been lost during wireframe rendering, we force the discriminator to reconstruct class-wise attribute information instead, i.e., the total area for each element type in the input layout. Specifically, the total area for element type c can be computed with: $S_c = \sum_{i=1}^N p_{i,c} s_i$, where $p_{i,c}$ is the class probability on element type c for element i . As a result, the discriminator outputs both a probability of the input belonging to real layouts, and an estimated area distribution over all element types.

3.4 Element-Dropout Wireframe Rendering

We adopt CNNs to better capture visual properties of an layout by rasterizing graphic elements to 2D wireframe images. To demonstrate this, we consider rasterizing a design comprised of N elements parameterized as $\{(p_1, \theta_1), \dots, (p_N, \theta_N)\}$ to a multi-channel output image I of dimension $W \times H \times M$, where each channel corresponds to one of M element types, and W and H denote the width and height of the layout canvas in pixels respectively. Each design element can be rendered into its own grayscale image $F_\theta(x, y)$ of dimension $W \times H$, as detailed in Li *et al.* [7]. The class activation vector for pixel (x, y) on I can be computed with:

$$I(x, y, c) = \max_{i \in [1 \dots N]} p_{i,c} F_{\theta_i}(x, y). \quad (2)$$

The wireframe discriminator mainly perceives global patterns of all elements. However, local behaviours such as alignment and non-overlapping are also important for a good design. To better capture those local rules, we further introduce element dropout strategy by randomly removing some input elements during the rendering process. Denote r as a N -dimension vector of independent Bernoulli random variables, each of which has a probability b of being 1. This vector is sampled and multiplied with the rendered grayscale image $F_\theta(x, y)$ to produce another rendered image I' :

$$r_i \sim \text{Bernoulli}(b), \quad (3)$$

$$I'(x, y, c) = \max_{i \in [1 \dots N]} r_i p_{i,c} F_{\theta_i}(x, y). \quad (4)$$

Fig. 3 depicts the architecture of the discriminator, which consists of both a global and a local branch. The global branch renders all input elements to wireframe images and uses CNNs to measure and optimize global patterns. Meanwhile, the local branch randomly removes some elements during rendering and adopts additional CNNs to focus on local patterns by looking at partial lists of elements. The discriminator is thus able to well capture both the global and the local layout patterns.

3.5 Discriminator Optimization

Denote D_{Θ_a} and $D_{\Theta'_a}$ as the global and the local branch with parameter Θ_a and Θ'_a respectively, and G_{Θ_g} as the generator with parameter Θ_g . The objective function of the discriminator consists of two terms. One is the log-likelihood for the input of being real layouts:

$$L_a = -\log D_{\Theta_a}(x) - \log(1 - D_{\Theta_a}(G_{\Theta_g}(z))) \\ - \log D_{\Theta'_a}(x) - \log(1 - D_{\Theta'_a}(G_{\Theta_g}(z))). \quad (5)$$

The other one is the L_1 distance between the predicted and the real area distribution over all element types:

$$L_r = \sum_{c=1}^M |S_c - S'_c|, \quad (6)$$

where S_c and S'_c represent the real and the predicted area for element type c respectively. The discriminator is trained to minimize $L_a + w_r L_r$, where w_r is a preset loss weight.

3.6 Generator Optimization

We introduce several losses based on different design principles for layout optimization.

3.6.1 Adversarial Loss

The quality of the layout depends, among other things, on overall composition on the page (e.g., use of positive and negative space, balance, etc). Manually designing losses to optimize such patterns is not trivial, and may lead a very complex energy function to optimize. We thus adopt adversarial training to learn target design patterns from real data automatically. Taking the generated layout as input, the discriminator predicts its probability of belonging to a real layout by assessing both the global and the local behaviours of design elements. Trying to fool the discriminator produces an adversarial loss L_{adv} , which encourages the generator to synthesize layouts similar as real ones:

$$L_{adv} = -\log D_{\Theta_a}(G_{\Theta_g}(z)) - \log D_{\Theta'_a}(G_{\Theta_g}(z)). \quad (7)$$

3.6.2 Margin Area Loss

We control the area of the generated bounding box for different elements by feeding their expected area attributes to the generator. We denote s_i and s'_i as the input expected area and the area of the predicted bounding box for element i respectively. s_i and s'_i are expected to be close in value, while remaining some degree of flexibility to better balance different design principles for layout optimization. We thus introduce a margin area loss:

$$L_{area} = \sum_{i=1}^N k\left(\frac{|s'_i - s_i|}{s_i} - \alpha\right), \quad (8)$$

where $k(x) = \max(0, x)$ (implemented as ReLU activation), and $\alpha \geq 0$ is a preset parameter to control the allowable relative area difference, which is set as 0.3 in our experiments.

