Bootstrapping Complete The Look at Pinterest

Eileen Li, Eric Kim, Andrew Zhai, Josh Beal, Kunlong Gu

Visual Search, Pinterest

{eileenli,erickim,andrew,jbeal,kgu}@pinterest.com

Figure 1: Complete The Look gives users outfit ideas and helps them find complementary products in related categories. The screen-captures above show outfit, jewelry, and shoes recommendations for the red striped blouse.

ABSTRACT

Putting together an ideal outfit is a process that involves creativity and style intuition. This makes it a particularly difficult task to automate. Existing styling products generally involve human specialists and a highly curated set of fashion items. In this paper, we will describe how we bootstrapped the Complete The Look (CTL) system at Pinterest. This is a technology that aims to learn the subjective task of "style compatibility" in order to recommend complementary items that complete an outfit. In particular, we want to show recommendations from other categories that are compatible with an item of interest. For example, what are some heels that go well with this cocktail dress? We will introduce our outfit dataset of over 1 million outfits and 4 million objects, a subset of which we will make available to the research community, and describe the pipeline used to obtain and refresh this dataset. Furthermore, we will describe how we evaluate this subjective task and compare model performance across multiple training methods. Lastly, we will share our lessons going from experimentation to working prototype, and how to mitigate failure modes in the production environment. Our work represents one of the first examples of an industrial-scale solution for compatibility-based fashion recommendation.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; • **Computing methodologies** → **Computer vision**; **Visual content-based**

indexing and retrieval; **Image representations**; **Machine learning**; **Neural networks**.

KEYWORDS

style modeling; embedding; visual search; recommender systems

ACM Reference Format:

Eileen Li, Eric Kim, Andrew Zhai, Josh Beal, Kunlong Gu. 2020. Bootstrapping Complete The Look at Pinterest. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '20)*, August 23–27, 2020, Virtual Event, CA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3394486.3403382>

1 INTRODUCTION

Since the number of online shoppers has grown exponentially in the past decade, platforms such as Amazon, Instagram, Taobao and Pinterest have all worked to create products that add value to the shopping experience.

In this crowded arena, Pinterest uniquely focuses on discovery and inspiration. Over 350M users visit the Pinterest website every month to discover new ideas in the realm of fashion, beauty, food and drinks, travel, home decor, and more. On Pinterest, they discover and save "pins", images with rich metadata attached, such as title, description, and url. Pinterest's ultimate goal is to turn inspiration into real-life actions and improvements. In recent years, visual search products such as Shop The Look [7] helped to close this loop by enabling users to make purchases directly from pins they have saved. Shop The Look uses the Unified Embedding model [37] trained on millions of pieces of Pinterest content to find products that are visually similar to the objects detected in the image.

As a follow-on to Shop The Look, we have been working on a novel shopping experience called Complete The Look (CTL) (see Figure 1). Whereas Shop The Look has the goal of visual exact match and retrieval, Complete The Look aims to find products that are visual complements. CTL helps users find ideas about how to

This work is licensed under a Creative Commons Attribution International 4.0 License.
KDD '20, August 23–27, 2020, Virtual Event, CA, USA
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7998-4/20/08.
<https://doi.org/10.1145/3394486.3403382>

style a particular product and gives them the option to continue their shopping experience in related categories. We hope to answer queries such as, “What are some hats that I might buy to wear with this jacket?”.

Completing an outfit is not a new problem, but existing solutions require fashion products that have been highly curated and matched by human stylists. There are also systems that use past engagement to build complementary recommendations (i.e., “You liked this item? Maybe you’ll like this other item too.”). In contrast, we describe a solution for when such explicit user engagement does not exist. Specifically, we offer our approach to **bootstrapping the Complete-the-Look system** and how we made it work with our diverse corpus of tens of millions of products. To do so requires answering some challenging questions, such as:

- How do we generate an outfit dataset that is high-quality and relevant to the shopping content on Pinterest?
- How do we evaluate this subjective task of “style compatibility”?
- How do we handle mislabeled or missing product metadata?
- How do we design a system that is performant, scalable, and easy to maintain?

In particular, our contributions are:

- (1) We describe an automatic way of generating an outfit dataset from the Pinterest platform of over 1M unique outfits and 4M fashion items, leveraging existing technologies in object detection, image style classification and attribute classification. To the best of our knowledge, this is the largest known dataset of fashion outfits, and we will release 100K outfits from our training set with our entire test set of 25K outfits.
- (2) We combine this dataset with a Convolutional Neural Net (CNN) to learn useful style embeddings. We present quantitative experimental results comparing performance across multiple training methodologies, including loss functions (contrastive, triplet, classification), data preprocessing, network architecture, and dataset collection.
- (3) We deploy and evaluate this model in an end-to-end recommendation system that performs retrieval from a diverse corpus of Pinterest shopping products, overcoming challenges in serving infrastructure, domain adaption, and metadata mislabeling.

