

ICAR: Image-Based Complementary Auto Reasoning

Xijun Wang^{1,2}, Anqi Liang², Junbang Liang², Ming Lin^{1,2}, Yu Lou², Shan Yang²

¹ University of Maryland, College Park, USA

² Amazon, USA

{xijun, lin}@umd.edu, {lianganq, junbangl, ylou, ssyang}@amazon.com

Abstract

Scene-aware Complementary Item Retrieval (CIR) is a challenging task which requires to generate a set of compatible items across domains. Due to the subjectivity, it is difficult to set up a rigorous standard for both data collection and learning objectives. To address this challenging task, we propose a *visual compatibility* concept, composed of *similarity* (resembling in color, geometry, texture, and etc.) and *complementarity* (different items like table vs chair completing a group). Based on this notion, we propose a *compatibility learning* framework, a category-aware Flexible Bidirectional Transformer (FBT), for visual “scene-based set compatibility reasoning” with the *cross-domain* visual similarity input and auto-regressive complementary item generation. We introduce a “Flexible Bidirectional Transformer (FBT),” consisting of an encoder with flexible masking, a category prediction arm, and an auto-regressive visual embedding prediction arm. And the inputs for FBT are cross-domain visual similarity invariant embeddings, making this framework quite generalizable. Furthermore, our proposed FBT model learns the inter-object compatibility from a large set of scene images in a self-supervised way. Compared with the SOTA methods, this approach achieves up to **5.3%** and **9.6%** in FITB score and **22.3%** and **31.8%** SFID improvement on fashion and furniture, respectively.

Introduction

Online shopping catalogs provide great convenience, such as searching and comparing similar items. However, when customers can compare similar items, they often miss the browsing of complementary items in the e-shopping experience. Millions of online images offer a new opportunity to shop with inspirational home decoration ideas or outfit matching. But, to retrieve stylistically compatible products from these online images for set matching can be an overwhelming process. The ability to recommend visually complementary items becomes especially important, when shopping for home furniture and clothings.

In this work, we aim to address the *visual scene-aware Complementary Item Retrieval* (CIR) (Sarkar et al. 2022; Kang et al. 2019) task. In this task (as shown in Figure 2), we attempt to model human’s ability to select a set of objects from cross-domain pools, given a scene image, objects

in the scene, and object categories. Therefore, we propose a *visual compatibility* concept, consisting of two key elements: *similarity* and *complementarity*. Visual similarity and complementarity, however, can contradict each other sometimes. Items that look similar (color, geometry, texture, and etc.) may not be complementary (different items like dinner table vs sofa) when putting them into a set. Items that complement each other do not necessarily look similar (e.g. an outfit set in contrasting colors). The ambiguous definition for visual complementarity is a major challenge. This ambiguity makes it difficult to rigorously define an objective and creates extra challenge for collecting such datasets, when designing a data-driven method.

To address these issues, we first propose a “compatibility leaning” framework to model the visual *similarity* and *complementarity*. To the best of our knowledge, we are among the first to show qualitatively that our model based on this framework can generalize to unseen domains (where the model is not trained with, as shown in Figure 1 and Figure 6). For the scene-based CIR task, it’s complex to learn both the cross-domain similarity and complementarity. Therefore, we use cross-domain visual similarity invariant embeddings in our framework. Many previous CIR works (Han et al. 2017; Kang et al. 2019) also start from some types of learned embedding. But failing to model the visual similarity creates extra complexity for the complementary learning. Secondly, we propose to use self-supervised learning for visual complementarity reasoning by introducing an auto-regressive transformer based architecture. Given the difficulty to define style complementarity mathematically, we propose a solution based on the assumption that the items exist in the inspirational scene images are compatible with each other.

Built upon the aforementioned premises, we present a novel self-supervised transformer-based learning framework (overview shown in Fig. 3). Our model effectively learns both the similarity and complementarity between a set of items. Our model does not require extra complementary labels. In addition, compared to the prior work that models complementary items as *pairs* or a *sequence of items*, we model them as *unordered sets*. We carefully design our compatibility leaning model. First, we ensure that the learned embedding both contains and extracts all the necessary information for the compatibility learning. Second, we make

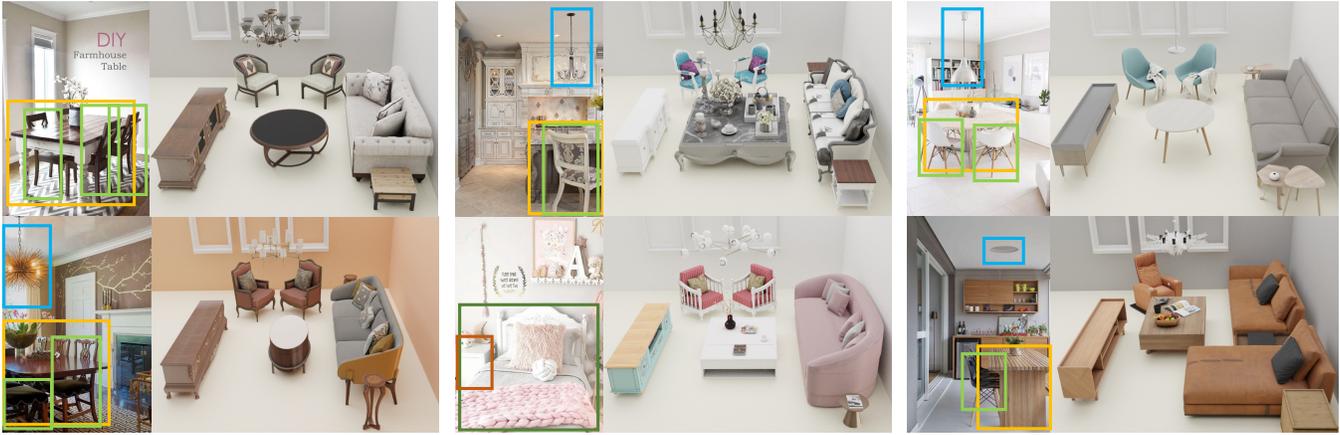


Figure 1: Cover Image: We present a new self-supervised model for scene-aware, visually compatible object retrieval tasks. In this example, given an inspirational home scene image (sampled from STL-home(Kang et al. 2019) column 1, 3, 5) with a pool of objects (3D-FRONT(Fu et al. 2021)) from an unseen domain, our model auto-regressively retrieves a set of stylistically compatible items (column 2, 4, 6).