3.6.3 Overlapping Loss

Well-designed layouts typically avoid overlapping elements. We offer an overlapping loss to penalize overlapping between any element pairs:

$$L_{over} = \sum_{i=1}^N \sum_{\forall j \neq i} \frac{s_i \cap s_j}{s_i}, \quad (9)$$

where $s_i \cap s_j$ denotes the overlapping area between element i and j .

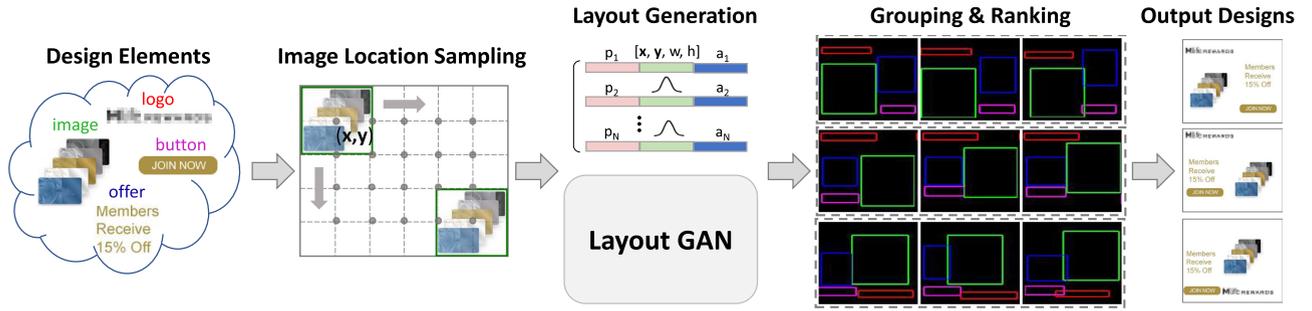


Fig. 4. Overall pipeline for automatic advertisement design. Given a set of design elements, the system first samples a number of locations for the product image, and then feeds the sampled locations and elements' attributes to layout GAN to generate a number of layouts, which are further grouped and ranked for final advertisement rendering and recommendation.

3.6.4 Alignment Loss

The alignment of design elements is crucial to how viewers perceive a layout. Adjacent elements are usually in six possible alignment types: Left, X-center, Right, Top, Y-center and Bottom aligned. Denote $\theta = (x^L, y^T, x^C, y^C, x^R, y^B)$ as the top-left, center and bottom-right coordinates of the predicted bounding box, we encourage pairwise alignment among elements by introducing an alignment loss:

$$L_{alg} = \sum_{i=1}^N \min \begin{pmatrix} g(\Delta x_i^L), g(\Delta x_i^C), g(\Delta x_i^R), \\ g(\Delta y_i^T), g(\Delta y_i^C), g(\Delta y_i^B) \end{pmatrix}, \quad (10)$$

where $g(x) = -\log(1 - x)$, and Δx_i^* ($*$ = L, C, R) is computed as:

$$\Delta x_i^* = \min_{\forall j \neq i} |x_i^* - x_j^*|. \quad (11)$$

Δy_i^* ($*$ = T, C, B) can be computed similarly.

3.6.5 Order Loss

The reading-order of elements is expected to be retained when adjusting an existing layout to new sizes. Following the left-to-right and top-to-down observation order commonly used in banner ads, a design element closer to the layout origin, i.e., the upper-left corner of the layout canvas, is expected to have a prior reading-order. As defined in Section 3.2.3, denote o_i as the pre-defined reading-order of element i and d_i as its distance to the layout origin. If element i has a prior reading-order compared to element j , i.e., $o_i < o_j$, the former is supposed to be closer to the layout origin than the latter accordingly, i.e., $d_i < d_j$. We thus introduce an order loss to penalize any element pairs whose relative distances to the layout origin are against their pre-defined reading-orders during layout adjustment:

$$L_{ord} = \sum_{i=1}^N \sum_{j=1}^N \mathbf{1}[o_i < o_j] k(d_i - d_j), \quad (12)$$

where $k(x) = \max(0, x)$. $\mathbf{1}[o_i < o_j]$ equals 1 when $o_i < o_j$, and 0 otherwise.

The final objective function is a weighted sum of all the losses above, whose weights are set as 0.6, 4.0, 8.0, 20.0 and 20.0 respectively in our experiments. We found the model is not very sensitive to different settings of these weights.

4 APPLICATIONS

A possible application of our model is in designing advertisements. Given a number of design elements of different types such as the logo, product image, headline, button, offer and disclaimer, the task of advertisement design is to specify the locations and sizes of different elements to form a well-organized layout. Fig. 4 illustrates our overall pipeline for automatic advertisement design, which consists of three basic steps: image location sampling, attribute-conditioned layout generation and layout grouping and ranking. We further illustrate how to adjust existing layouts to new sizes by incorporating additional reading-order constraints.

