

Visual Discovery at Pinterest

Andrew Zhai^{1*}, Dmitry Kislyuk^{1*}, Yushi Jing^{1*}, Michael Feng¹
Eric Tzeng^{1,2}, Jeff Donahue^{1,2}, Yue Li Du¹, Trevor Darrell²

¹ Visual Discovery, Pinterest ² University of California, Berkeley
{andrew,dkislyuk,jing,m,etzeng,jdonahue,shirleydu}@pinterest.com
trevor@eecs.berkeley.edu

ABSTRACT

Over the past three years Pinterest has experimented with several visual search and recommendation systems, from enhancing existing products such as Related Pins (2014), to powering new products such as Similar Looks (2015), Flashlight (2016), and Lens (2017). This paper presents an overview of our visual discovery engine powering these services, and shares the rationales behind our technical and product decisions such as the use of object detection and interactive user interfaces. We conclude that this visual discovery engine significantly improves engagement in both search and recommendation tasks.

Keywords

visual search; recommendation systems; convnets; object detection

1. INTRODUCTION

Visual search and recommendations [8], collectively referred to as *visual discovery* in this paper, is a growing research area driven by the explosive growth of online photos and videos. Commercial visual search systems such as Google Goggles and Amazon Flow are designed to retrieve photos with the exact same object instance as the query image. On the other hand, recommendation systems such as those deployed by Google Similar Images [42], Shopping [7] and Image Swirl [22] display a set of visually similar photos alongside the query image without the user making an explicit query.

At Pinterest we experimented with different types of visual discovery systems over a period of three and half years. We benefited from the confluence of two recent developments – first, advances in computer vision, especially the use of convolutional networks and GPU acceleration, have led to significant improvements in tasks such as image classification and object detection. Second, a substantial number

*indicates equal contribution.

©2017 International World Wide Web Conference Committee (IW3C2), published under Creative Commons CC BY 4.0 License. WWW'17 Companion, April 3–7, 2017, Perth, Australia. ACM 978-1-4503-4914-7/17/04. <http://dx.doi.org/10.1145/3041021.3054201>




Figure 1: Pinterest Flashlight: User can select any objects in the image (e.g. lamp, desk, shelf) as a visual search query.




Figure 2: If objects are automatically detected, Flashlight displays a clickable “dot” for faster navigation.

of users prefer using discovery systems to *browse* (finding inspirational or related content) rather than to *search* (finding answers). With hundreds of millions of users (who are often browsing for ideas in fashion, travel, interior design, recipes, etc.), Pinterest is a unique platform to experiment with various types of visual discovery experiences.

Our previous work [21] presented preliminary results showing that convnet features and object detection can be used effectively to improve user engagement in visual search systems. We subsequently launched Pinterest Flashlight [2], an interactive visual discovery tool that allows user to select any object in the image (e.g. lamp, desk, shelf) as visual queries, as shown in Figure 1. If the system is certain about the location of an object, either through automatic object




Figure 3: By indexing objects instead of whole images in our visual search system, Pinterest Lens finds objects *within* the images.

detection [14] or collaborative filtering, Flashlight displays a clickable “dot” on the image for faster navigation, as shown in Figure 2. Recently we launched Pinterest Lens [1], a visual discovery experience accessed via a mobile phone camera, as shown in Figure 3. In this release we also applied object detection to the billions of images on Pinterest so that query objects can be matched against other visually similar objects *within* the catalog images.

This paper gives a general architectural overview of our visual discovery system and shares the lessons we learned from scaling and integrating visual features into products at Pinterest. For example, we investigate the performance of popular classification networks for retrieval tasks and evaluate a binarization transformation [4] to improve the retrieval quality and efficiency of these classification features. This feature transformation enables us to drastically reduce the retrieval latency while improving the relevance of the results in large-scale settings. We also describe how we apply object detection to multiple visual discovery experiences including how to use detection as a feature in a general image recommendation system and for query normalization in a visual search system.