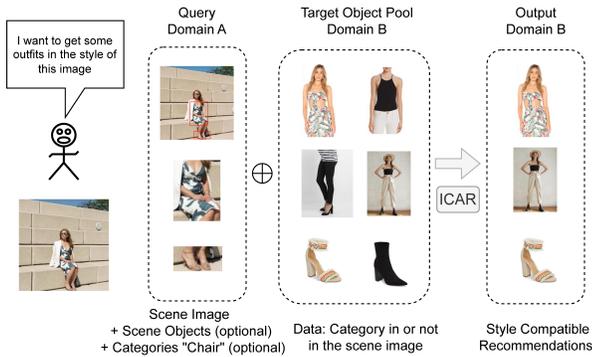


Figure 2: Scene-aware Complementary Item Retrieval Task Illustration. Given a query scene image, (optional) scene objects and item categories, the task goal is to generate a cross-domain set of stylistically compatible items.

full use of Transformer’s ability in reasoning about the interactions between these learned embeddings. To model flexible-length unordered set generation with cross-domain retrieval, we propose a new Flexible Bidirectional Transformer (FBT). In this FBT model, we model the unordered set generation using random shuffle and masking technique. In addition, we introduce a category prediction arm and a cross-domain retrieval arm to the transformer encoder. The added category prediction branch helps the model to reason about the complementary item types.

More importantly, we notice that most of the CIR prior work are evaluated via the Fill-In-The-Blank (FITB) (Han et al. 2017) metric or human in the loop. The FITB metric can reflect the model’s ability in cross-domain retrieval but it does not measure the complementarity as a set. Human in the loop evaluation, however, is both limited to scale and biases, if not conducted thoughtfully. To address these issues, we propose a new CIR evaluation metric: “Style Fréchet In-

ception Distance” (SFID) (see supplementary for details).

In summary, the key contributions of this work include:

- **Visual compatibility** is defined based on *similarity* and *complementarity* for the Scene-aware Complementary Item Retrieval (CIR) task and a new *compatibility learning framework* is designed to solve this task.
- For the **compatibility learning framework**, a category-aware *Flexible Bidirectional Transformer (FBT)* is introduced for visual scene-based set compatibility reasoning with the *cross-domain* visual similarity input and auto-regressive complementary item generation.

Related Work

Visual Similarity Learning Visual similarity learning has been a main computer vision topic. The goal is to mimic human’s ability in finding visually similar objects or scenes. This is particularly studied in image retrieval (El-Nouby et al. 2021; Radenović, Tolias, and Chum 2018; Teh, DeVries, and Taylor 2020; Cheng et al. 2021) - finding images with a certain definition of similarity and so on. Prior to retrieving similar clothing, researchers also studied how to detect and segment clothing from real-life images (Yamaguchi et al. 2012; Yang and Yu 2011; Gallagher and Chen 2008). With clothing detection or segmentation, similar clothing retrieval is explored via style analysis (Hsiao and Grauman 2017; Kiapour et al. 2014; Simo-Serra and Ishikawa 2016; Yu et al. 2012; Yamaguchi, Hadi Kiapour, and Berg 2013; Kiapour et al. 2014; Di et al. 2013). Recent fashion retrieval tasks can be further categorized based on the input information, such as images (Liu et al. 2016; Simo-Serra and Ishikawa 2016; Zhai et al. 2017; Tran et al. 2019), clothing attributes (Ak et al. 2018), videos (Cheng et al. 2017).

Visual complementarity Learning Visual complementarity learning, unlike visual similarity learning, is much more ambiguous. There are a couple of research directions: pairwise complementary item set complementary prediction (no

cross domain retrieval) (Tangseng, Yamaguchi, and Okatani 2017; Li et al. 2017; Han et al. 2017; Hsiao and Grauman 2018; Tangsen, Yamaguchi, and Okatani 2017; Shih et al. 2018; Li et al. 2020), set complementary item retrieval (Hu, Yi, and Davis 2015; Huang et al. 2015; Liu et al. 2012), personalized set complementary item prediction (requires user input) (Taraviya et al. 2021; Chen et al. 2019; Li et al. 2020; Su et al. 2021; Zheng et al. 2021; Guan et al. 2022b,a) and multi-modal complementary item prediction (Guan et al. 2021). All these prior work focus on feature representation learning. Another line of works (Chen et al. 2015; Song et al. 2017b; Tan et al. 2019; Lin, Tran, and Davis 2020) focus on learning multiple sub-embedding based on different properties for both similarity and compatibility. Later on, Sarkar (Sarkar et al. 2022) uses Transformer and CNN-based image classification features for compatibility learning. Unlike all the work above, we build our visual compatibility model which focuses on both similarity and complementarity.

Learning Framework Many researchers have studied and explored building a cascaded learning framework. The cascaded method here means learning how to encode the data then modeling the statistics of this encoding. Many of the methods proposed for CIR task can also be categorized as two-stage models. But almost all of them use the image classification training target as the first-stage feature extractor (Han et al. 2017; Sarkar et al. 2022; Kang et al. 2019; Chen et al. 2019). Taraviya (Taraviya et al. 2021) propose a two-stage model for personalized pairwise complementary item recommendation where they learn a feature embedding specially designed for customer preferences in their first stage. In our compatibility learning, we set cross-domain visual similarity embedding as input, and design FTB for complementary set generation. We show empirically that the visual similarity feature, compared to image classification learned features, are better suited for CIR. Our design makes our model surpass the prior work in scene aware CIR task.

Method

Problem Statement Given a scene image \mathcal{I} , a set of unordered objects $\mathcal{O} = \{o_i\}_{i=0}^N, o_i \in \mathcal{D}_A$ in the scene and a set of unordered object categories $\mathcal{C} = \{c_i\}_{i=0}^L$, the problem is to retrieve cross-domain a set of complementary objects $\mathcal{X} = \{x_i\}_{i=0}^L, x_i \in \mathcal{D}_B$. This generated set of objects needs to be both visually compatible with each other and of visually similar style to the input scene image \mathcal{I} . Here we use \mathcal{D}_A and \mathcal{D}_B to denominate the two different visual domains, L to represent the number of objects to retrieve during inference and N is the number of scene objects. The difference between the two domains \mathcal{D}_A and \mathcal{D}_B can be quantified as the Fréchet distance \mathcal{F} larger than a certain threshold θ .