Image Location Sampling. Advertisement designs always use the product image as a starting point. Its location and size could greatly determine the layout style, and are thus supposed to be specified as design prior so that other elements can be arranged appropriately. To this end, the geometric parameters of the product image can be given as input conditions, while those for other element types are randomly sampled from a Gaussian distribution and to be refined accordingly to assemble a reasonable layout. Specifically, the width and height of the output bounding box for a given product image is computed from its expected area and aspect ratio. As for center coordinates, we apply grid sampling to obtain a number of candidate locations uniformly distributed inside the layout canvas, at which the product image can be accommodated.

Attribute-Conditioned Layout Generation. The layout arrangement of different design elements would highly depend on their content-based attributes. For text elements such as headlines, buttons, offers and disclaimers, the expected area determined by the text length is an essential condition for generating bounding boxes of proper scale so that text can be rendered with appropriate font size. For image elements such as logos, the aspect ratio is supposed to be preserved, and thus should be taken as an additional input attribute. Given the above content-based attributes, the model produces a number of layouts based on different sampled locations for the product image. We trained three models for generating layouts of different aspect ratios, i.e., portrait, square and landscape.

Grouping and Ranking. A concrete automatic design system should be able to assess the quality of different generated layouts for recommendation. Given a number of generated layouts, we first feed them into the trained discriminator and extract features from the last convolutional layer. We visualize

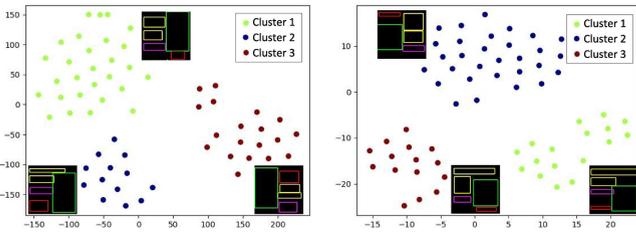


Fig. 5. Visualizing layout features using t-SNE. Given a set of design elements, the model generates a number of layouts with different patterns, which can be grouped into clusters based on the feature similarity. The representative layout for each clustered group is displayed nearby.

the features for different layouts using t-SNE [32]. As shown in Fig. 5, presence of clusters can be clearly found, which provides evidence that there exist similar patterns among different generated layouts. We thus divide all generated layouts into groups based on the feature similarity using k-means clustering, and assess the quality of each generated layout using the following cost function:

$$E = \omega_{adv}L_{adv} + \omega_{over}L_{over} + \omega_{alg}L_{alg}, \quad (13)$$

where ω_{adv} , ω_{over} and ω_{alg} denote the weights for the adversarial, overlapping and alignment loss terms mentioned in Section 3.6 respectively. These weights are set as the same values as used for model optimization. Finally, we rank layouts in each clustered group according to their computed cost values in descending order, and output the first-ranked layout in each group as the recommended result.

Graphic Layout Adjustment. As designs are usually presented in various displays, we sometimes need to adjust existing generated layouts to new sizes and aspect ratios while retaining elements' reading-orders. Though size free, the layout generation pipeline can only produce layouts of certain aspect ratios. We thus further train another generative model for layout adjustment by including reading-order as additional input attributes and by applying the order loss in Section 3.6.5 to retain such information in the adjusted layouts. In addition, we set the location of the

product image to be predictable without input prior to better meet the reading-order constraint. For inference, given a generated layout to be adjusted, we first extract elements' attributes and apply transformation to the aspect ratio of the logo and the product image according to the source and the target layout canvas size. Then, we feed the transformed attributes to the trained layout adjustment model to generate intermediate layouts, which are further resized to the target canvas size as final adjustment results.

5 EXPERIMENTS

The implementation is based on TensorFlow [33]. The network parameters are initialized from zero-mean Gaussian with standard deviation of 0.02. All network parameters are optimized using Adam [34] with a fixed learning rate of 0.00001. In our experiments, we focus on advertisement layouts with two to six bounding boxes that may belong to six most frequent element types appeared in real designs, as mentioned earlier. For design data, we collect totally around 17,000 advertisement layouts designed by professional designers, which are of three different aspect ratios, i.e., portrait, square and landscape. We randomly split these collected layouts into a training and a testing set and extract elements attributes in each layout as design conditions.

5.1 Ablation Studies

We investigate the effectiveness of different model components.