The rest of the paper is organized as follows: Section 2 and 3 give a brief introduction on the Pinterest image collection essential to our visual discovery experience and survey related works. Section 4 and 5 present the visual features and object detectors used in our system. Section 6 describes how visual features can be used to enhance a generic image recommendation system. Section 7 and 8 presents our experience launching Pinterest Flashlight and our latest application, Pinterest Lens.

2. PINTEREST IMAGES

Our work benefited significantly from having access to billions of catalog photos. On Pinterest, hundreds of millions of users organize images and videos around particular topics into boards as shown in Figure 4, which result in a very large-scale and hand-curated collection with a rich set of metadata. Most Pin images are well annotated: when a person saves an image onto a board, a Pin is created around an image and a brief text description supplied by the user. When the same Pin is subsequently saved to a new board by a different user, the original Pin gains additional metadata



Figure 4: Pinterest users collect images into varieties of themed boards.

that the new user provides. This data structure continues to expand every time an image is shared and saved. Such richly annotated datasets gave us the ability to generate training sets that can potentially scale to hundreds of thousands of object categories on Pinterest.

3. RELATED WORK

Visual discovery systems: Amazon Flow, Google and Bing Similar Images are examples of widely used visual search and recommendation systems. Various works [19] [12] have been proposed to improve the ranking of the image search results using visual features. In recent years, research has also focused on domain-specific image retrieval systems such as fashion recommendation [43] [28] [36], product recommendation [26] [23] and discovery of food images [5]. Compared with existing commercial visual discovery engines, our system focuses more on *interactive* retrieval of objects within the images.

Representation learning: Over the past few years, convnet architectures such as AlexNet [24], GoogLeNet [39], VGG and ResNet have continuously pushed the state of the art on large-scale image classification challenges. Though trained for classification in one particular domain, the visual features extracted by these models perform well when transferred to other classification tasks [34], as well as related localization tasks like object detection and semantic segmentation [29]. This paper presents empirical evaluation of widely used classification-based image features in retrieval settings.¹

Detection: For web-scale services, the part-based model [11] approach was a well studied for detection, but recently deep learning methods have become more prevalent, with applications such as face detection [41], street number detection [15], and text detection [17]. Recent research focuses on application of detection architectures such as Faster R-CNN [32], YOLO [31] and the Single Shot Detector [27]. In this work we present, to the best of our knowledge, the first end-to-end object detection system for large-scale visual discovery systems.

¹ Recent work has also demonstrated the effectiveness of learning embeddings or distance functions directly from ranking labels such as relative comparisons [33] [20] and variations of pairwise comparisons [6] [38], or using Bilinear features [13] for fine-grained recognition. We will report our empirical findings in our future work.

Table 1: Precision@K on Pinterest evaluation dataset.

Model	Layer	Type	Dist.	P@1	P@5	P@10
AlexNet	fc6	raw	L2	0.093	0.045	0.027
AlexNet	fc6	raw	L1	0.099	0.047	0.027
AlexNet	fc7	raw	L2	0.104	0.052	0.031
AlexNet	fc7	raw	L1	0.106	0.052	0.031
AlexNet	fc8	raw	L2	0.088	0.047	0.028
AlexNet	fc8	raw	L1	0.090	0.046	0.028
GoogleNet	loss3/classifier	raw	L2	0.095	0.050	0.032
GoogleNet	loss3/classifier	raw	L1	0.098	0.050	0.032
VGG16	fc6	raw	L2	0.108	0.051	0.030
VGG16	fc6	raw	L1	0.118	0.057	0.035
VGG16	fc7	raw	L2	0.116	0.058	0.036
VGG16	fc7	raw	L1	0.113	0.060	0.038
VGG16	fc8	raw	L2	0.104	0.054	0.034
VGG16	fc8	raw	L1	0.106	0.054	0.034
ResNet101	pool5	raw	L2	0.160	0.080	0.050
ResNet101	pool5	raw	L1	0.149	0.073	0.045
ResNet101	fc1000	raw	L2	0.133	0.068	0.042
ResNet101	fc1000	raw	L1	0.139	0.067	0.041
ResNet152	pool5	raw	L2	0.170	0.083	0.050
ResNet152	pool5	raw	L1	0.152	0.077	0.047
ResNet152	fc1000	raw	L2	0.149	0.073	0.045
ResNet152	fc1000	raw	L1	0.148	0.073	0.044
AlexNet	fc6	binary	H.	0.129	0.065	0.039
AlexNet	fc7	binary	H.	0.110	0.054	0.033
AlexNet	fc8	binary	H.	0.089	0.046	0.027
VGG16	fc6	binary	H.	0.158	0.081	0.049
VGG16	fc7	binary	H.	0.133	0.068	0.044
VGG16	fc8	binary	H.	0.110	0.055	0.035
ResNet101	fc1000	binary	H.	0.125	0.062	0.039
ResNet101	pool5	binary	H.	0.055	0.025	0.014
ResNet152	fc1000	binary	H.	0.133	0.065	0.041
ResNet152	pool5	binary	H.	0.057	0.026	0.015
VGG16 (Pin.)	fc6	binary	H.	0.169	0.089	0.056