Conditional Compatibility Auto Reasoning

We formulate the problem of generating a set of objects $\mathcal{X} = \{x_i\}_{i=0}^L$ conditioned on the scene image \mathcal{I} and a specified set of categories as how to compute likelihood (Eq. 1) of creating the object set \mathcal{X} given the scene image \mathcal{I} , objects in the scene \mathcal{O} , and set of categories \mathcal{C} . We model the probability of generating the unordered set \mathcal{X} as the sum of

generating the set in any permutation $\hat{\mathcal{X}}$:

$$p(\mathcal{X}_i|\mathcal{I}, \mathcal{O}, \mathcal{C}) = \sum_{\hat{\mathcal{X}} \in \Phi(\mathcal{X}_i)} p(x_i|x_0, \dots, x_{i-1}, \mathcal{I}, \mathcal{O}, \mathcal{C}), i \leq L \quad (1)$$

where $\Phi(\mathcal{X})$ includes all the permutation of the target object set \mathcal{X} given all the permutation of the categories \mathcal{C} , and L is the maximum number of items to compose a set. For each permutation of \mathcal{X} , the set generation becomes a sequence generation problem. We model the sequence generation as an auto-regressive process. In the auto-regressive process, the next item in the set is generated conditioned on the prior items. This auto-regressive process statistically formulated as the multiplication of the probabilities:

$$p(x_i|x_0, \dots, x_{i-1}) = \prod_j^{j < (i-1)} p(x_j|x_0, \dots, x_{j-1}). \quad (2)$$

To learn to conditionally generate the best set of objects, our model learns to maximize the log likelihood of the probability, $p(\mathcal{X}|\mathcal{I}, \mathcal{O}, \mathcal{C})$,

$$\log p(\mathcal{X}_i|\mathcal{I}, \mathcal{O}, \mathcal{C}) = \sum_{\hat{\mathcal{X}} \in \Phi(\mathcal{X}_i)} \sum_j^{j < i} \log p(x_i|x_{<j}, \mathcal{I}, \mathcal{O}, \mathcal{C}) \quad (3)$$

To approximate the log likelihood Eq. 3, we propose a two-stage learning framework.

Compatibility Learning Framework

Visual Similarity Learning To relieve the complexity of learning both the visual complementarity and similarity directly from pixel domain, we propose to separate them into two stages. In the first stage, our model focuses on the visual similarity learning. As shown in Figure 4, we apply a CNN-based (ResNet50) visual similarity model (Jun et al. 2019) with normalized softmax loss (Zhai and Wu 2018) and soft-margin triplet loss (Hermans, Beyer, and Leibe 2017) (refer to Supplementary for more details). With this model, we project the scene image \mathcal{I} , objects in the image \mathcal{O} , and the item images in retrieval pool \mathcal{X} onto this embedding:

$$\{\mathbf{I}, \mathbf{O}, \mathbf{X}\} = g(\{\mathcal{I}, \mathcal{O}, \mathcal{X}\}), \mathcal{I} \in \mathbb{R}^3, \mathcal{O} \in \mathbb{R}^3, \mathcal{X} \in \mathbb{R}^3 \quad (4)$$

where g is our visual similarity model. This projection helps our second stage model to converge faster, similar in spirit to sequential optimization. We also show empirically (refer to Sec. Similarity Learning Results for details) that the visual similarity embedding is best suited for learning cross-domain visual compatibility.

Complementarity Reasoning with Flexible Bidirectional Transformer At the second stage, we propose a new Flexible Bidirectional Transformer (FBT) (see Figure 5 for *conditional cross-domain unordered set generation*). We choose Transformer model (Vaswani et al. 2017) as the core architecture to learn the inter-object compatibility. The vanilla Transformer model (Vaswani et al. 2017), as it is originally proposed for modeling ordered sequence structured data, such as languages and images, is insufficient for our task.

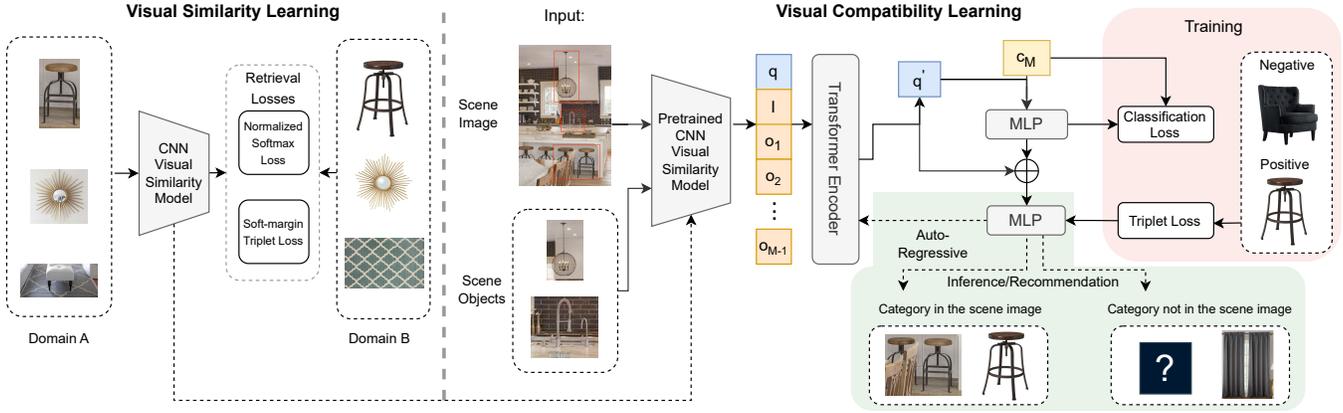


Figure 3: ICAR Model Overview. In similarity learning, we apply a CNN-based model (Jun et al. 2019) to learn the visual similarity features across two domains. The learned features are required for both complementary reasoning in the complementarity learning and the cross-domain retrieval. With the learned features, in the complementarity learning, we propose a Flexible Bidirectional Transformer (FBT) model to learn the multi-object visual compatibility.