5.1.1 Element Dropout Strategy

We incorporate a local branch with element dropout into the discriminator to better capture local patterns of layouts. For validation, we compare the proposed model with the model variant without such local branch. Note that in this experiment, we do not apply the overlapping and the alignment loss when optimizing the above two models to purely compare discriminator behaviors. The first two rows in Fig. 6 provide some example layouts generated by the above

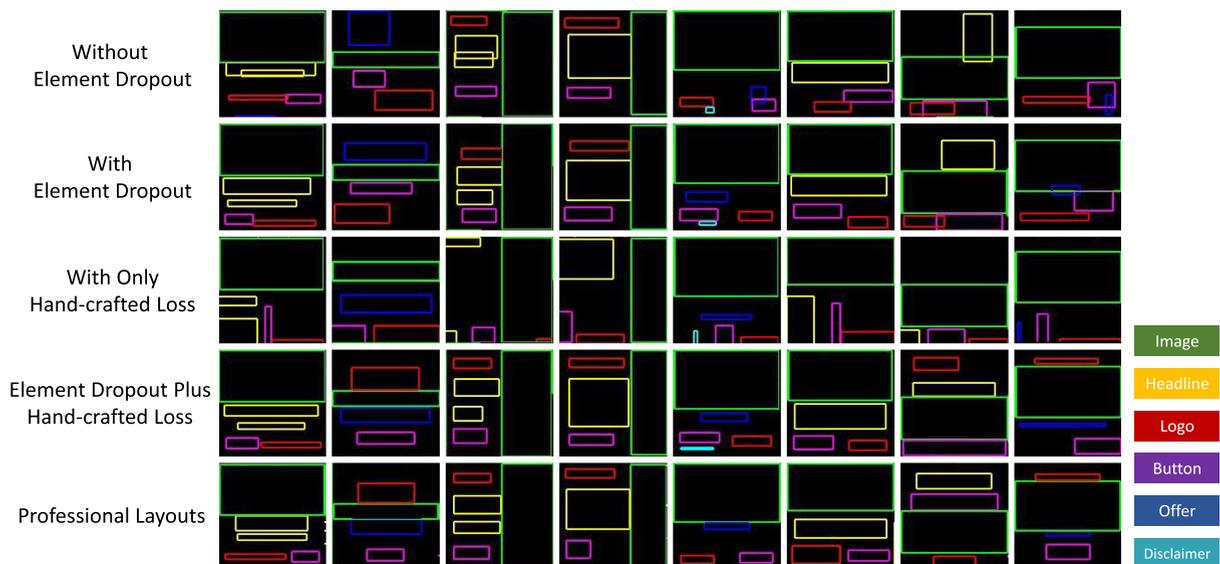


Fig. 6. Qualitative comparisons of generated layouts from different model variants.

TABLE 1
Spatial Analysis of Layouts From Different Models

Methods	Overlap	Alignment
Without element dropout	0.68	0.09
With element dropout	0.39	0.08
Element dropout + overlapping loss	0.09	0.08
Element dropout + alignment loss	0.36	0.04
Element dropout + both losses	0.08	0.04
Professional layouts	0.00	0.04

two models respectively. We also retrieve similar professional layouts that best match the input elements' attributes for reference, as shown in the last row. One can see that the additional local branch with element dropout helps better optimize elements' pairwise relations, thus alleviating overlapping and misalignment problems.

For quantitative evaluation of the generated layouts, we propose two metrics, overlapping and alignment index (the lower the better), as formulated in Equations (9) and (10) respectively. The first two rows in Table 1 provide quantitative comparisons of both models. We also include metric values for professional layouts in the last row for reference. They demonstrate that incorporating element dropout achieves lower values in both metrics, validating its superiority in optimizing local layout patterns. In addition, we compare results of adopting different dropout probabilities, referred to as b in Equation (3). Table 2 shows both metrics achieve the lowest when b equals 0.5. We use such settings in all experiments.

5.1.2 Hand-Crafted Loss

Considering the complexity of layout patterns, we further introduce hand-crafted losses, i.e., the overlapping and the alignment loss for optimization. For validation, we compare the results from model variants trained with only the adversarial loss, with only the hand-crafted loss and with both loss terms. As depicted in the second to fourth rows of Fig. 6, optimized layouts by a single adversarial loss suffer some degree of misalignment and overlapping. While using only the hand-crafted loss leads to degeneration. By combing both loss terms together, the model produces satisfactory results.

Similar conclusions can be drawn from quantitative comparisons. As shown in the middle four rows of Table 1, introducing hand-crafted loss terms incrementally based on adversarial training offers a consistent decrease in both metrics. To prove that introducing the hand-crafted loss would not affect overall layout balance, we further provide a supplementary metric to measure design symmetry in this experiment. Specifically, symmetric balance is a common design principle for organizing elements in conventional

TABLE 2
Comparisons of Using Different Dropout Probabilities

Probability	Overlap	Alignment
$b = 0.25$	0.13	0.05
$b = 0.50$	0.08	0.04
$b = 0.75$	0.09	0.04
$b = 1.00$	0.11	0.05

TABLE 3
Comparisons of Using Different Attribute Conditioning Methods

Conditioning	Overlap	Alignment
$w_r = 0.0$	0.10	0.05
$w_r = 0.5$	0.08	0.04
$w_r = 1.0$	0.12	0.04
Feeding	0.15	0.05

