

Staging E-Commerce Products for Online Advertising using Retrieval Assisted Image Generation

Yueh-Ning Ku^{**}
Snap Inc., USA

Shaunak Mishra[‡]
Amazon, USA

Mikhail Kuznetsov[†]
Amazon, USA

Paloma de Juan
Yahoo Research, USA

ABSTRACT

Online ads showing e-commerce products typically rely on the product images in a catalog sent to the advertising platform by an e-commerce platform. In the broader ads industry such ads are called dynamic product ads (DPA). It is common for DPA catalogs to be in the scale of millions (corresponding to the scale of products which can be bought from the e-commerce platform). However, not all product images in the catalog may be appealing when directly re-purposed as an ad image, and this may lead to lower click-through rates (CTRs). In particular, products just placed against a solid background may not be as enticing and realistic as a product staged in a natural environment. To address such shortcomings of DPA images at scale, we propose a generative adversarial network (GAN) based approach to generate staged backgrounds for un-staged product images. Generating the entire staged background is a challenging task susceptible to hallucinations. To get around this, we introduce a simpler approach called copy-paste staging using retrieval assisted GANs. In copy paste staging, we first retrieve (from the catalog) staged products similar to the un-staged input product, and then copy-paste the background of the retrieved product in the input image. A GAN based in-painting model is used to fill the holes left after this copy-paste operation. We show the efficacy of our copy-paste staging method via offline metrics, and human evaluation. In addition, we show how our staging approach can enable animations of moving products leading to a video ad from a product image.

ACM Reference Format:

Yueh-Ning Ku, Mikhail Kuznetsov, Shaunak Mishra, and Paloma de Juan. 2023. Staging E-Commerce Products for Online Advertising using Retrieval Assisted Image Generation. In *Proceedings of (AdKDD 2023)*. ACM, New York, NY, USA, 6 pages. <https://doi.org/XXXXX>

1 INTRODUCTION

The choice of image for an online ad can have a significant impact on the online user exposed to the ad. If the ad image is enticing enough, it can not only create brand awareness among online users but also

^{*}Work done while at Yahoo Research.

[†]Work done while at Yahoo Research.

[‡]Work done while at Yahoo Research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

AdKDD 2023, Long Beach, CA

© 2023 Association for Computing Machinery.

<https://doi.org/XXXXX>



Figure 1: Un-staged and staged version of a chair sold by an e-commerce vendor (staged version is likely to have better CTR).

drive them to click the ad and make subsequent purchases (conversions) [2, 20]. However, if the ad image is not properly designed to capture the user’s attention, it would lead to poor user interactions and adversely affect the advertising platform (by lowering revenue) and the advertiser (by lowering conversion rate). In this context, a common observation [10] is that ad images with products in a natural or real world setting (lifestyle images) tend to have better online performance. For example, an ad selling a chair is expected to perform better if the image shows a chair in a living room versus a chair against a solid (synthetic) background (as shown in Figure 1). However, such staging of products may be expensive and time consuming, specially when a vendor is selling multiple products at the same time. In DPA offerings from ad platforms (*e.g.*, Yahoo), the catalog images from an e-commerce vendor (*e.g.*, Walmart, Amazon) are typically used directly as ad images. As described later in our data analysis (based on data from an ad platform), a major fraction of such images are not staged, and hence there is a scope to enhance such images (*e.g.*, by generating a suitable background for the product).

Image generation has been an actively studied topic for the past few years. Diffusion models [13] and GANs [3, 6] have been powering the state-of-the-art results in this area. Image in-painting [11] is a slightly easier version of the problem where only parts of the image need to be generated as opposed to the whole image. To the best of our knowledge, we are the first to study GAN based image generation approaches for enhancing product images to serve as ad images. Our main contributions can be summarized as follows.

- (1) We study three tasks (as outlined below): (i) vanilla staging, (ii) copy-paste staging, and (iii) image-to-parallax animation.