4. FEATURE REPRESENTATION

We adopted and evaluated several popular classification models such as AlexNet[24], GoogLeNet [40], VGG16 [37], and variants ResNet101 and ResNet152[16]. In addition to the raw features, we also explored binarized [4] representations of these features which are of interest to us due to their much smaller memory footprint. We also compare the Euclidean (L2) and Manhattan (L1) distance metrics for the raw features. For these features, training and inference are done through the open-source Caffe [18] framework on multi-GPU machines. We are separately investigating on-device visual discovery experience to augment the server-based approach using the mobile implementation of TensorFlow [3].

Beyond the base models trained for ImageNet classification [9], we also fine-tune on our own Pinterest training dataset generated in a similar fashion as the evaluation dataset while ensuring no overlap. This is because Pinterest images, many of which are high-quality photography and professional product images, have different statistics than ImageNet images, many of which are personal photos focused on individual generic objects. Models are fine-tuned by replacing the softmax classification layer of a pre-trained base model with a new classification layer, trained to classify Pinterest images, initialized using the same intermediate and lower level weights.

The training dataset consists of millions of images distributed over roughly 20,000 classes while the evaluation dataset, visualized in Figure 6, consists of 100,000 images distributed over the same classes. Both datasets were collected from a corpus of Pinterest images that are labeled with annotations. We limit the set of annotations by taking the top 100,000 text search queries on Pinterest and ran-




Figure 5: VGG16 fc6 layer raw features (left) vs binary features (right). PASCAL VOC 2011 classification images are shown as colored dots. White noise images are shown as black dots. One can see that binarization separates out the noise from the Pascal images while retaining the label clusters.

Table 2: Applying binary transformation to the VGG16 fc6 feature improves precision@k on PASCAL VOC 2011 images

Type	Dist	P@1	P@5	P@10
raw	L2	0.588	0.544	0.506
raw	L1	0.658	0.599	0.566
binary	H.	0.752	0.718	0.703

domly sampling 20,000 of them. After filtering, we retain a corpus of images containing only annotations matching the sampled queries and randomly sample from this corpus to create balanced training and evaluation datasets. The query annotation used to generate the image becomes its class label.

To evaluate our visual models, from our evaluation dataset we use 2,000 images as query images while the rest are indexed by our visual search system using the image representation being evaluated. We then retrieve a list of results for each query image, sorted by distance. A visual search result is assumed to be relevant to a query image if the two images share the same class label, an approach that is commonly used for offline evaluation of visual search systems [30] in addition to human evaluation. From the list of results with the associated class labels, we compute Precision @ K metrics to evaluate the performance of the visual models.

The results of these experiments for P@1, P@5, and P@10 performance are shown in Table 1. We see that when using raw features, intermediate features (pool5, fc6, fc7) result in better retrieval performance than more semantic features (fc8, fc1000). Among the raw feature results, ResNet152 pool5 features using L2 distance perform the best. For scalability, however, we are most interested in performance using binary features as our system must scale to billions of images. When binarized, VGG16 fc6 features perform best among the features we tested. These binarized fc6 features are 16x smaller than the raw pool5 features, with a trade-off in performance. Fine-tuning the VGG16 weights on Pinterest data, however, makes up for this performance difference.