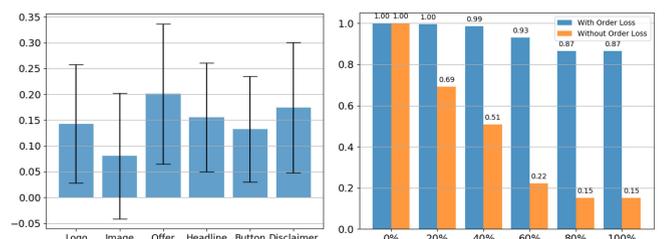
layouts. As the real layouts in our collected data mostly contain horizontal symmetry, we adopt the balance formulation inspired by O'Donovan *et al.* [16] to obtain symmetric scores, which compute the fraction of pixels having symmetric counterparts along the horizontal axis in the rendered layout. Incorporating the hand-crafted loss improves the symmetric score from 74.33 to 76.47 percent, which is closer to the score for collected real layouts, i.e., 76.51 percent. This suggests that the proposed hand-crafted loss would not affect global balance of layouts.

5.1.3 Attribute Reconstruction

We build the conditional discriminator by forcing it to reconstruct attributes of input layouts. Table 3 provides quantitative comparisons of using different loss weights for the attribute reconstruction task. All models with conditional discriminator ($w_r > 0$) achieve lower values in both metrics compared to the model variant without performing attribute reconstruction ($w_r = 0$). This suggests that the conditional discriminator effectively leverages attribute information through reconstruction, and thus improves discrimination for attribute conditional layouts. In addition, a proper loss weight, i.e., $w_r = 0.5$, contributes to performance boost. Moreover, the last row in Table 3 depicts that incorporating attributes through reconstruction outperforms the alternative manner, i.e., simply feeding attribute information as discriminator input [3], which further validates our design choice.

5.1.4 Attribute Preserving

To investigate the effectiveness of preserving input attributes, we analyze the area and the reading-order of output elements to compare them with input attributes based on the layout adjustment task. Specifically, we compute the relative difference between the generated bounding box area and the input expected area for each element, and analyze the mean along with the standard deviation of such area difference for each element type. As shown in Fig. 7a, the mean area



(a) Relative area difference. (b) Reading-order retention.

Fig. 7. Analysing the capability of preserving attributes.

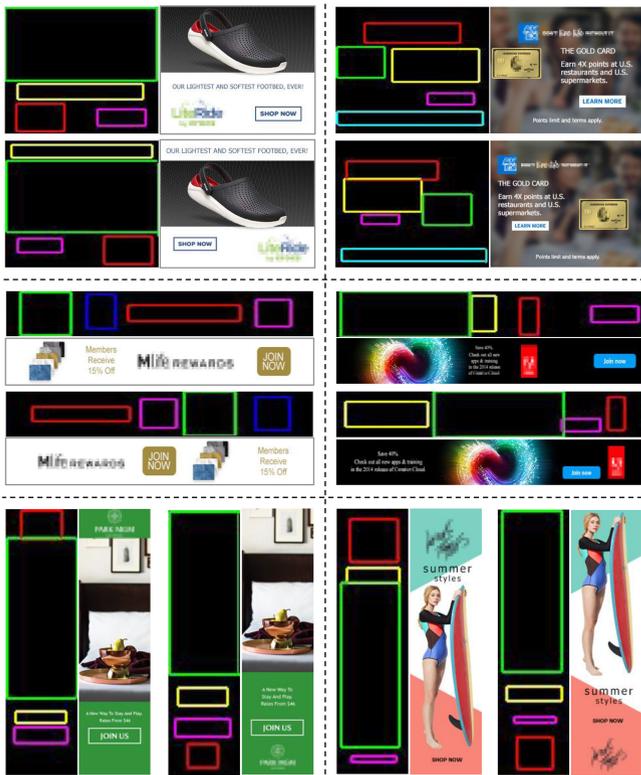


Fig. 8. Synthesized layout samples and corresponding rendered designs of different aspect ratios.

difference for all element types falls below the preset allowable threshold, i.e., $\alpha = 0.3$ as described in Section 3.6.2. To validate retention of the reading-order, we analyze the proportion of generated layouts whose percentage of design elements meeting the input reading-order constraint exceeds specific thresholds. Fig. 7b depicts that incorporating the reading-order loss effectively preserves the input reading-order and achieves higher proportion values consistently under different thresholds compared to the model variant without using such loss. The analysis of both area and reading-order retention demonstrates that our model can effectively incorporate and preserve input attributes to generate attribute conditional layouts.