- (2) In task 1, we aim to generate the entire background for a product. We use pix2pix [6] to train a GAN model with pairs of images (input: segmented out product image, output: staged product image with ground truth background).
- (3) In task 2, our goal is to retrieve a similar product image (with staging) and copy-paste the background while filling in gaps (holes) created in the process of swapping products. We leverage GAN-based in-painting to fill in the gaps mentioned above. We also introduce a weighted boundary loss for in-painting to focus on the image generation quality at product boundaries. Through Fréchet inception distance (FID) score, and human evaluations we show that copy-paste staging is significantly better than the vanilla staging baseline.
- (4) In task 3, we use GAN-based in-painting to create a sequence of images simulating the main product’s movement against the staged background as in a parallax animation. The foreground and background both move, but at different speeds, creating the illusion of depth. This is to show how our approach can lead to video ads from product images.

Our retrieval based approach (second task above) shares the intuition common in text generation: retrieval augmented generation (RAG) has better context understanding and generation quality. The remainder of this paper is organized as follows: related work in Section 2, problem formulation in Section 3, relevant data in Section 4, and our proposed approaches in Section 5. We go over our experimental results in Section 6, and end with a discussion in Section 7.

2 RELATED WORK

Online advertising: In online advertising, the ad creative (text and image) plays an important role in influencing online users towards brand awareness, clicks and purchases [8, 16, 20]. Studying ad images and text using state-of-the-art deep learning models in computer vision and natural language processing (NLP) is an emerging area of research. In [5], ad image content was studied using computer vision models, and their dataset had manual annotations for: ad category, reasons to buy products advertised in the ad, and expected user response given the ad. Using this dataset, [9, 21] used ranking models to recommend themes for ad creative design using a brand’s Wikipedia page. In [10], object tag recommendations for improving an ad image was studied using data from A/B tests. Although related to ads, the above methods are not applicable in our setup since none of them are image generation methods.

GANs for image generation. Generative Adversarial Networks (GANs) are a popular approach for image generation. While vanilla GANs [3] generate images from random noise, Conditional GANs [7] also allow to use extra information in a generative process. Recently developed pix2pix method [6] and its successors [18] use Conditional GANs for image-to-image generation optimizing GAN objective together with distance to a target image. While we focus on GANs in this paper, our copy-paste staging approach can be generalized to more recent diffusion models [13] (discussed in Section 7).

Saliency detection. Saliency detection in product images is needed to understand which parts of an image correspond to the main product being advertised (as opposed to the background). Salient object detection (SOD) aims to detect the most visually attractive objects

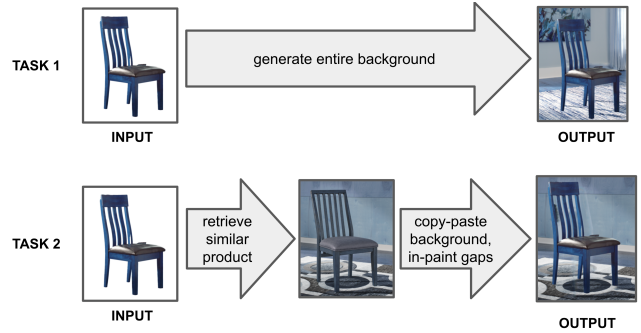


Figure 2: Tasks 1 and 2 explained with examples. In task 2, we need to retrieve a related product image with staging and then copy paste its background onto the input image, whereas in task 1 the goal is to generate the entire background.

with precise boundaries in images (*i.e.*, it returns a boundary map which can be used to segment out objects from the image). With the introduction of convolutional neural networks (CNNs) in computer vision, SOD accuracy has witnessed remarkable improvements. Recently, U²-Net [12] achieved state-of-the-art results for saliency detection by using a nested version of U-Net [14]-like architecture to capture richer local and global information. In our proposed approaches, we use saliency detection via U²-Net as a building block.

3 PROBLEM FORMULATION

We study three tasks in this paper as outlined below.

Task 1: vanilla staging. In task 1, the input is a product image (without staging), and the desired output is a product image with a model generated *relevant* background (stage). The model should generate the entire background as shown in Figure 2.

Task 2: retrieval assisted copy-paste staging. Task 2 is a simpler version of task 1. Here, we are given a pool of existing product images \mathcal{P} , and we need to retrieve a similar product image with staging such that we can copy-paste the staged background from the retrieved image onto the input image as shown in Figure 2. Image generation (in-painting) is needed to fill in the gaps after copy-pasting (since the input product and the product in the retrieved image are not identical, gaps will be created when we swap products).