We conjecture that the drop in performance of the ResNet pool5 feature when binarized is due to how the pool5 features are average pooled from ReLU activations. ReLU activations are either zero or positive and averaging these activation will bias towards positive pool5 features that are each binarized to one. We see that for ResNet101 pool5 fea-

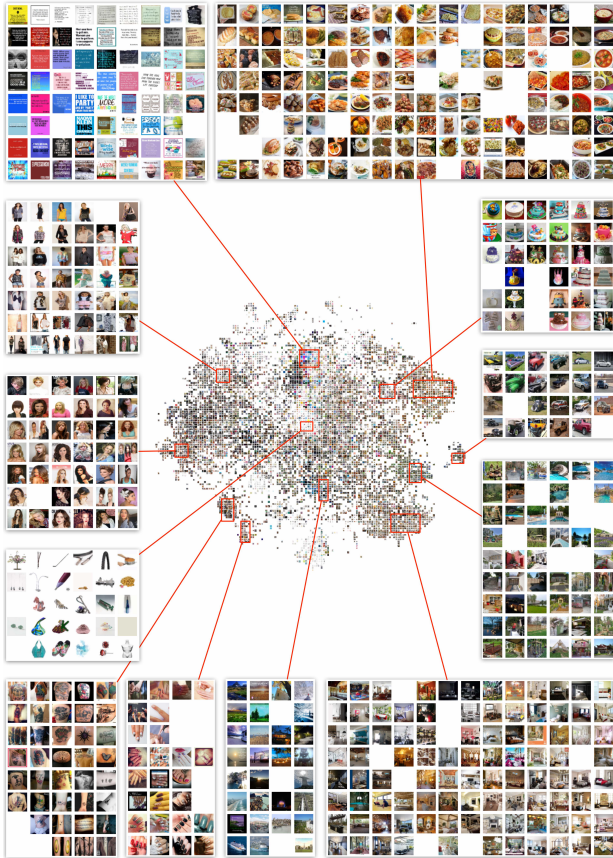


Figure 6: Visualization of binarized embeddings of Pinterest images extracted from fine-tuned VGG16 FC6 layer.

tures on our evaluation dataset, on average 83.86% of the features are positive. Additionally, we speculate that intermediate features of AlexNet and VGG16 are better suited for the current binarization scheme. By training the intermediate features with a ReLU layer, the network will learn features that ignore the magnitude of negative activations. We see that this is also exactly what the binarization post-processing does for negative features.

One interesting observation is that binarization improves the retrieval performance of AlexNet and VGG16. We reproduce this behavior in Table 2 on the Pascal VOC 2011 [10] classification dataset. In this experiment, we built our index using fc6 features extracted from VGG16 from the train and validation images of the Pascal classification task. Each image is labeled with a single class with images that have multiple conflicting classes removed. We randomly pick 25 images per class to be used as query images and measure precision@K with $K = 1, 5, 10$. With the Pascal dataset, we can also qualitatively see that one advantage of binarization is to create features that are more robust to noise. We see this effect in Figure 5 where binarization is able to cleanly separate noisy images from real images while the raw features cluster both together.

5. OBJECT DETECTION

One feature that is particularly relevant to Pinterest is the presence of certain object classes (such as shoes, chairs, tables, bags, etc.). Extracting objects from images not only al-

Table 3: Object detection performance.

	Faster R-CNN		SSD	
	precision	recall	precision	recall
Fashion	0.449	0.474	0.473	0.387
Home decor	0.413	0.466	0.515	0.360
Vehicles	0.676	0.625	0.775	0.775
Overall	0.426	0.470	0.502	0.371
Latency	272 ms		59 ms	

lows us to build novel discovery experiences (e.g. Section 8), but also improves user engagement metrics as detailed in Sections 6 and 7. This section covers our iterations of object detection at Pinterest, starting from deformable parts models described in our prior work [21] to Faster R-CNN and Single Shot Detection (SSD) described below.