2 RELATED WORKS

2.1 Visual Similarity

Fine-grained visual similarity and retrieval is a well-studied problem. The earlier approaches for modeling visual similarity relied mostly on hand-crafted features and attributes [19] [34]. More recent approaches using CNNs have been able to achieve state-of-the-art results [3] [24] [26] [15]. The challenge falls on curating a high-quality matching dataset (i.e., scene with bounding box and object pairs). The model trained on this dataset, usually by way of metric learning, learns to transform an image to its embedding representation. During retrieval, a query embedding is compared with many candidate embeddings to find the most similar results. This technology powers much of Pinterest’s visual search [17], and is the backbone to the system on which CTL is built.

2.2 Style Modeling

Earlier approaches in style modeling relied on crowdsourcing hand-labeled datasets [25]. This involved having explicit buckets for each style (e.g., “Bohemian” vs. “Classy”), which is limiting due to its subjective nature. Since then, there have been efforts in scaling the annotations for fashion datasets [31] [27] in order to enable much more fine-grained attribute classification. [13] uses the topic modeling approach to discover latent “styles”, but still requires a hand-crafted set of attributes. More recently, researchers have used outfits from the popular site Polyvore to train models such as SiameseCNN [9], sequence models (e.g., LSTM [8]), and graph-inspired networks [5] for learning style [20] [12] [33] [10]. In contrast to our system, most of these approaches are not optimized for retrieval but are instead trained for fashion compatibility classification.

2.3 Recommender Systems

Recommendation systems for fashion have grown in importance as as increasing percentage of shopping moves online. The earlier systems used traditional machine learning methods (e.g., SVM, logistic regression) trained on a hand-annotated dataset and retrieved recommendations from a limited corpus [23] [28]. More recently, Pinterest published a paper that focused on scene-to-product recommendation [18]—e.g., given a beach scene, recommend products that are complementary. Here we present a modified approach that handles product-to-product recommendation. Many recommendation systems, such as Alibaba’s iFashion [4], use past user engagement to train complementary models [21] [9]. We attempt the challenging task of bootstrapping such a system without explicit user data.

3 SYSTEM OVERVIEW

We define Complete The Look (CTL) as the problem of matching a single product to multiple complementary products. While a user is viewing an apparel product on Pinterest, we aim to offer them complementary products to help them “complete” their look. The serving system is real-time, generating recommendations in a fraction of a second through four stages: *query understanding*, *candidate generation*, *full scoring*, and *blending* as shown in Figure 2.

Query understanding takes the apparel product a user is currently viewing and enriches the query with additional features. For CTL, the core features are inferred from the CTL model, discussed more in Section 5.1, which predicts the product category of the query along with the style embedding. The product category prediction is then expanded into the outfit apparel categories that are complementary. As an example, for the “Shirts & Tops” query category, complementary categories to “complete” the look may be “Shoes”, “Sunglasses”, and “Jewelry & Watches”.

Candidate generation leverages the complementary categories from *query understanding* to restrict our apparel corpus of millions of products to the items that match the given categories. Each category has its own list of candidates. Note that the corpus categories are also generated from offline batch inference of the CTL model and the inverted indices per corpus categories are built offline.

Full scoring takes each complementary category’s candidate lists and runs approximate nearest neighbor (ANN) search of the

Figure 2: Serving system overview for Complete The Look

query’s style embedding against the complementary item’s style embeddings to rank the results within a category. These ANN indices are built offline for each inverted index and contain the CTL model predicted style embeddings for the candidates.

Blending merges each complementary category’s full-scored candidate lists and selects the final ordering to present to users.

4 DATASET

Fashion datasets such as Fashion136K [16] and StreetStyle [27] generally included “in-the-wild” street photos of people wearing clothes. More recently, composition outfit datasets from Polyvore website have been used for learning fashion tasks, such as FashionVC (20,726 outfits) [32], Maryland Polyvore (33,375 outfits) [8], and Polyvore Outfits (68,306 outfits) [33]. Polyvore was a popular website where people created and shared collages of outfits. Each outfit image consisted of multiple fashion items that came together to form a cohesive style. Unfortunately, the site was shut down in 2018.

Pinterest is a visual discovery engine with billions of pieces of diverse content saved by people around the world. In order to build a product that works well in this varied ecosystem, we need a reliable way to gather a large-scale outfit dataset for training that can be refreshed periodically. By using an Image Style classifier and object detector, we have implemented an extraction flow to decompose outfit images into their set of fashion objects. We chose to filter for images in collage-like sets similar to those that were popular on Polyvore. These “polyvore”-style images include objects that better match the domain of our product corpus, which typically consists of images with a clean background.