Task 3: image-to-parallax animation. In this task, the goal is to take an input image (as shown in Figure 2 for task 2), and create an animation (*i.e.*, sequence of images), where the object in the input image (as in Figure 2) appears to be moving against a stationary but staged background. Such animations are expected to lead to higher user engagement [1].

4 DATA

For our experiments, we sampled data from Yahoo Gemini DPA (spanning November-December 2020). With an impression threshold of 10,000, we collected a sample of $\sim 18,899$ ads (each corresponding to a product), from which 11,927 had solid backgrounds, and the rest 6,972 had staged backgrounds. Each product in the sample

	cat	count
Bedroom > Bedroom Sets		117
Furniture > Bedroom > Beds > Queen		107
Furniture > Living Room > Entertainment Centers		100
Furniture > Living Room > Tables > End Tables		86
Furniture > Living Room > Tables > Coffee Tables		83
Furniture > Kitchen and Dining Room > Dining R...		78
Furniture > Kitchen and Dining Room > Dining R...		73
Furniture > Bedroom > Dressers > Dressers		71
Furniture > Bedroom > Dressers > Chests of Dra...		69

Figure 3: Subcategories for the “Furniture” category and counts.

belongs to a certain hierarchical category, e.g. "Kids > Kids Nightstands", "Furniture > Bedroom > Headboards > Queen", "Shoe > Heel > Oxford Heel". For our experiments, we considered only furniture images with staged backgrounds, which left us with 2071 images. Top-frequent subcategories with instance counts for the “Furniture” category (as in our sample) are shown in Figure 3.

5 PRODUCT STAGING USING GANS

In this section, we first explain the salient object detection model which is common to our proposed approaches for all the tasks outlined in Section 3. Next, we cover our proposed approaches for tasks 1, 2, and 3 (tasks explained in Section 3) in Sections 5.2, 5.3, and 5.4 respectively. Our major modeling contributions are in the approaches for tasks 2 and 3; for task 1, although we define the task, we use existing approaches (pix2pix [6]) to solve it. Task 2 is an easier version of task 1, and our proposed approach leads to more realistic staged product images compared to pix2pix for task 1.

5.1 Product segmentation via saliency

For tasks 1, 2 and 3, we use U²-Net [12] as our saliency object detector. The saliency object detector plays a crucial role in the first step of our approaches since it separates the main product(s) that will be replaced or copy-pasted versus the background in a product image. Once saliency probability maps are obtained from U²-Net, we set the threshold at 0.5 to generate binary masks and separate foreground pixels from the background pixels. For each task, we use product segmentation in a different manner as explained below.

5.2 Vanilla staging

For vanilla staging, we use the pix2pix method to generate an image background for a product which is segmented by a saliency mask. The algorithm is a conditional GAN optimizing the loss combining 1) regular GAN objective, and 2) ℓ_1 distance between original and restored images. We use product segmentation to prepare pairs of images to train pix2pix for stage (background generation). In particular, given an image with a staged product, we remove the background via product segmentation and use this as the input image to pix2pix. Figure 4 shows an example for this approach (original image in the middle, segmented product on the left, and the version with generated staging on the right).



Figure 4: Example of vanilla staging using pix2pix.

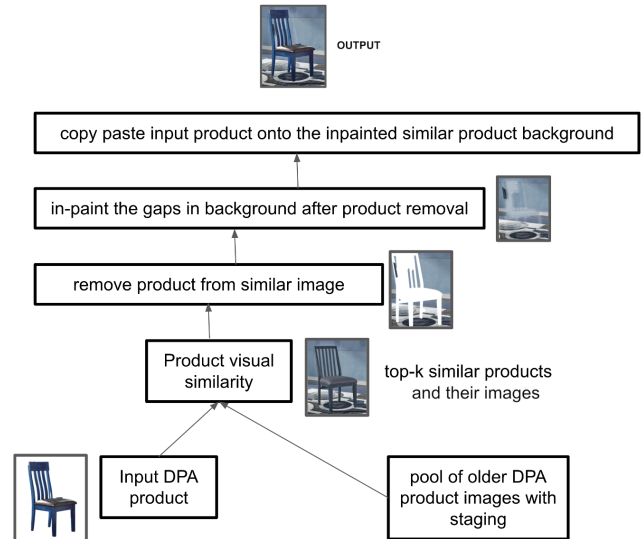


Figure 5: Proposed copy-paste staging algorithm.