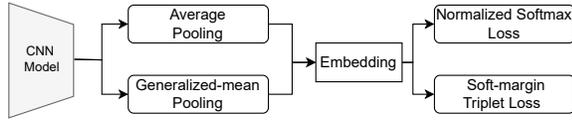


Figure 4: VSIM: Visual Similarity Model.

We introduce: (1) random shuffling together with random length sequence masking for set generation; (2) category prediction arm to better model the category distribution for a set of objects; and (3) visual embedding prediction arm for visual compatibility modeling. During inference, our FBT model generates an unordered set auto-regressively (Eq. 2

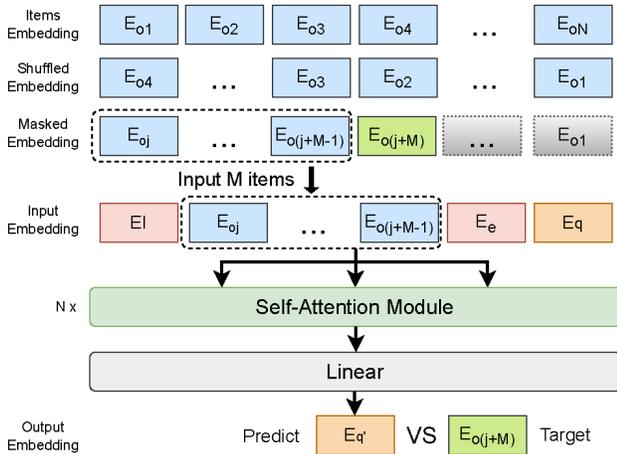


Figure 5: FBT: Flexible Bidirectional Transformer. We randomly sample $M \in [0, N]$ items from the total item number N of items (in a scene) as input set, and the $(M + 1)_{th}$ item not in the input set as output target.

and shown in the green part of Figure 3). Inspired from the CLS token proposed in the Vision Transformer (ViT) (Dosovitskiy et al. 2020), we also use a trainable variable denoted as \mathbf{q} to extract inter-token relation,

$$\begin{aligned} \mathbf{q}' &= e(\mathbf{E}\mathbf{I}, \Phi(\mathbf{E}\mathbf{O}); \mathbf{E}\mathbf{q}) \\ &= \text{MLP}(\text{MSA}(\mathbf{E}\mathbf{I}, \Phi(\mathbf{E}\mathbf{O}), \mathbf{E}\mathbf{e}, \mathbf{E}\mathbf{q})) \end{aligned} \quad (5)$$

$$\mathbf{E}\mathbf{O} = [\mathbf{E}\mathbf{o}_1, \mathbf{E}\mathbf{o}_2, \dots, \mathbf{E}\mathbf{o}_M, \text{MASK}],$$

where \mathbf{q}' denotes the corresponding output of the trainable input token \mathbf{q} , e is the end token, Φ is the masking operation, $e(\cdot)$ represents the Transformer encoder with MSAs (the Multi-headed Self-Attention layers), MLPs (Multi-Layer Perception), \mathbf{E} is the linear projection and M is the unmasked sequence length. The output \mathbf{q}' is then be used for predicting both the category c_{M+1} (Eq. 6) and the visual embedding of the next item \mathbf{x}_{M+1} (Eq. 7).

$$\hat{c} = \text{MLP}(\mathbf{q}') \quad (6)$$

$$\hat{\mathbf{x}}_{M+1} = \text{MLP}[\mathbf{q}', \hat{c}] \quad (7)$$

The output category embedding \hat{c} is supervised via the Cross-Entropy loss. And the visual feature embedding $\hat{\mathbf{x}}$ is supervised using a triplet loss (Yang et al. 2019). To form a triplet, the anchor is the predicted embedding $\hat{\mathbf{x}}_{M+1}$ with the target item's embedding \mathbf{x}_{M+1} as the positive and randomly selected same category object's embedding as the negative. One challenge in the feature learning is the space collapsing, where points in the embedding space get too close. This space collapse can lower the representation capacity. To avoid such an issue, we apply a differential entropy regularizer (Girdhar et al. 2019) to maximize the distance between each point and its closest neighbors in the embedding space. The regularizer is defined as follows:

$$\text{Reg} = -\frac{1}{N} \sum_i \log(D_{\min_{i \neq j}}(z_i, z_j)), \quad (8)$$

where we divide the L2 distance between sample i, j by 4 as $D_{i,j}$ to make the distance between 0 and 1. z is the input feature maps.

Experiments

Setup

Benchmark Datasets: In the following experiments, we evaluate our proposed ICAR using four datasets. DeepRooms (Gadde, Feng, and Martinez 2021) is a large-scale (1.4 million), high-quality human annotated object detection dataset covering a total of 81 fine-grained furniture and home product categories with 210K room-scene images. STL-Home (Kang et al. 2019) includes 24,022 interior home design images and 41,306 home decor items, which can generate 93,274 scene-product pairs. STL-Fashion (Kang et al. 2019) contains 72,189 fashion-product pairs from its 47,739 fashion images and 38,111 product images. And Exact Street2Shop (Hadi Kiapour et al. 2015) provides fashion-product 10,608 pairs from its 10,482 fashion images and 5,238 product images with the bounding box of products in scene images.

Implementation: All the models in our experiments are trained using AdamW (Loshchilov and Hutter 2017) with cosine learning scheduler from an initial learning rate of $2e - 4$ to 0. Models are trained for 500 epochs with a batch size of 256. We choose 1 negative sample in the triplet loss. And we use 1.0, 1.0, and 0.05 as the weights for cross-entropy loss, triplet loss, and regularizer loss respectively.