At a high level, the dataset extraction flow is comprised of the following steps:

- (1) We trained an Image Style classifier to identify “polyvore”-style images on Pinterest. These images are not necessarily from polyvore.com, but simply follow the same visual style of an outfit collage (Figure 3).
- (2) We ran our object detection model to gather bounding boxes and category labels for items on these “polyvore”-style images.
- (3) We cleaned up this dataset using a number of post-processing criteria for higher quality images. See Tables 2 and 3 for the dataset distributions broken down by category and number of items per outfit, respectively.

Figure 3: Pinterest images are funneled through steps 1) fashion and “polyvore”-style classification; 2) object detection and labeling; and 3) heuristics post-processing, resulting in high quality, high volume outfit dataset.

4.1 Image Style classifier

Image Style is a signal consisting of a family of multi-class and multi-label visual classifiers. Each classifier uses the Unified Embedding [37] as input and outputs a score between 0 and 1 for every style class. These classifiers are trained on human-curated image datasets with low label noise. There are five labels of interest for the CTL model: Polyvore, Product Shot, Stock Photo, Full Outfit, and Cropped Outfit (see Figure 4). These labels are used to distinguish common types of shopping content in the fashion domain. By filtering on the Polyvore score at a high-precision threshold of 0.9, around 35M images in the fashion domain are obtained.

Figure 4: Examples of Image Style categories for CTL model, from left to right: (1) Polyvore, (2) Product Shot, (3) Stock Photo, (4) Full Outfit, and (5) Cropped Outfit.

4.1.1 Domain difference. We use the “polyvore”-style to collect training data because objects detected from these images resemble products on solid backgrounds (i.e., Product Shot). As expected, quality decreases dramatically for out-of-domain products (i.e., Stock Photo, Full Outfit, Cropped Outfit). To ensure quality, we restrict CTL queries and results using the Image Style classifier to filter for Product Shots only (see Figure 5).

Figure 5: (Left) Product Shot pins that match the image domain of our training set. We restrict CTL to only these images. (Right) Cropped Outfit pins that we filter out for CTL.

4.2 Object Detection

To decompose each “polyvore”-style image into its individual articles of clothing, e.g., “Shirts & Tops” and “Pants” (Figure 3), we utilize an object detection model that is trained on fashion categories. Specifically, we use a Faster-RCNN [30] with a ResNeXt101 [35] backbone and Feature Pyramid Networks [22], trained using the Detectron [6] framework.

Our detection training set consists of 251,000 training images with 727,000 bounding boxes in the home decor and fashion domains. We set aside a validation set consisting of 10,000 images with 30,000 bounding boxes.

At inference time, we use a whitelist of 21 fashion categories (Table 2) to restrict the bounding boxes to the CTL categories. Of the 35M “polyvore”-style images, about 10M of them have at least one CTL object detected.

mAP	Precision	Recall
77.2	76.0	82.4

Table 1: Detection evaluation metrics on our in-house test set of 10k fashion images. “mAP” is mean average precision. Precision and recall are obtained by choosing the operating point that maximizes the F1 score.

To assess the quality of the detection model, we collected an in-house test set of 10k images in the fashion category (see Table 1).

We found that the performance of the object detection model was satisfactory for the CTL dataset generation pipeline.

Label	#Items(full)	#Items(100K)	#Items(test)
Shoes	856,775	94,059	22,728
Handbags	810,501	88,955	20,729
Shirts & Tops	519,625	57,250	14,451
Pants	468,845	51,706	11,666
Coats & Jackets	397,353	43,552	9,698
Dresses	254,806	28,066	6,903
Jewelry	222,743	24,409	7,813
Hats	168,394	18,581	4,343
Skirts	138,343	15,152	3,630
Sunglasses	86,802	9,343	2,174
Shorts	84,743	9,265	2,166
Scarves & Shawls	52,233	5,751	1,218
Watches	51,613	5,605	1,313

Table 2: Number of items in full, released, and test datasets from top 13 out of 21 total categories.

#Items	#Outfits(full)	#Outfits(100K)	#Outfits(test)
3	223,240	24,545	5,383
4	421,815	46,545	9,288
5	251,835	27,643	6,616
6	72,916	7,927	2,786
7	10,767	1,127	722
8	938	102	165

Table 3: Number of clothing articles in an outfit; in full, 100K and test datasets.

4.3 Dataset Cleanup

Figure 6: We clean up the dataset using a number of heuristics to ensure clean bounding boxes with a cohesive set of fashion items.