5.3 Retrieval assisted copy-paste staging

For copy paste staging, we bypass the problem of generating the entire image background by using backgrounds from other relevant images. Our method consists of the following steps:

- (1) For a given segmented product (Figure 1, left), retrieve top- k similar products from a training collection (Fig. 6a). Similarity measure is a cosine distance between embeddings of corresponding product images provided by Inception-V3 [15].
- (2) For the top- k similar images, segment out the original products (Fig. 6b) and fill in the holes by inpainting using EdgeConnect [11] (GAN based model) and a new loss function that we introduce in Section 5.3.1.
- (3) Copy-paste the original product image to the inpainted top- k similar images, aligning shape and center mass for the corresponding product masks (Fig. 6c).

The above algorithm is illustrated with examples in Figure 5. We provide additional examples in Figure 6. After completing steps 1-3, we generate k product images with various backgrounds, only small parts of which (holes around the product before/after) are generated by GANs, which makes the images look more real if comparing with vanilla staging. For better background generation we introduce a new loss function as described below.

5.3.1 Weighted Boundary Loss. Recent works [11] and [17] explore coarse-to-fine inpainting approaches, since the structures

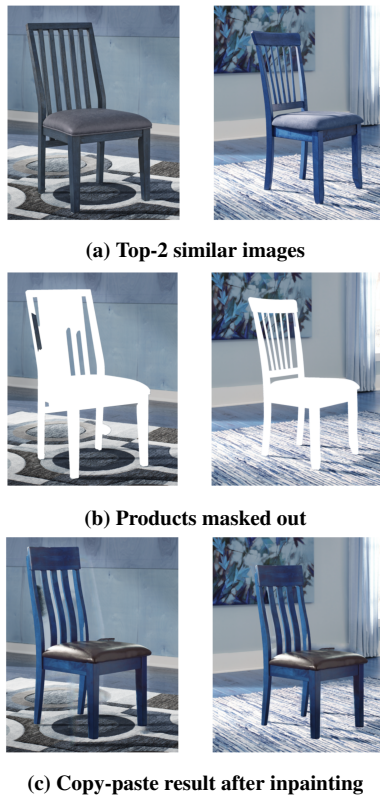


Figure 6: End-to-end result of copy-paste staging for top-2 similar images

of objects are complex and diverse, adding an intermediate step, like edge maps or monochromic images, can help models to learn progressively and eventually generate better final inpainted outputs. We propose a weighted boundary loss (WBL) to not only simplify the learning process (since the model needs to focus on lesser area), but also mimic the end application use case. Following prior work [11], our total generator loss consists of a conventional adversarial loss L_{ADV} and a feature-matching loss L_{FM} . In addition to these two losses, since our goal is to make the model learn better at the boundary of the masked area, we add weighted boundary loss L_{WBL} to amplify the loss penalty at the boundary area pixels. WBL is:

$$L_{WBL} = W_{map} * L_{\ell_1-norm}(E_{GT}, E_{pred}), \quad (1)$$

where E_{GT} is ground truth edge map of input images, E_{pred} is predicted edge map generated by the generator. The W_{map} is a pixel-wise weighted map and has the same size as input masked images and ground truth. To be more specific, the W_{map} has $\lambda_{boundary}$ for pixels around the boundary between masked area and unmasked area, and $\lambda_{non-boundary}$ for pixels away from the boundary, the pixel-wise L_{ℓ_1-norm} will multiply the corresponding λ as we calculate L_{WBL} . As Figure 7 illustrates, for each training sample, we create free-form dense masks by the method proposed by [19]. Then, we find the boundary area of the free-form mask and assign $\lambda_{boundary}$ (white area in Figure 7c) and $\lambda_{non-boundary}$ (gray area in Figure 7c). For experiments, we fixed $\lambda_{boundary} = 0.9$ and $\lambda_{non-boundary} = 0.1$.