5.2 Design Results

For qualitative results, we provide example generated layouts and corresponding rendered advertisements of three aspect ratios (portrait, square and landscape). As illustrated in Fig. 8, our method can well capture various layout patterns and generate aesthetically pleasing designs of different aspect ratios. In addition, we introduce a template-based baseline for attribute-conditioned layout retrieval. Specifically, given a set of design elements, we first construct a query vector similar as the generator input but comprising only elements' attributes. Then we compute the Cosine similarity between the query vector and the vectors derived from existing professional layouts in the training data, and retrieve the most similar professional layout directly as output. The first and the third column in Fig. 9 provide results from the template-based baseline and our method respectively. The retrieved layouts, though derived from professional designs, poorly meet

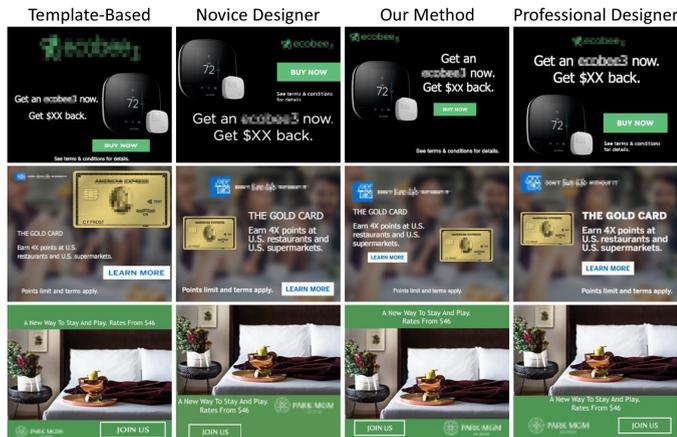


Fig. 9. Comparisons of designs from different methods.

elements' attribute conditions. As shown in the first two rows, mismatched aspect ratio for logos would ruin a design. Fig. 10 also provides the area difference analysis as introduced in Section 5.1.4. Our learning-based model can better meet the area constraint compared to the template-based baseline. We further compare our generated results with designs created by professional designers and those by novice designers who lack much design experience and complete designs by simply translating and scaling given elements. Fig. 9 demonstrates that our results are generally better than or at least as good as those produced by novice designers, and sometimes comparable to professional designs.

For quantitative evaluation, we conduct a user study to subjectively evaluate results from different methods mentioned above. The user study involves two user groups, i.e., 22 normal users with certain design background and 2 professional users with rich design experience (different from designers creating the professional designs to be evaluated). Specifically, we collect 20 sets of design elements and compare six design results for each set, which are produced by our method, the template-based baseline, professional designers and three novice designers, respectively. A subject is asked to rank the given six designs in descending order by following two criteria: 1) whether they are aesthetically displeasing, 2) whether they can clearly convey information. Fig. 11 depicts the histogram of rankings for designs from different methods and Table 4 provides computed average rankings (the lower the better). One can see that the ranking results from the normal and the professional users are consistent. Concretely, though not at the level of professional designs, our automatic

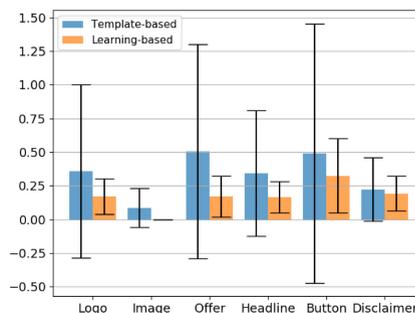


Fig. 10. Comparisons of relative area difference.

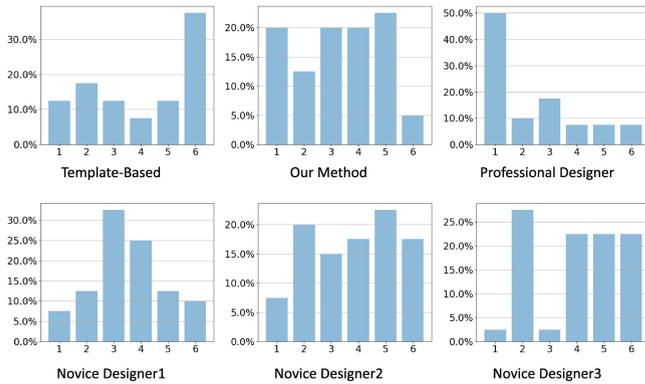


Fig. 11. Histograms of rankings for different methods.

TABLE 4
Average Rankings for Different Methods

Methods	Ranks (Normal)	Ranks (Professional)
Template-based	4.10	4.03
Novice Designer	3.58	3.78
Our Method	3.48	3.28
Professional Designer	2.66	2.35



Fig. 12. Comparisons of retargets from different methods.

designs achieve better rankings compared to designs produced by the template-based baseline and also those by novice designers. This suggests that our generated designs are generally preferred by participants.