We want outfits that have good style cohesion and high diversity in category and color. For example, training on a corpus consisting mostly of outfits with white tops and blue jeans may lead to a model that heavily biases towards recommending jeans for every query. We also want crops with clean backgrounds so that they can better match the conditions of our product images. To help curate a higher quality dataset, we applied a series of post-processing steps:

- Discard overlapping bounding boxes by applying non-maximum suppression¹.
- Discard bounding boxes that are too small, specifically less than 5% of total image area.
- Enforce that each image has between 3 and 8 clothing items. We found that outfits outside of this range tend to be noisy and lack cohesion, e.g., consisting of several different outfits in the same image.
- Enforce at least 3 different types of clothing articles per outfit. This further guarantees that the image contains a complete outfit.
- Discard outfits where all items are of the same color. Many outfit images are monochromatic, resulting in models biased for matching color instead of style. Spot-checking results seem to show improvements with this filter, though the problem still exists (see qualitative feedback in Section 6.5).

We tried curating by user engagement (e.g., clicks and repins) and image age (i.e., number of days on Pinterest) to get popular, trending outfits, but we did not find a noticeable difference in the generated outfits.

The final dataset has **1,006,519 outfits and 4,246,430 fashion objects**. To the best of our knowledge, this is the largest known dataset of fashion outfits. We keep 24,960 outfits (109,471 items) as a holdout set and train on the remaining 981,559 outfits (4,136,959 items). The female:male ratio is about 10:1. 100K outfits from the training set and the entire test set are made publicly available².

5 METHODOLOGY

5.1 Model Architecture

The CTL model is comprised of three modules: the *visual featurizer*, *category classifier*, and *style extractor*. We try to leverage as much existing Pinterest technology as possible, lowering the risk and maintenance cost of the system.

The visual featurizer uses the same SE-ResNext101 backbone [14] as Pinterest’s Unified Embedding [37]. Unified Embedding is a multi-task learning model that has been trained for three visual discovery tasks: Flashlight [2], Lens [36], and Shop The Look [7]. The Unified Embedding is the primary image-based embedding used by many systems at Pinterest.

“Layer4” of the featurizer backbone is fed into two separate networks: one for category classification, and another for style learning.

The category classifier predicts the top category label for each input image, from a total of 21 fashion CTL categories. These predictions are important because we need to determine the query category (e.g., “Dresses”) and also filter CTL results for particular candidate categories (e.g., “Shoes” and “Coats & Jackets”). At

Figure 7: The CTL model has three components: visual featurizer, category predictor, and style . We compare three different ways to train: A) classification “proxy” loss, B) contrastive loss, and C) triplet loss.

serving time, products are stored in our retrieval index keyed by the category prediction from this classifier, so we can efficiently retrieve and filter results on any set of product categories.

The ground truth labels for this classifier come from the detection model (see Table 1 for detection evaluation metrics). The category classifier is implemented as two fully connected (FC) layers, achieving accuracy of 98.15% across all categories on the test set.

The style extractor is a neural network with two FC-layers that outputs a 128-dimensional style embedding. We use batch normalization and a dropout rate of 0.5. We experimented with three different training methods: a classification network with “proxy” loss, contrastive loss (siamese) [9] and triplet loss [11]. Our production CTL system uses the best performing model with triplet loss. See section 6 for experiment results.

The motivation for the classification network is its similarity to the Unified Embedding model [37], if for example, we decide to add CTL as an additional task to our Unified Embedding training. In the classification model, each outfit has a unique instance label. At training time, we feed into the model the image of a fashion item. From the output layer of the style extractor, we calculate the softmax cross entropy loss for predicting into a subset of “proxy” instances. We sample the number of proxies to be 2048 which is consistent with Unified Embedding training.

To sample pairs for training a siamese network, we follow the sampling methodology in [9]. Similarly, we guarantee that each of our positive pairs comes from different categories (“heterogeneous dyads”) while negatives are randomly sampled regardless of category. This is done to dissuade learning visual similarity in place of style compatibility. We then optimize using contrastive loss:

$$\mathcal{L}_{contrast}(i, j) := y_{ij}D_{ij}^2 + (1 - y_{ij})[\alpha - D_{ij}]_+^2$$

¹We use an IOU threshold of 0.1

²<https://github.com/eileenforwhat/complete-the-look-dataset>

Figure 8: We take this figure directly from the Unified Embedding paper [37]. It depicts how we get a subset of “proxy” instances for calculating classification loss.

where D_{ij} denotes the distance between samples i and j , and y_{ij} is 1 if samples i and j have the same label and 0 otherwise.

We also train a triplets network by sampling (*anchor_image*, *pos_image*, *neg_image*). This is similar to the scene-based CTL paper [18], where they trained a triplet network using (*anchor_scene*, *pos_product*, *neg_product*). In our case, however, every item in the triplet is a crop of a product. We optimize using the loss function:

$$\mathcal{L}_{\text{triplet}}(a, p, n) := [D_{ap}^2 - D_{an}^2 + \alpha]_+$$

where a, p, n denote the query, positive, and negative samples respectively; D_{ap}^2 denotes the distance between the anchor and positive samples, D_{an}^2 denotes the distance between the anchor and negative samples; and α is the margin term.