Figure 7: Illustration of weighted boundary loss. 7a: Ground-truth training data. 7b: Free-form masked training data. 7c: Weighted mask.



Figure 8: Illustration of image parallax effect. 8a: original image 8b: image after shifting all the salient pixels. 8c: gap in-painted.

5.4 Image-to-animation

Parallax effect happens when the background pixels move slower than foreground objects in an animation, thereby creating an illusion of depth in a two-dimensional image. Generally, parallax effect requires independent foreground images and background images, and proper technique to make transparent backgrounds. In our proposed approach, by leveraging the power of salient object detection and in-painting, a parallax effect animation can be generated from a 2D image. Practically, we run salient object detection to define foreground pixels, then gradually move the foreground object around creating empty gaps between the current position of the object and original position. To fill the empty gap, we then use image in-painting model to in-paint those pixels and create serial realistic images. We illustrate the sample results of the above approach in Figure 8.

6 RESULTS

We first go over some sample results for tasks 1-3 followed by offline metrics (retrieval performance, generation quality) and human perceptual study results.


6.1 Sample results for tasks 1, 2 and 3

Task 1 (vanilla staging): sample results for task 1 (obtained via pix2pix) are shown in Figures 9 (b) and 10 (b). Compared to the original image with the background, the generated image has a lot of artifacts, and does not look so realistic. As we discuss below, the copy-paste staging results look more realistic.

Task 2 (copy-paste staging): Figures 9 (e) and 10 (e) show sample results for task 2 using the proposed copy-paste staging approach.



(a) Original image from DPA catalogs (left) and the input product image for staging (right).



(a) Original image from DPA catalogs (left) and the input product image for staging (right).



(b) Result of vanilla staging (pix2pix)



(b) Result of vanilla staging (pix2pix)



(c) Top-2 similar images



(c) Top-2 similar images



(d) Products masked out



(d) Products masked out



(e) Results of Copy-paste staging



(e) Results of Copy-paste staging

Figure 9: Example 1 for copy-paste staging.

Overall, the copy-paste staging results look much more realistic compared to pix2pix results.

Task 3 (parallax animation). A sample for the generated image-to-parallax animation: https://www.dropbox.com/s/9at5gz24ukhf2gi/product_staging_image_to_parallax_demo.mp4?dl=0.

Figure 10: Example 2 for copy-paste staging.

Table 1: Precision/Recall results for image retrieval

	@1	@3	@5
precision@k	0.468	0.409	0.374
recall@k	0.468	0.664	0.734

Table 2: FID scores for various staging and training options

	EdgeConnect (baseline)	EdgeConnect + WBL
vanilla staging	127.77	122.22
copy-paste staging	38.44	37.44

6.2 Offline metrics

Similar product image retrieval performance: a retrieved image is considered similar, if it belongs to the same subcategory as the input product. For example, if the input is a queen bed of subcategory "Furniture > Bedroom > Headboards > Queen", the retrieved product should fall into the same subcategory. Table 1 shows precision-recall results for similar products retrieval. For evaluation, we considered furniture subcategories (42 in total) with number of images ≥ 20 .

Generation quality: we measure the performance of our copy-paste staging results by evaluating Frechet inception distance (FID) [4]. FID is a popular metric for evaluating the quality of images created by GANs. The Wasserstein-2 distance in FID is calculated by comparing the features distribution of in-painted images with the distribution of real images, where the features are generated by a pre-trained InceptionV3 model. The comparison results are shown in Table 2. Since the copy-paste method in-paints only small regions of image around an object, it achieves much better FID score than vanilla staging. WBL further improves FID score in both methods.

6.3 Human evaluation

For a human perceptual study, we performed pairwise comparison tests, where experts were given two images at once, and the task was to determine which image appears more realistic (natural). For each comparison, we used three independent expert judgements, and decided the winner by majority voting. We performed three such tests: 1) 100 comparisons of vanilla staging images vs. ground truth images; 2) 100 copy-paste staging (with WBL) vs. ground truth; 3) 100 copy-paste (with WBL) vs. vanilla staging. The study showed:

- 0% of pix2pix images were better than ground truth;
- 3% of copy-paste images were better than ground truth;
- 76% of copy-paste images were better than pix2pix.