Evaluation Metrics

To evaluate our proposed model’s performance for scene-aware cross-domain CIR, we benchmark our method using previously proposed Fill-In-The-Blank (FITB) (Han et al. 2017) and our newly proposed SFID metrics. FITB accuracy is measured via counting percentage of times the model correctly picked the ground-truth object from the candidate pool. Following (Kang et al. 2019), we set the number of candidate to 2 for STL and Street2Shop and 3 for DeepRooms. For fair comparison, we apply the same method in setting up the FITB candidate pool across all the experiments. While FITB estimates the model’s ability to retrieve the best object, it fails to capture the compatibility among a set of items. To compensate, we propose a new distribution distance-based metric: **Style FID (SFID)**. It is difficult to define visual style, as it covers various aspects, including color, texture, shape and so on. Instead, we use the *visual similarity distribution* between the generated set and the designed sets of items as the measurement of stylistic compatibility. Similar to the FID (Heusel et al. 2017), we apply the Fréchet distance \mathcal{F} for distribution distance measurement. Instead of directly estimating the pixel value distribution, we apply a feature extractor that can project the pixel values onto an embedding. With the feature extractor, the *computed distribution* can focus on style related features, including colors, edges, textures, patterns, and shapes. We define

$$\text{SFID Score} = \mathcal{F}(f(\mathcal{X}), f(\mathcal{Y})), \quad (9)$$

$\mathcal{F}(\mathbf{X}, \mathbf{Y}) = |\mu_X - \mu_Y| + \text{tr}(\sigma_X + \sigma_Y - 2\sqrt{\sigma_X * \sigma_Y})$, where \mathcal{X} is the generated set of objects, \mathcal{Y} is a set of well designed (ground-truth) objects, function f is a feature extractor, μ and σ are the mean and variance. Please refer supplementary for more details about SFID.

Method	FITB \uparrow	Condition
CSA-Net w. TL	60.0	/
CSA-Net w. OL	57.4	/
Visual Similarity Learning	62.0	/
OutfitTransformer	77.6	/
ICAR (Ours)	83.9	Predict category
ICAR (Ours)	87.1	Given category

Table 1: FITB Results on DeepRooms (Gadde, Feng, and Martinez 2021). Our approach improved FITB accuracy by 9.5% over Visual Similarity Learning, CSA-Net (Lin, Tran, and Davis 2020), and OutfitTransformer (Sarkar et al. 2022). TL: triplet loss, OL: outfit loss.

Compatibility Learning Results

Quantitative Results Here we present quantitative evaluation results on four different scales for benchmark datasets. We first evaluate our algorithms on the largest DeepHomes (Gadde, Feng, and Martinez 2021) dataset. As shown in Table 1 and 3, ICAR improves performance over the state-of-the-art by **9.5%** for FITB and 11.2 (about **23.3%**) for SFID score. To further test the effectiveness of our category embedding, we compare the FITB results when specify or not the category during inference time. If category is not specified, the FITB score drops from 87.1% to 83.9%.

Method	STL-F \uparrow	STL-H \uparrow	S2S \uparrow
IBR	58.5	57.0	56.5
Siamese Nets	67.1	72.4	63.0
BPR-DAE	61.1	64.2	59.3
Complete the Look	70.0	75.0	63.1
OutfitTransformer	65.0	77.0	67.2
ICAR (Ours)	75.3	86.6	71.4

Table 2: FITB Results on STL and Street2Shop (S2S). Our approach can improve FITB accuracy by 4.2%–9.6% over IBR (McAuley et al. 2015), Siamese Nets (Veit et al. 2015), BPR-DAE (Song et al. 2017a), Complete the Look (Kang et al. 2019), & OutfitTransformer (Sarkar et al. 2022). STL-F: STL-Fashion, STL-H: STL-Home.

We then evaluate our method on more existing datasets: STL-Home (Kang et al. 2019), STL-Fashion (Kang et al. 2019) and Street2Shop (Hadi Kiapour et al. 2015). For these three datasets, the human annotators also label products that have a similar style to the observed product and are compatible with the scene. This study suggests that the scene images and the products may not have one-to-one correspondence, thereby increasing the difficulty of style matching. We show in Table 2 and Table 3, that ICAR outperforms SOTA by 5.3% on STL-Fashion and 9.6% on STL-Home in FITB metric, and outperforms SOTA by 3.4 (22.3%) on STL-Fashion and 2.9 (31.8%) on STL-Home in terms of SFID metric.

Qualitative Results Next we present qualitative results on scene-based cross-domain CIR results. We sample scene im-