5.2 Training and deployment details

We train our models using the PyTorch framework [29] with 8 GPUs and the model takes about 10 hours to converge. We use Apex [1] mixed precision to increase training efficiency. We use the Adam optimizer with learning rate of 0.048. For deployment, we deploy the PyTorch model to C++ directly by serializing to TorchScript.

6 EVALUATION

6.1 Methods

Evaluating style compatibility is challenging because of its subjective nature. We compare our model performance using multiple methods to get a comprehensive assessment. We have a test set of 25K outfits (109,847 items), which we use to measure retrieval recall and Fill-in-the-Blank (FITB) accuracy. We also perform end-to-end evaluation on Pinterest’s real product corpus, leveraging our in-house fashion specialists for labeling. Lastly, CTL is launched internally, allowing us to conduct user studies and gain valuable insights about how real users think when they use the product.

6.1.1 Recall@{1, 5, 10}. We measure the retrieval recall on a sub-sampled corpus. Since CTL is ultimately used in a retrieval setting, measuring Recall@K (R@K) most closely mirrors the production task.

In order to calculate exact recall, we limit the number of items per outfit in the test set to be exactly five, removing outfits that do not meet this criteria. For each item of an outfit, we calculate the R@{1,

5, 10} by retrieving the top matches using k-nearest neighbors, with euclidean distance between style embeddings as the distance measure. The retrieval corpus includes the other items from the same outfit (positives) and randomly sampled items (negatives). We calculate R@K according to two ways of generating the corpus: 1) sampled across all categories, and 2) restricted to one category at a time (and taking the mean). In both cases, the total size of the retrieval corpus is $N=200$. Thus, R@K is computed by the following:

$$R@K = \frac{(\# \text{ positives in top } K)}{\min[(\text{total num items in outfit}) - 1, K]}$$

Since metrics gathered from both methods of sampling are highly correlated, we only report on the latter in this paper for simplicity. In Figure 9, we visualize some results from the retrieval task.

Figure 9: Example of the retrieval evaluation. For each item in an outfit, we calculate R@{1,5,10}. The result is considered correct if it belongs to the same outfit as the query.

6.1.2 Fill-in-the-Blank (FITB). For the FITB task, we remove one item from each outfit. The model then has to pick from the positive item and three randomly sampled items that belong to the same category. Unlike the retrieval task, FITB uses multiple products as queries rather than just a single product. Current SOTA methods for this task (e.g., [8], [33]) are not suited for our large-scale one product to multi-complementary product retrieval.

A downside of both R@K and FITB evaluation tasks is that they only consider clothing articles from the same “polyvore”-style image to be positive examples for a given query clothing article. In fact, it’s likely that clothing articles from other “polyvore”-style images in the test corpus could be compatible with the query; for instance, blue jeans tend to be compatible with many other items. We address this shortcoming by using human evaluation (see Section 6.1.3).

6.1.3 Human Judgment. Human judgment is extremely valuable because it directly evaluates how our models perform compared with a human stylist. The labeling template we developed is shown in Figure 10. Since this is a highly specialized task, we decided to only use in-house fashion specialists who follow a comprehensive set of labeling guidelines. These labeling guidelines describe each failure mode (e.g. mismatch color, print, season) with examples and attempt to reduce any subjectivity. We have picked the best performing variant for each of {Siamese, Triplets, and Classification} (refer to Section 5.1) for comparison.

The questions for this task were generated by sampling 120 products from Pinterest and running the CTL system end-to-end

(see Figure 2). We compute precision by evaluating whether each (*query, candidate*) pair is a compatible match. In addition, we also asked for qualitative feedback such as common failures and most jarring mistakes. Both quantitative and qualitative feedback can be found in Section 6.5.

Figure 10: Human evaluation task: for each model, we retrieve CTL results for a sampled set of products, and ask whether each (*query, candidate*) pair is compatible.

6.2 Comparing style extractor training methods

We explored several approaches to training the CTL style extractor for the fashion visual complements task (refer to 5.1). We report metrics on R@{1,5,10} and FITB accuracy using our test set.

Scene-based CTL. This is the same dataset and model as [18], and we include this evaluation as a baseline (See Table 4). This model was trained for the task of retrieving complementary products given a scene image (e.g., the beach). The training uses (*anchor_scene, positive_product, negative_product*) with triplet loss, and a separate network for scene and product images. For this evaluation, we only use the product network since there are no scenes in our dataset. We see that style embeddings trained in this manner do not generalize well to the new task of product-to-products recommendation.

Classification with "proxy" loss. We found that training for the classification task as described in Section 5.1 showed promising qualitative results, but performed worse than metric learning methods (See Table 4).

Siamese. We train with contrastive loss. We try the sampling strategy described in [9] of 16:1 negative-to-positive ratio, but we find that a simple 1:1 strategy performs better (See Table 4).