We further adapt our method to adjusting existing layouts to new sizes while retaining elements' reading-orders. The first two rows in Fig. 12 provide retargets produced by our method and the model variant without considering the reading-order constraint respectively when given a specific design (middle in red) as input. We also compare with results from the energy-based model proposed by O'Donovan *et al.* [16], as shown in the third row. One can see that our method can properly specify the locations and sizes of different elements while retaining original reading-orders, thus generating more appealing and adequate retargets compared to the other two counterparts.

At last, we compare the average inference time of producing an automatic graphic layout by using different methods,

TABLE 5
Comparisons of Average Inference Time

Methods	Template-based	Energy-based	Our method
Average inference time (s)	0.32	~2400.00	0.69

i.e., the template-based baseline, the energy-based model [16] and our method. As shown in Table 5, our method is greatly faster or at least comparable in terms of speed compared to the alternatives, while creating better designs.

6 CONCLUSION

In this paper, we have proposed a novel attribute-conditioned layout GAN to synthesize graphic layouts by incorporating element attributes as conditions into both the generator and the discriminator. In addition, we introduce a novel element dropout strategy to enhance the capability of the discriminator in capturing local patterns, as well as several hand-crafted loss designs to facilitate layout optimization. We apply our model to generate attribute-conditioned layouts and adjust existing ones to new sizes. Experimental results well demonstrate the effectiveness of our method. In this work, we define the reading-order by simply following the left-to-right and top-to-down viewing order commonly used in banner ads. However, for designs containing many elements, defining the reading-order could be a complex problem. Future work include determining more content-adaptive reading-orders using theories such as visual attention and salience for automatic graphic design.

REFERENCES

- [1] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Advances Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [2] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [3] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," 2015, *arXiv:1511.06434*.
- [4] H. Zhang *et al.*, "Stackgan++: Realistic image synthesis with stacked generative adversarial networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 8, pp. 1947–1962, Aug. 2019.
- [5] A. Brock, J. Donahue, and K. Simonyan, "Large scale gan training for high fidelity natural image synthesis," 2018, *arXiv:1809.11096*.
- [6] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," 2018, *arXiv:1812.04948*.
- [7] J. Li, J. Yang, A. Hertzmann, J. Zhang, and T. Xu, "LayoutGAN: Generating graphic layouts with wireframe discriminators," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [8] A. Odena, C. Olah, and J. Shlens, "Conditional image synthesis with auxiliary classifier gans," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 2642–2651.
- [9] N. Hurst, W. Li, and K. Marriott, "Review of automatic document formatting," in *Proc. 9th ACM Symp. Document Eng.*, 2009, pp. 99–108.
- [10] Z. Bylinskii *et al.*, "Learning visual importance for graphic designs and data visualizations," in *Proc. 30th Annu. ACM Symp. User Interface Softw. Technol.*, 2017, pp. 57–69.
- [11] B. Deka *et al.*, "Rico: A mobile app dataset for building data-driven design applications," in *Proc. 30th Annu. ACM Symp. User Interface Softw. Technol.*, 2017, pp. 845–854.
- [12] B. Kovacs, P. Donovan, K. Bala, and A. Hertzmann, "Context-aware asset search for graphic design," *IEEE Trans. Vis. Comput. Graphics*, vol. 25, no. 7, pp. 2419–2429, Jul. 2019.