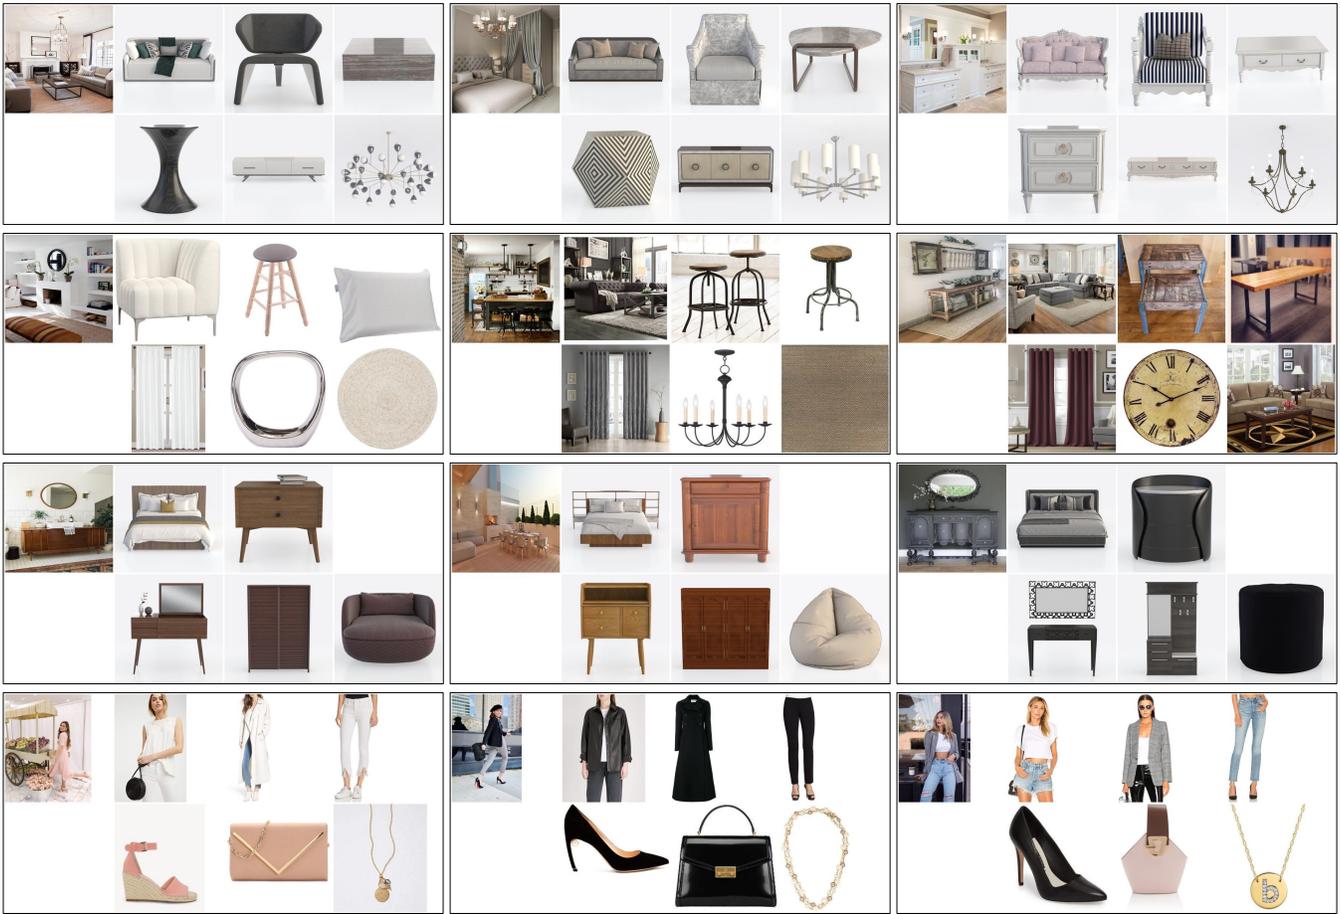


Figure 6: Scene-aware Cross-Domain CIR Qualitative Results. We show qualitatively our model is capable of retrieving stylistically compatible items from both seen (Row 2 and 4) and unseen domains (Row 1 and 3), given a home (Row 1-3) or fashion (Row 4) scene image. Column 1, 5, 9 are the input scene images. See supplementary materials for more examples.

Method	DR ↓	STL-F ↓	STL-H ↓	S2S ↓
GT VS Neg	49.0	15.8	10.2	9.1
OutfitTransformer	48.0	15.2	9.1	8.7
ICAR (Ours)	36.8	11.8	6.2	7.1

Table 3: SFID Results on DeepRooms, STL(F: fashion, H: home) and Street2Shop(S2S). Our approach can improve SFID accuracy by 11.2 (23.3%, DeepRooms) and 2.9 (31.8%, STL-Home) on furniture and by 3.4 (22.3%, STL-F) and 1.6 (18.4%, S2S) on fashion images, respectively, over OutfitTransformer (Sarkar et al. 2022). DR: DeepRooms.

ages from the STL-home and STL-fashion dataset test-split. Then we specify the item categories. As shown in Figure 6, from left to right and from top to bottom, our model generates diverse style recommendations, including but not only Contemporary (e.g. 1st, 2nd, 4th and 9th set), Classic (e.g. 3rd, 7th and 8th set), Industrial (e.g. 5th set), Rustic (e.g. 6th set). Refer to supplementary materials for more examples.

Method	FITB ↑	Condition
ICAR - fixed length masking	82.2	Predict category
ICAR - random length masking	83.9	Predict category
ICAR - fixed length masking	84.8	Given category
ICAR - random length masking	87.1	Given category

Table 4: Masking Method Comparison. Our random length masking outperforms fixed length masking.

Sequence Masking Validation In our FBT model, we propose a random shuffle and random length masking technique for unordered set generation. Here we ablate our proposed random length masking technique. In this experiment, we compare the results when we apply random length masking vs. fixed length masking to the input sequence. As shown in Table 4, we show that our random length masking approach outperforms fixed length masking for unordered set generation with or without target category been given.

Visual Compatibility Embedding Validation To validate our model’s ability in learning the visual style implicitly,



Figure 7: Learned Scene Image Embedding Clustering Results. To validate the style implicitly learned through our network, First column is the t-SNE on the randomly sampled 2k STL-home and STL-fashion test in split scene images (Columns 2-5).

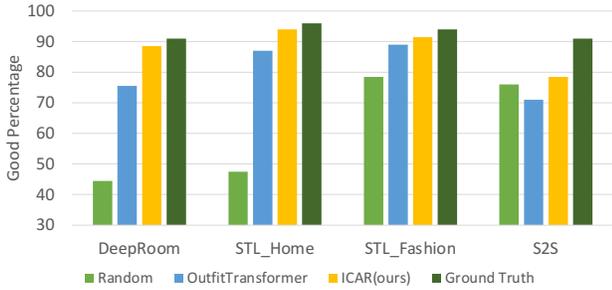


Figure 8: Human Ratings on Different Datasets. Our SFID score correlates better than SOTA with human judgement.

we perform t-SNE analysis on the learned embedding of the scene images sampled from the STL-home and STL-fashion datasets. In Figure 7, we show the clustering results on the left and the scene images from the clusters in column 2-5. The cluster labels are computed using the k-means ($K = 6$) method. We find some typical interior design styles such as the industrial style, rustic style and contemporary style in the clusters. More interestingly, we also find that our model pays attention to color. For example, in the C home scene-image cluster, our model learns to extract the Morandi type of color scheme. Given that our model is trained in a self-supervised way, we observe some mix in style within some of the clusters.

Human Perception Validation We conduct user studies by asking interior design experts to rate (Good: items are compatible, Neutral: one or two items are incompatible) 100 generated sets of furniture item images from the Random sample, OutfitTransformer(Sarkar et al. 2022), our method (ICAR), and Ground Truth for different datasets. As shown in Figure 8, ICAR outperforms all other methods, and those results are consistent with our SFID score (Table 3). We further compute the Pearson correlation coefficient to measure the association between SFID and human rating score (normalized as Ground truth score/Method score). We obtain a 0.7 average Pearson correlation value ($[-1, 1]$). This demonstrates that there is a strong positive association between our SFID score and human perception.