Triples. We generate triplets (*anchor, positive, negative*) from our outfits by taking any two items from the same outfit as anchor

and positive. We compare performance when sampling negatives randomly or from the same category as positives. We found that triplets training while restricting by category out-performs all other methods (See Table 4).

Method	R@1	R@5	R@10	FITB
CTL(scene-based)	3.9	16.5	29.6	40.1
Classification	15.8	41.6	58.8	71.9
Siamese(16:1)	15.2	40.7	57.6	71.9
Siamese(1:1)	16.0	46.0	64.3	74.3
Triples(random)	18.5	49.4	67.5	77.6
Triples(cat)	20.3	51.4	68.8	78.5

Table 4: Recall@1,5,10 and FITB metrics for different CTL training methods.

6.3 Comparing visual featurizers

In this section we compare the effects of using Unified Embedding, which has the advantage of additional training on Pinterest data, versus ImageNet for weight initialization of the visual featurizer. These experiments were conducted using triplet loss with category-restricted sampling, which is our best performing model (see Table 4).

We find that pretraining on Pinterest data yields substantial gains relative to using default ImageNet weights, with an R@1 gain from 17.6% to 20.3% (see Table 5). This suggests that there is a substantial domain shift from ImageNet images to Pinterest images.

Method	R@1	R@5	R@10	FITB
ImageNet	17.6	47.2	64.0	75.5
Unified	20.3	51.4	68.8	78.5

Table 5: A comparison of ImageNet vs. Unified Embedding weight initialization for the visual featurizer.

6.4 Comparing training set sizes

In this section we compare how decreasing training set size affects performance (see Table 6). This helps us answer the question, "How much training data is enough?". We use triplet loss with category-restricted sampling, which is our best performing model in Table 4. Performance on the test set continues to increase as we add training examples, although 10K to 100K outfits sees a significant improvement (R@1 12.1% -> 16.4%) while 100K to 1M outfits sees a lesser gain (R@1 16.4% -> 20.3%).

#Outfits	#Objects	R@1	R@5	R@10	FITB
10K	46K	12.1	37.0	54.1	66.0
100K	453K	16.4	44.3	62.6	74.8
200K	835K	16.8	46.3	64.5	75.1
500K	2.1M	19.1	49.2	66.6	77.1
1M (full)	4.13M (full)	20.3	51.4	68.8	78.5

Table 6: An ablation of model performance versus training dataset size. The 100K dataset is released.

6.5 Human evaluation and qualitative feedback

We asked our in-house fashion specialists to evaluate the end-to-end CTL quality through the eyes of a stylist (refer to Section 6.1.3). The results are shown in Table 7 and Figure 11. The quantitative results are consistent with our other evaluations, and shows that the triplet loss with category-restricted sampling is the best performing model.

Method	Overall Precision
Siamese (1:1)	47.7
Classification	50.6
Triplets (cat)	55.0

Table 7: Human evaluation results comparing best performing variants from Table 4.

Figure 11: Example of correct vs. incorrect evaluation results for a given query, by the standards of a human stylist.

We also asked our in-house specialists for qualitative feedback, and it was clear from their feedback that—when compared with a human stylist—CTL relies more on color and pattern matching. This can result in recommendations that either lack diversity or the “human touch” if matches are always the same color or pattern, or are extremely off if the color or pattern does not match entirely. This suggests additional investigation into ensuring that there is sufficient clothing diversity in the training set, as Section 4.3 describes.

6.6 User studies

As of early 2020, we have a prototype of CTL that is released internally. Qualitative results can be found in Figure 12. Although we expect more improvements before general launch, this prototype enabled us to conduct user studies to get early feedback.

We invited 5 Pinterest users to come in for hour-long sessions during which we guided them through using the CTL product feature. Users had high expectations for CTL results, and they wanted to see matches that resembled curation from a human stylist. Overall, the user studies demonstrated that users do see the value of the CTL product, which confirms our continued investment in building it out.

7 CONCLUSION AND FUTURE WORK

We shared our approach and results from bootstrapping to productionizing the Complete The Look (CTL) system.

We implemented an automated data pipeline that generates a labeled image dataset for the fashion outfit recommendation task. For our model, we ran comprehensive sets of experiments comparing multiple methods, relying on offline metrics such as R@K that best mirror our product use case. Serving the model in production proved challenging as the problem scaled to tens of millions of unique products. To keep recommendation quality high, we trained and utilized image classifiers that help us trigger CTL only on images that match the training set’s domain.

In the future, we hope to add more training data from different image styles (e.g., cropped outfit, full outfit, stock photo) to make the model more robust to domain variations. We also hope to incorporate price as an input, since it is a poor user experience to match a \$8,000 ring with a \$20 dress. Furthermore, we hope to add user personalization, so that we can tailor recommendations to a particular user’s style.