- [13] D. Ren, B. Lee, and M. Brehmer, "Charticulator: Interactive construction of bespoke chart layouts," *IEEE Trans. Vis. Comput. Graphics*, vol. 25, no. 1, pp. 789–799, Jan. 2019.
- [14] S. J. Harrington, J. F. Naveda, R. P. Jones, P. Roetling, and N. Thakkar, "Aesthetic measures for automated document layout," in *Proc. ACM Symp. Document Eng.*, 2004, pp. 109–111.
- [15] C. G. Morcillo, V. J. Martin, D. V. Fernandez, J. J. C. Sanchez, and J. A. Albusac, "Gaudi: An automated graphic design expert system," in *Proc. Nat. Conf. Artif. Intell.*, 2010.
- [16] P. O'Donovan, A. Agarwala, and A. Hertzmann, "Learning layouts for single-page graphic designs," *IEEE Trans. Vis. Comput. Graphics*, vol. 20, no. 8, pp. 1200–1213, Aug. 2014.
- [17] P. O'Donovan, A. Agarwala, and A. Hertzmann, "Designscape: Design with interactive layout suggestions," in *Proc. 33rd Annu. ACM Conf. Human Factors Comput. Syst.*, 2015, pp. 1221–1224.
- [18] X. Pang, Y. Cao, R. W. Lau, and A. B. Chan, "Directing user attention via visual flow on web designs," *ACM Trans. Graph.*, vol. 35, no. 6, 2016, Art. no. 240.
- [19] A. Swearngin, M. Dontcheva, W. Li, J. Brandt, M. Dixon, and A. J. Ko, "Rewire: Interface design assistance from examples," in *Proc. CHI Conf. Human Factors Comput. Syst.*, 2018, Art. no. 504.
- [20] P. O'Donovan, J. Libeks, A. Agarwala, and A. Hertzmann, "Exploratory font selection using crowdsourced attributes," *ACM Trans. Graph.*, vol. 33, no. 4, 2014, Art. no. 92.
- [21] S. Qi, Y. Zhu, S. Huang, C. Jiang, and S.-C. Zhu, "Human-centric indoor scene synthesis using stochastic grammar," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 5899–5908.
- [22] L.-F. Yu, S.-K. Yeung, C.-K. Tang, D. Terzopoulos, T. F. Chan, and S. J. Osher, "Make it home: Automatic optimization of furniture arrangement," *ACM Trans. Graph.*, vol. 30, no. 4, 2011, Art. no. 86.
- [23] P. Merrell, E. Schkufza, Z. Li, M. Agrawala, and V. Koltun, "Interactive furniture layout using interior design guidelines," *ACM Trans. Graph.*, vol. 30, no. 4, 2011, Art. no. 87.
- [24] M. Fisher, D. Ritchie, M. Savva, T. Funkhouser, and P. Hanrahan, "Example-based synthesis of 3D object arrangements," *ACM Trans. Graph.*, vol. 31, no. 6, 2012, Art. no. 135.
- [25] K. Wang, M. Savva, A. X. Chang, and D. Ritchie, "Deep convolutional priors for indoor scene synthesis," *ACM Trans. Graph.*, vol. 37, no. 4, 2018, Art. no. 70.
- [26] H. Fan, H. Su, and L. J. Guibas, "A point set generation network for 3D object reconstruction from a single image," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, Art. no. 6.
- [27] R. Kumar, J. O. Talton, S. Ahmad, and S. R. Klemmer, "Bricolage: Example-based retargeting for web design," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2011, pp. 2197–2206.
- [28] S. Baluja, "Browsing on small screens: Recasting web-page segmentation into an efficient machine learning framework," in *Proc. 15th Int. Conf. World Wide Web*, 2006, pp. 33–42.
- [29] D. Cho, J. Park, T.-H. Oh, Y.-W. Tai, and I. So Kweon, "Weakly-and self-supervised learning for content-aware deep image retargeting," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 4558–4567.
- [30] M. Rubinstein, D. Gutierrez, O. Sorkine, and A. Shamir, "A comparative study of image retargeting," *ACM Trans. Graph.*, vol. 29, no. 6, 2010, Art. no. 160.
- [31] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, Art. no. 4.
- [32] L. V. D. Maaten and G. Hinton, "Visualizing data using T-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, 2008.
- [33] M. Abadi *et al.*, "Tensorflow: A system for large-scale machine learning," in *Proc. 12th USENIX Conf. Operating Syst. Des. Implementation*, 2016, pp. 265–283.
- [34] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.



Jianan Li received the BS and PhD degrees from the School of Optoelectronics, Beijing Institute of Technology, Beijing, China. From July 2015 to July 2017, he worked as a joint training PhD student at National University of Singapore. From October 2017 to April 2018, he worked as an intern at Adobe Research. His research interests mainly include computer vision and real-time image/video processing.



Jimei Yang received the BS degree in electrical engineering and information science from China Agricultural University, in 2006, the MEng degree in pattern recognition and intelligent systems from the University of Science and Technology of China, in 2009, and the PhD degree in computer science from the University of California, Merced. He is currently a senior research scientist at Adobe. He was a visiting PhD student in Artificial Intelligence Lab at University of Michigan, Ann Arbor in 2015. From June 2007 to June 2009, he worked as a research assistant at Institute of Automation, Chinese Academy of Sciences.



Jianming Zhang received the BS and MS degrees in mathematics from Tsinghua University, in 2008 and 2011, respectively, and the PhD degree in computer science from Boston University, in 2016. He is currently a research scientist at Adobe. His research interests include computer vision and machine learning with a focus on deep learning, visual saliency and image editing.



Chang Liu received the BS degree from the University of Science and Technology of China, in 2014, and the MS degree in computer science from the University of California, Davis, in 2016. Currently he is a machine learning engineer in Adobe Inc, working on Photoshop/Lightroom family products. His primary research interests include primarily in computer vision and data analytics fields.



Christina Wang received the bachelors degree from the Beijing Institute of Technology, in 1992, and the masters degree in computer engineering from the University of Kansas, in 1996. She is currently a senior computer scientist from Adobe Inc. where she works in various areas of software development, including web application design and implementation, and transferring research projects from Adobe researchers to product teams where the research results are developed to commercial products or services.



Tingfa Xu received the PhD degree from the Changchun Institute of Optics, Fine Mechanics and Physics, Changchun, China, in 2004. He is currently a professor with the School of Optoelectronics, Beijing Institute of Technology, Beijing, China. His research interests mainly include computer vision and optoelectronic imaging.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.