Faster R-CNN: One approach of interest for object detection is Faster R-CNN, given its state-of-the-art detection performance, reasonable latency for real-time applications [32], and favorable scalability with a high number of categories (since the vast majority of parameters are shared). We experimented with Faster R-CNN models using both VGG16 [37] and ResNet101 [16] as base architectures and the ResNet101 variant is currently one of the detection models, along with SSD, used in production at Pinterest.

When training Faster R-CNN models, we make use of a few differences from the original presentation of the method. First, we train our models using direct end-to-end optimization, rather than the alternating optimization presented in the original paper. Additionally, for the ResNet101 architecture of Faster R-CNN, we found that significant computation savings could be achieved by decreasing the number of proposals considered from 300 to 100 (or dynamically thresholded based on the objectness score of the ROI) during inference without impacting precision or recall. The shared base convolution features are pre-trained on a fine-tuned Pinterest classification task, and then trained for 300k iterations using an internal Pinterest detection dataset.

Single Shot Detection: Although Faster R-CNN is considerably faster than previous methods while achieving state-of-the-art precision, we found that the latency and GPU memory requirements were still limiting. For example, the productionized implementation of Faster R-CNN used in Flashlight, as described in Section 7, relies on aggressive caching, coupled with dark reads during model swaps for cache warming. We also index objects into our visual search system to power core functionalities of Lens, as described in Section 8, which requires running our detector over billions of images. Any speedups we can get in such use cases can lead to drastic cost savings. As a result, we have also experimented with a variant of the Single Shot Multibox Detector (SSD) [27], a method known to give strong detection results in a fraction of the time of Faster R-CNN-based models.

Because speed is one of our primary concerns, we perform detection with SSD at a relatively small resolution of 290×290 pixels. Our VGG-based architecture resembles that of the original authors very closely, so we refer the reader to the original work for more details. Here, we outline a few key differences between our model and that of the original authors.




Figure 7: Sample object detection results from the Faster R-CNN model (labeled as red) and SSD model (labeled as blue)

First, we train our models with a 0.7 IoU threshold instead of a 0.5 threshold, to ensure that the resulting detections tightly bound the objects in question. The original SSD architecture uses additional convolutional layers with strides of 2 to detect larger objects—however, we found that this led to poor localization performance under our desired 0.7 IoU threshold. Instead, our model uses a stride of 1 for all additional convolutions, which greatly improves the quality of bounding box localization. Due to the size of our detection dataset (74,653 images), we also found that randomly sampling positive and negative anchors in a 1:3 ratio, as in the original paper, led to poor convergence. Instead, we make use of the recent Online Hard Example Mining (OHEM) method to sample anchors by always selecting the anchors that the models incurs the largest loss for [35]. However, despite our success with OHEM, we note that OHEM actually led to overfitting on a smaller detection dataset (19,994 images), indicating that OHEM is perhaps most useful on large and difficult datasets.

With these changes, we were able to train an SSD-based model that can perform object detection on 290×290 images in just 93 ms on an NVIDIA GRID K520 (an older generation GPU supported by Amazon Web Services), while achieving an F1 score on our internal evaluation dataset comparable to the previously described ResNet Faster R-CNN model.

Table 3 provides precision, recall, and latency figures on a Pinterest detection evaluation dataset, using our Faster R-CNN and SSD implementations. This dataset contains 79 categories, broadly categorized by *Fashion*, *Home Decor*, and *Vehicles*. This evaluation was performed at a 0.7 IoU threshold, and although the precision results look favorable for SSD, we did observe, qualitatively, that the localization quality was worse (one reason being that SSD uses a smaller, warped input size). Nevertheless, SSD, owing to its simpler and more streamlined architecture, and superior latency, warrants close investigation as the favored object detection model for production applications.