The above results clearly demonstrate the superiority of copy-paste staging and are in line with offline FID scores.

7 DISCUSSION

Our proposed approach provides low budget advertisers a way to stage products digitally without having to spend on the physical resources needed for staging. Staging a room can easily cost up to few hundred dollars for an advertiser, and with image generation methods like the ones we have proposed, this would be basically free of cost (except for the legalities around copying backgrounds

from other images). Leveraging the recent progress in prompt based image generation models, our approach can be further improved along the following lines: backgrounds from similar images could be used to generate prompts which then generate the background of the original product image. In addition, staged ads and parallax animations are expected to drive user engagement, and validating such hypothesis via an A/B test is one of our next steps.

REFERENCES

- [1] [n. d.]. New Study by Verizon Media, Magna, & IPG Media Lab Finds Interactive Ad Formats Engage Hard-to-Convince Audiences. <https://www.verizonmedia.com/press/2021/04/12/new-study-by-verizon-media-magna-ipg-media-lab>.
- [2] Narayan Bhamidipati, Ravi Kant, and Shaunak Mishra. 2017. A Large Scale Prediction Engine for App Install Clicks and Conversions. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (CIKM '17)*.
- [3] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).
- [4] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* 30 (2017).
- [5] Zaem Hussain, Mingda Zhang, Xiaozhong Zhang, Keren Ye, Christopher Thomas, Zuha Agha, Nathan Ong, and Adriana Kovashka. 2017. Automatic Understanding of Image and Video Advertisements. In *CVPR*.
- [6] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
- [7] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).
- [8] Shaunak Mishra, Mikhail Kuznetsov, Gaurav Srivastava, and Maxim Sviridenko. 2021. VisualTextRank: Unsupervised Graph-Based Content Extraction for Automating Ad Text to Image Search (*KDD '21*).
- [9] Shaunak Mishra, Manisha Verma, and Jelena Gligorijevic. 2019. Guiding Creative Design in Online Advertising. In *Proceedings of the 13th ACM Conference on Recommender Systems (RecSys '19)*.
- [10] Shaunak Mishra, Manisha Verma, Yichao Zhou, Kapil Thadani, and Wei Wang. 2020. Learning to Create Better Ads: Generation and Ranking Approaches for Ad Creative Refinement. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (CIKM '20)*.
- [11] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Qureshi, and Mehran Ebrahimi. 2019. EdgeConnect: Generative Image Inpainting with Adversarial Edge Learning. *arXiv preprint*.
- [12] Xuebin Qin, Zichen Zhang, Chenyang Huang, Masood Dehghan, Osmar R Zaiane, and Martin Jagersand. 2020. U2-Net: Going deeper with nested U-structure for salient object detection. *Pattern Recognition* 106 (2020), 107404.
- [13] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. [n. d.]. Zero-Shot Text-to-Image Generation (*ICML 2021*).
- [14] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.
- [15] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *IEEE CVPR*. 2818–2826.
- [16] Manisha Verma and Shaunak Mishra. 2022. Recommendation Systems for Ad Creation: A View from the Trenches (*RecSys '22*).
- [17] Tengfei Wang, Hao Ouyang, and Qifeng Chen. 2021. Image Inpainting with External-internal Learning and Monochromatic Bottleneck. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5120–5129.
- [18] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2018. High-resolution image synthesis and semantic manipulation with conditional gans. In *IEEE CVPR*. 8798–8807.
- [19] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. 2019. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4471–4480.
- [20] Yichao Zhou, Shaunak Mishra, Jelena Gligorijevic, Tarun Bhatia, and Narayan Bhamidipati. 2019. Understanding Consumer Journey Using Attention Based Recurrent Neural Networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19)*.
- [21] Yichao Zhou, Shaunak Mishra, Manisha Verma, Narayan Bhamidipati, and Wei Wang. 2020. Recommending Themes for Ad Creative Design via Visual-Linguistic Representations. In *Proceedings of The Web Conference 2020 (WWW '20)*.