Similarity Learning Results

Here we validate the effectiveness of focusing on visual similarity learning in the first stage. We compare our method against four different types of learning target i.e. image classification, image reconstruction, image generation and image representation. And for each type of learning target, we choose the state-of-the-art method (shown in the first column in Table 5) except the image representation. For image representation, we train the model in a contrastive learning manner. All the models in this experiment are trained using the DeepRooms (Gadde, Feng, and Martinez 2021) dataset. As shown in Table 5, our model performs the best when the first stage focus on the visual similarity learning.

Method	FITB \uparrow	Learning Target Type
ICAR - VQGAN	73.1	Reconstruction
ICAR - Swin	80.5	Classification
ICAR - BEiT	73.3	Generation
ICAR - CL	75.0	Image Representation
ICAR - VS	87.1	Retrieval

Table 5: Similarity Learning: Visual similarity learning is the best for scene-based CIR. VQGAN (Esser, Rombach, and Ommer 2021), Swin (Liu et al. 2022), BEiT (Bao, Dong, and Wei 2021). CL: Contrastive Learning, VS: Visual Similarity.

Conclusion

In this paper, we introduce a *compatibility learning framework* using a novel category-aware “Flexible Bidirectional Transformer” (FBT), based on *Visual similarity* and *complementarity* to effectively retrieve a set of stylistically compatible items across domains, given a scene-image query. This learning framework is also generalizable and can be extended to other types of conditional cross-domain CIR tasks.

While our results show a promising direction, there is more to be explored, like text descriptors and video contents, to be used as compatibility signals.

Acknowledgments

Lin is supported partly by B. Mersky and Capital One E-NNovate Endowed Professorship.

References

- Ak, K. E.; Kassim, A. A.; Lim, J. H.; and Tham, J. Y. 2018. Learning attribute representations with localization for flexible fashion search. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7708–7717.
- Bao, H.; Dong, L.; and Wei, F. 2021. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*.
- Chen, Q.; Huang, J.; Feris, R.; Brown, L. M.; Dong, J.; and Yan, S. 2015. Deep domain adaptation for describing people based on fine-grained clothing attributes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5315–5324.
- Chen, W.; Huang, P.; Xu, J.; Guo, X.; Guo, C.; Sun, F.; Li, C.; Pfadler, A.; Zhao, H.; and Zhao, B. 2019. POG: personalized outfit generation for fashion recommendation at Alibaba iFashion. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2662–2670.
- Cheng, W.-H.; Song, S.; Chen, C.-Y.; Hidayati, S. C.; and Liu, J. 2021. Fashion meets computer vision: A survey. *ACM Computing Surveys (CSUR)*, 54(4): 1–41.
- Cheng, Z.-Q.; Wu, X.; Liu, Y.; and Hua, X.-S. 2017. Video2shop: Exact matching clothes in videos to online shopping images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4048–4056.
- Di, W.; Wang, C.; Bhardwaj, A.; Piramuthu, R.; and Sundaresan, N. 2013. Style finder: Fine-grained clothing style detection and retrieval. In *Proceedings of the IEEE Conference on computer vision and pattern recognition workshops*, 8–13.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- El-Nouby, A.; Neverova, N.; Laptev, I.; and Jégou, H. 2021. Training vision transformers for image retrieval. *arXiv preprint arXiv:2102.05644*.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12873–12883.
- Fu, H.; Cai, B.; Gao, L.; Zhang, L.-X.; Wang, J.; Li, C.; Zeng, Q.; Sun, C.; Jia, R.; Zhao, B.; et al. 2021. 3d-front: 3d furnished rooms with layouts and semantics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10933–10942.
- Gadde, R.; Feng, Q.; and Martinez, A. M. 2021. Detail Me More: Improving GAN’s Photo-Realism of Complex Scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13950–13959.
- Gallagher, A. C.; and Chen, T. 2008. Clothing cosegmentation for recognizing people. In *2008 IEEE conference on computer vision and pattern recognition*, 1–8. IEEE.
- Girdhar, R.; Carreira, J.; Doersch, C.; and Zisserman, A. 2019. Video action transformer network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 244–253.
- Guan, W.; Jiao, F.; Song, X.; Wen, H.; Yeh, C.-H.; and Chang, X. 2022a. Personalized Fashion Compatibility Modeling via Metapath-guided Heterogeneous Graph Learning. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 482–491.
- Guan, W.; Song, X.; Zhang, H.; Liu, M.; Yeh, C.-H.; and Chang, X. 2022b. Bi-directional Heterogeneous Graph Hashing towards Efficient Outfit Recommendation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 268–276.
- Guan, W.; Wen, H.; Song, X.; Yeh, C.-H.; Chang, X.; and Nie, L. 2021. Multimodal compatibility modeling via exploring the consistent and complementary correlations. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2299–2307.
- Hadi Kiapour, M.; Han, X.; Lazebnik, S.; Berg, A. C.; and Berg, T. L. 2015. Where to buy it: Matching street clothing photos in online shops. In *Proceedings of the IEEE international conference on computer vision*, 3343–3351.
- Han, X.; Wu, Z.; Jiang, Y.-G.; and Davis, L. S. 2017. Learning fashion compatibility with bidirectional lstms. In *Proceedings of the 25th ACM international conference on Multimedia*, 1078–1086.
- Hermans, A.; Beyer, L.; and Leibe, B. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Hsiao, W.-L.; and Grauman, K. 2017. Learning the latent” look”: Unsupervised discovery of a style-coherent embedding from fashion images. In *Proceedings of the IEEE International Conference on Computer Vision*, 4203–4212.
- Hsiao, W.-L.; and Grauman, K. 2018. Creating capsule wardrobes from fashion images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7161–7170.
- Hu, Y.; Yi, X.; and Davis, L. S. 2015. Collaborative fashion recommendation: A functional tensor factorization approach. In *Proceedings of the 23rd ACM international conference on Multimedia*, 129–138.
- Huang, J.; Feris, R. S.; Chen, Q.; and Yan, S. 2015. Cross-domain image retrieval with a dual attribute-aware ranking network. In *Proceedings of the IEEE international conference on computer vision*, 1062–1070.
- Jun, H.; Ko, B.; Kim, Y.; Kim, I.; and Kim, J. 2019. Combination of multiple global descriptors for image retrieval. *arXiv preprint arXiv:1903.10663*.
- Kang, W.-C.; Kim, E.; Leskovec, J.; Rosenberg, C.; and McAuley, J. 2019. Complete the look: Scene-based complementary product recommendation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10532–10541.
- Kiapour, M. H.; Yamaguchi, K.; Berg, A. C.; and Berg, T. L. 2014. Hipster wars: Discovering elements of fashion styles. In *European conference on computer vision*, 472–488. Springer.
- Li, X.; Wang, X.; He, X.; Chen, L.; Xiao, J.; and Chua, T.-S. 2020. Hierarchical fashion graph network for personalized outfit recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 159–168.
- Li, Y.; Cao, L.; Zhu, J.; and Luo, J. 2017. Mining fashion outfit composition using an end-to-end deep learning approach on set data. *IEEE Transactions on Multimedia*, 19(8): 1946–1955.
- Lin, Y.-L.; Tran, S.; and Davis, L. S. 2020. Fashion outfit complementary item retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3311–3319.