8 ACKNOWLEDGMENTS

The authors would like to thank the rest of the visual search team, especially Chuck Rosenberg for his leadership, Kofi Boakye for his thoughtful editing, and Angela Guo for her feedback and guidance. We would also like to thank Mariellen Barros and Marta Scotto for labeling efforts; Joyce Zha, Claire Li, and Suzy Kim for frontend support; Cherrylyn Cawit and Shilpa Banerjee for conducting user studies.

REFERENCES

- [1] [n.d.]. apex.amp — Apex 0.1.0 documentation. <https://nvidia.github.io/apex/amp.html>. (Accessed on 06/12/2020).
- [2] [n.d.]. Flashlight for Pinterest – Visual Discovery on Steroids. <http://www.mcngmarketing.com/flashlight-pinterest-visual-discovery-steroids/>. (Accessed on 02/05/2020).
- [3] Sean Bell and Kavita Bala. 2015. Learning Visual Similarity for Product Design with Convolutional Neural Networks. *ACM Trans. Graph.* 34, 4, Article 98 (July 2015), 10 pages. <https://doi.org/10.1145/2766959>
- [4] Wen Chen, Pipei Huang, Jiaming Xu, Xin Guo, Cheng Guo, Fei Sun, Chao Li, Andreas Pfadler, Huan Zhao, and Binqiang Zhao. 2019. POG: Personalized Outfit Generation for Fashion Recommendation at Alibaba iFashion. (05 2019).
- [5] Zeyu Cui, Zekun Li, Shu Wu, Xiaoyu Zhang, and Liang Wang. 2019. Dressing as a Whole: Outfit Compatibility Learning Based on Node-wise Graph Neural Networks. (02 2019). <https://doi.org/10.1145/3308558.3313444>
- [6] Ross Girshick, Ilija Radosavovic, Georgia Gkioxari, Piotr Dollár, and Kaiming He. 2018. Detectron. <https://github.com/facebookresearch/detectron>.
- [7] Kunlong Gu. 2019. Automating Shop the Look on Pinterest - Pinterest Engineering Blog - Medium. <https://medium.com/pinterest-engineering/automating-shop-the-look-on-pinterest-a17aeff0eae2>. (Accessed on 02/02/2020).
- [8] Xintong Han, Zuxuan Wu, Yu-Gang Jiang, and Larry S. Davis. 2017. Learning Fashion Compatibility with Bidirectional LSTMs. In *Proceedings of the 25th ACM International Conference on Multimedia* (Mountain View, California, USA) (MM '17). Association for Computing Machinery, New York, NY, USA, 1078–1086. <https://doi.org/10.1145/3123266.3123394>
- [9] Ruining He, Charles Packer, and Julian McAuley. 2016. Learning Compatibility Across Categories for Heterogeneous Item Recommendation. 937–942. <https://doi.org/10.1109/ICDM.2016.0116>
- [10] Ruining He, Charles Packer, and Julian McAuley. 2016. Learning Compatibility Across Categories for Heterogeneous Item Recommendation. 937–942. <https://doi.org/10.1109/ICDM.2016.0116>
- [11] Elad Hoffer and Nir Ailon. 2014. Deep Metric Learning Using Triplet Network. https://doi.org/10.1007/978-3-319-24261-3_7
- [12] Wei-Lin Hsiao and Kristen Grauman. 2017. Creating Capsule Wardrobes from Fashion Images. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2017), 7161–7170.
- [13] Wei-Lin Hsiao and Kristen Grauman. 2017. Learning the Latent “Look”: Unsupervised Discovery of a Style-Coherent Embedding from Fashion Images. (07 2017).
- [14] Jie Hu, Li Shen, and Gang Sun. 2017. Squeeze-and-Excitation Networks. *CoRR abs/1709.01507* (2017). arXiv:1709.01507 <http://arxiv.org/abs/1709.01507>

Figure 12: Complete The Look retrieval results from Pinterest product corpus. The first column is the query image and the four images to the right are recommendations from different categories.