6. ENHANCING RELATED PINS

Related Pins is a pin recommendation system that leverages the vast amount of human-curated content on Pinterest to provide personalized recommendations of pins based on a given query pin. It is most heavily used on the pin closeup view shown in Figure 8, which is known as the Related Pins feed. Additionally, Related Pins recommendations have been incorporated into several other parts of Pinterest, including the home feed, pin pages for unauthen-

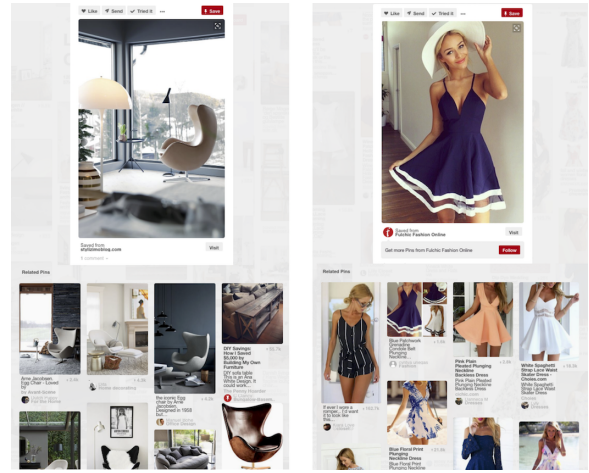


Figure 8: Related Pins is an item-to-item recommendation system. The results are displayed below the currently viewing image.

ticated visitors, emails, and certain curated pin collections (such as the Explore tab).

User engagement on Pinterest is defined by the following actions. A user closeups on a pin by clicking to see more details about the pin. The user can then click to visit the associated Web link; if they remain on-site for an extended period of time, it is considered a long click. Finally, the user can save pins onto their own boards. We are interested in driving “Related Pins Save Propensity” which is defined as the number of users who have saved a Related Pins recommended pin divided by the number of users who have seen a Related Pins recommended pin. Liu et al. [25] presented a more detailed architecture overview and evolution of the related pins feature. This section focuses on how convnet features and object detection can improve the engagement of Related Pins.

Convnet features for recommendations

Related Pins are powered through collaborative filtering (via image-board co-occurrences) for candidate generation and the commonly used Rank-SVM to learn a function where various input features are used to generate a score for re-ranking. To conduct the experiment, we set up a series of A/B experiments, where we selected five million popular Pins on Pinterest as queries, and re-ranked their recommendations using different sets of features. The control group re-ranked Related Pins using a linear model with the existing set of features. The treatment group re-ranked using fine-tuned VGG *fc6* and *fc8* visual similarity features along with indicator variables (in addition to the features used in control).

Across the 5M query Pins, the treatment saw a 3.2% increase in engagement (click and repin)². After expanding the treatment to 100M query Pins, we observed a net gain of 4.0% in engagement with Related Pins, and subsequently launched this model into production. Similar experiments with a fine-tuned AlexNet model yielded worse results (only

²This metric was measured across a 14 day period in Sep. 2015.

0.8% engagement gain) as expected from our offline evaluation of our visual feature representation.

When broken down by category, we noted that the engagement gain was stronger in predominantly visual categories, such as art (8.8%), tattoos (8.0%), illustrations (7.9%), and design (7.7%), and lower in categories which primarily rely on text, such as quotes (2.0%) and fitness planning (0.2%). Given the difference in performance among categories, we performed a follow-up experiment where we introduced a cross feature between the category vector of the query and the scalar *fc6* visual similarity feature (between the query and candidate) to capture the dependency between category and usefulness of the visual features in our linear model. This introduces 32 new features to the model, one for each of our site-wide categories (these features are sparse, since the Pinterest category vector thresholds most values to zero). The result from this was a further 1.2% engagement increase in addition to the gains from the initial visual re-ranking model.

Object detection for recommendation

To validate the effectiveness of our detection system in production, one of our first experiments was to further improve the use of visual features in Related Pins (described in the previous section). Our primary observation was that users are sometimes interested in the *objects* in the Pin’s image, instead of the full image, and we therefore speculated that object detection could help compute *targeted* visual similarity features. For the below experiments, we focus on fashion queries as only our fine-tuned fashion Faster R-CNN model was available at the time of this experiment.

Table 4: Engagement results when adding cross features and object detection to visual similarity feature, measured over a 7 day period in Oct. 2015.