- Liu, S.; Feng, J.; Song, Z.; Zhang, T.; Lu, H.; Xu, C.; and Yan, S. 2012. Hi, magic closet, tell me what to wear! In *Proceedings of the 20th ACM international conference on Multimedia*, 619–628.
- Liu, Z.; Hu, H.; Lin, Y.; Yao, Z.; Xie, Z.; Wei, Y.; Ning, J.; Cao, Y.; Zhang, Z.; Dong, L.; et al. 2022. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12009–12019.
- Liu, Z.; Luo, P.; Qiu, S.; Wang, X.; and Tang, X. 2016. Deep-fashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1096–1104.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- McAuley, J.; Targett, C.; Shi, Q.; and Van Den Hengel, A. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, 43–52.
- Radenović, F.; Tolia, G.; and Chum, O. 2018. Fine-tuning CNN image retrieval with no human annotation. *IEEE transactions on pattern analysis and machine intelligence*, 41(7): 1655–1668.
- Sarkar, R.; Bodla, N.; Vasileva, M.; Lin, Y.-L.; Beniwal, A.; Lu, A.; and Medioni, G. 2022. OutfitTransformer: Outfit Representations for Fashion Recommendation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2263–2267.
- Shih, Y.-S.; Chang, K.-Y.; Lin, H.-T.; and Sun, M. 2018. Compatibility family learning for item recommendation and generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Simo-Serra, E.; and Ishikawa, H. 2016. Fashion style in 128 floats: Joint ranking and classification using weak data for feature extraction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 298–307.
- Song, X.; Feng, F.; Liu, J.; Li, Z.; Nie, L.; and Ma, J. 2017a. Neurostylist: Neural compatibility modeling for clothing matching. In *Proceedings of the 25th ACM international conference on Multimedia*, 753–761.
- Song, Y.; Li, Y.; Wu, B.; Chen, C.-Y.; Zhang, X.; and Adam, H. 2017b. Learning unified embedding for apparel recognition. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2243–2246.
- Su, T.; Song, X.; Zheng, N.; Guan, W.; Li, Y.; and Nie, L. 2021. Complementary factorization towards outfit compatibility modeling. In *Proceedings of the 29th ACM International Conference on Multimedia*, 4073–4081.
- Tan, R.; Vasileva, M. I.; Saenko, K.; and Plummer, B. A. 2019. Learning similarity conditions without explicit supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10373–10382.
- Tangseng, P.; Yamaguchi, K.; and Okatani, T. 2017. Recommending outfits from personal closet. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2275–2279.
- Taraviya, M.; Beniwal, A.; Lin, Y.-L.; and Davis, L. 2021. Personalized compatibility metric learning. In *KDD Workshop*, volume 1.
- Teh, E. W.; DeVries, T.; and Taylor, G. W. 2020. Proxynca++: Revisiting and revitalizing proxy neighborhood component analysis. In *European Conference on Computer Vision*, 448–464. Springer.
- Tran, S.; Du, M.; Chanda, S.; Manmatha, R.; and Taylor, C. 2019. Searching for Apparel Products from Images in the Wild. *arXiv preprint arXiv:1907.02244*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Veit, A.; Kovacs, B.; Bell, S.; McAuley, J.; Bala, K.; and Belongie, S. 2015. Learning visual clothing style with heterogeneous dyadic co-occurrences. In *Proceedings of the IEEE International Conference on Computer Vision*, 4642–4650.
- Yamaguchi, K.; Hadi Kiapour, M.; and Berg, T. L. 2013. Paper doll parsing: Retrieving similar styles to parse clothing items. In *Proceedings of the IEEE international conference on computer vision*, 3519–3526.
- Yamaguchi, K.; Kiapour, M. H.; Ortiz, L. E.; and Berg, T. L. 2012. Parsing clothing in fashion photographs. In *2012 IEEE Conference on Computer vision and pattern recognition*, 3570–3577. IEEE.
- Yang, M.; and Yu, K. 2011. Real-time clothing recognition in surveillance videos. In *2011 18th IEEE international conference on image processing*, 2937–2940. IEEE.
- Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R. R.; and Le, Q. V. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Yu, L.-F.; Yeung, S. K.; Terzopoulos, D.; and Chan, T. F. 2012. DressUp!: outfit synthesis through automatic optimization. *ACM Trans. Graph.*, 31(6): 134–1.
- Zhai, A.; Kislyuk, D.; Jing, Y.; Feng, M.; Tzeng, E.; Donahue, J.; Du, Y. L.; and Darrell, T. 2017. Visual discovery at pinterest. In *Proceedings of the 26th International Conference on World Wide Web Companion*, 515–524.
- Zhai, A.; and Wu, H.-Y. 2018. Classification is a strong baseline for deep metric learning. *arXiv preprint arXiv:1811.12649*.
- Zheng, N.; Song, X.; Niu, Q.; Dong, X.; Zhan, Y.; and Nie, L. 2021. Collocation and try-on network: Whether an outfit is compatible. In *Proceedings of the 29th ACM International Conference on Multimedia*, 309–317.