- [15] Junshi Huang, Rogerio Feris, Qiang Chen, and Shuicheng Yan. 2015. Cross-Domain Image Retrieval with a Dual Attribute-Aware Ranking Network. (05 2015). <https://doi.org/10.1109/ICCV.2015.127>
- [16] Vignesh Jagadeesh, Robinson Piramuthu, Anurag Bhardwaj, Wei di, and Neel Sundaresan. 2014. Large Scale Visual Recommendations From Street Fashion Images. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (01 2014). <https://doi.org/10.1145/2623330.2623332>
- [17] Yushi Jing, David C. Liu, Dmitry Kislyuk, Andrew Zhai, Jiajing Xu, Jeff Donahue, and Sarah Tavel. 2015. Visual Search at Pinterest. *CoRR* abs/1505.07647 (2015). [arXiv:1505.07647](http://arxiv.org/abs/1505.07647) <http://arxiv.org/abs/1505.07647>
- [18] W. C. Kang, E. Kim, J. Leskovec, C. Rosenberg, and J. McAuley. 2018. Complete the Look: Scene-based Complementary Product Recommendation. *arXiv:1812.01748* (2018). [arXiv:1812.01748](http://arxiv.org/abs/1812.01748)
- [19] Adriana Kovashka, Devi Parikh, and Kristen Grauman. 2012. WhittleSearch: Image Search with Relative Attribute Feedback. *Proceedings / CVPR, IEEE Computer Society Conference on Computer Vision and Pattern Recognition. IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 115, 2973–2980. <https://doi.org/10.1109/CVPR.2012.6248026>
- [20] Hanbit Lee, Jinseok Seol, and Sang-goo Lee. 2017. Style2Vec: Representation Learning for Fashion Items from Style Sets. (08 2017).
- [21] Yuncheng Li, Liangliang Cao, Jiang Zhu, and Jiebo Luo. 2016. Mining Fashion Outfit Composition Using An End-to-End Deep Learning Approach on Set Data. *IEEE Transactions on Multimedia* PP (08 2016). <https://doi.org/10.1109/TMM.2017.2690144>
- [22] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. 2016. Feature Pyramid Networks for Object Detection. *CoRR* abs/1612.03144 (2016). [arXiv:1612.03144](http://arxiv.org/abs/1612.03144) <http://arxiv.org/abs/1612.03144>
- [23] Si Liu, Tam Nguyen, Jia Shi Feng, Meng Wang, and Shuicheng Yan. 2012. Hi, magic closet, tell me what to wear! 1333–1334. <https://doi.org/10.1145/2393347.2396470>
- [24] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. 2016. DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations. 1096–1104. <https://doi.org/10.1109/CVPR.2016.124>
- [25] Alexander C. Berg Tamara L. Berg M. Hadi Kiapour, Kota Yamaguchi. 2014. Hipster Wars: Discovering Elements of Fashion Styles. In *European Conference on Computer Vision*.
- [26] Svetlana Lazebnik Alexander C. Berg Tamara L. Berg M. Hadi Kiapour, Xufeng Han. 2015. Where to Buy It: Matching Street Clothing Photos in On-line Shops. In *International Conference on Computer Vision*.
- [27] Kevin Matzen, Kavita Bala, and Noah Snavely. 2017. StreetStyle: Exploring world-wide clothing styles from millions of photos. *arXiv preprint arXiv:1706.01869* (2017).
- [28] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Hengel. 2015. Image-Based Recommendations on Styles and Substitutes. (06 2015). <https://doi.org/10.1145/2766462.2767755>
- [29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*. 8024–8035.
- [30] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *CoRR* abs/1506.01497 (2015). [arXiv:1506.01497](http://arxiv.org/abs/1506.01497) <http://arxiv.org/abs/1506.01497>
- [31] Edgar Simo-Serra and Hiroshi Ishikawa. 2016. Fashion Style in 128 Floats: Joint Ranking and Classification Using Weak Data for Feature Extraction. 298–307. <https://doi.org/10.1109/CVPR.2016.39>
- [32] Xueming Song, Fuli Feng, Jinhuan Liu, Zekun Li, Liqiang Nie, and Jun Ma. 2017. NeuroStylist: Neural Compatibility Modeling for Clothing Matching. 753–761. <https://doi.org/10.1145/3123266.3123314>
- [33] Mariya Vasileva, Bryan Plummer, Krishna Dusad, Shreya Rajpal, Ranjitha Kumar, and David Forsyth. 2018. Learning Type-Aware Embeddings for Fashion Compatibility. (03 2018).
- [34] Xianwang Wang and Tong Zhang. 2011. Clothes search in consumer photos via color matching and attribute learning. 1353–1356. <https://doi.org/10.1145/2072298.2072013>
- [35] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. 2016. Aggregated Residual Transformations for Deep Neural Networks. *CoRR* abs/1611.05431 (2016). [arXiv:1611.05431](http://arxiv.org/abs/1611.05431) <http://arxiv.org/abs/1611.05431>
- [36] Andrew Zhai. [n.d.]. Building Pinterest Lens: a real world visual discovery system. <https://medium.com/pinterest-engineering/building-pinterest-lens-a-real-world-visual-discovery-system-59812d8cbfbc>. (Accessed on 02/05/2020).
- [37] Andrew Zhai, Hao-Yu Wu, Eric Tzeng, Dong Huk Park, and Charles Rosenberg. 2019. Learning a Unified Embedding for Visual Search at Pinterest. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Anchorage, AK, USA) (KDD '19)*. Association for Computing Machinery, New York, NY, USA, 2412–2420. <https://doi.org/10.1145/3292500.3330739>