Features	Queries	Engagement
FT-VGG (control)	-	-
FT-VGG + category cross features	5M	+1.2%
FT-VGG + object detection <i>variant A</i>	315k fashion	+0.5%
FT-VGG + object detection <i>variant B</i>	315k fashion	+0.9%
FT-VGG + object detection <i>variant C</i>	315k fashion	+4.9%

After applying non-maximum suppression (NMS) to the proposals generated by our Faster R-CNN detector, we considered query Pins where the largest proposal either occupies at least 25% of the Pin’s image or if the confidence of the proposal passes a threshold of 0.9. We categorize these images as containing a *dominant visual object*, and using the best-performing fine-tuned VGG re-ranking variant from the previous section as our control, we experimented with the following treatments:

- *Variant A*: if a dominant visual object is detected in the query Pin, we compute visual features (VGG *fc6*) on just that object.
- *Variant B*: same as *variant A*, but we also hand-tune the ranking model by increasing the weight given to visual similarity by a factor of 5. The intuition behind this variant is that when a dominant visual object is present, visual similarity becomes more important for recommendation quality.
- *Variant C*: Both control and treatment have the same visual features. If the query contains a dominant vi-

sual object however, we increase the weight of the visual similarity features by a factor of 5, as in *variant B*. In this variant, we assume that the presence of detected visual objects such as bags or shoes indicates that visual similarity is more important for this query.

Results for these variants are listed in Table 4. Variants A and B of the object detection experiments show that directly computing visual features from object bounding boxes slightly increases engagement. One suspicion for the lack of more significant performance increases may be due to our detector returning tight bounding boxes that do not provide enough context for our image representation models as we are matching a query object to whole image candidates. Variant C shows a more significant 4.9% engagement gain over the VGG similarity feature control, demonstrating that the presence of visual objects alone indicates that visual similarity should be weighed more heavily.

Our object detection system underwent significant improvements following the promising results of this initial experiment, and was applied to substantially improve engagement on Flashlight, as described in the next section.

7. PINTEREST FLASHLIGHT

Flashlight, as shown in Figure 9 is a visual search tool that lets users search for object within any image on Pinterest. It supports various applications including Shop-the-Look³, which launched in 2017 and shown in Figure 10.

The input to Flashlight is an object proposal, which is generated either (1) using the detection features described in Section 5 or (2) directly from our users via the flexible cropping mechanism as shown in Figure 9. Visual search on the object is then powered by the retrieval features described in Section 4 for candidate generation with light weight reranking using the full image metadata as context. Flashlight returns both image results and clickable tf-idf weighted annotations to allow users to narrow down on results. These annotations are generated by aggregating the annotations of the image results.

We implemented Flashlight as an iteration of Similar Looks, a fashion specific visual search tool described in our prior work [21]. We learned from Similar Looks that being able to find similar results of objects within an image can increase engagement. However, after running the experiment, we saw challenges in launching the Similar Looks product into wider deployment. First, due to the offline nature of our older parts-based detection pipeline, the coverage of objects was too low. Only 2.4% of daily active users saw a detected object and user research revealed that inconsistency with how some images had objects while other images did not confused users, who expected interactive dots on every image, including new content.

The cropping mechanism was our solution to this problem, giving users the flexibility to manually select any object in any image and get real-time visually similar results for the selected object. By not restricting what a user can crop, any object can be searched. A few months after the initial launch, we introduced object detection to Flashlight, generating clickable dots to simplify the user interface. At the time of the publication, more than 1 billion instances of objects were detected and shown to the user.

³Shop-the-Look combines visual search results with expert curation to improve the accuracy of the results.

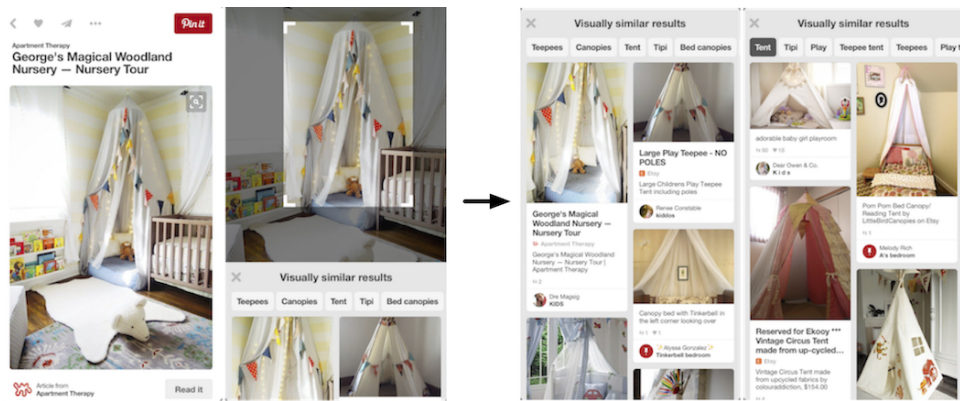


Figure 9: Pinterest Flashlight supports interactive retrieval of objects in the images through the use of “cropper” tool (left) and query-refinement suggestions (top right).

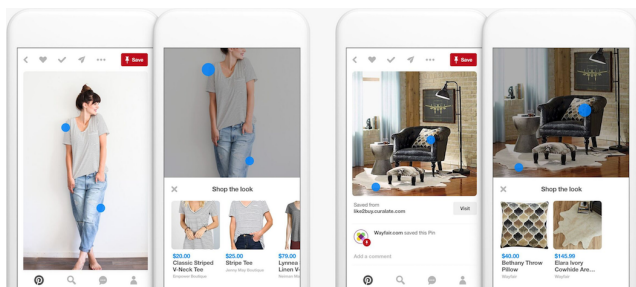


Figure 10: Shop-the-Look uses a variant of Flashlight returning only buyable content along with manual tagging to achieve a high quality visual shopping experience.

Convnet Features for Retrieval

The initial Flashlight launch did not include automatic object detection. We introduced a semi-transparent search icon on the top right of a Pin image that a users can press to use the cropping mechanism. As the user manipulates the crop bounding box, real time visual search results and annotations are generated.

Initially after the launch we saw that 2% of users who viewed a pin closeup will press the search button, which corresponds to roughly 4 million visual search request per day. Because search is done on crops of an image, when retrieving visually similar results, the re-ranking function normally applied after the candidate retrieval step is very minimal⁴ as we have no metadata for the crop. Therefore, the visual search results returned are mostly determined by the initial nearest neighbor retrieval on our deep learning embeddings. With the context that it is mostly deep learning that is powering this new feature, we are able to achieve 2/3 the engagement rate of text search, an existing product whose quality has been in development for many years.

Object Detection for Retrieval

One crucial improvement to Flashlight that we implemented is real-time object detection using Faster R-CNN, which of-

⁴Some components of the re-ranking function such as near-duplicate image scoring still apply

Table 5: Object Detection in Flashlight (Faster R-CNN ResNet101), detection threshold 0.7

Visual Thresh.	Annot. Thresh.	Categ. Conformity	Engagement Gain
0.0	0	0.0	-3.8%
1.0	0	0.0	-5.3%
0.0	1000	0.0	-5.2%
1.0	1000	0.0	+2.1%
1.0	1000	0.8	+4.9%

fers several advantages to the product. First, we’re able to simplify the user interface of Flashlight by placing predetermined *object dots* for categories we detect so users could click on the dot instead of manually moving a cropping box. Second, a useful goal for any retrieval system is to construct an aggregated user engagement signal, one use case being to boost up known highly engaged results for a query. This aggregation was infeasible in our initial version of Flashlight as most queries are unique due to the use of a manual cropper.

Though intuitively, object detection seemed to be a simple engagement improvement to Flashlight as users would have a much easier interface to use, we learned otherwise through our A/B experiments. Upon initial launch of the object detection experiment, where the control displayed a default bounding box, and treatment displayed a clickable dot over each detected object (as in Figure 2) we found that engagement metrics decreased (specifically, we were interested in “Flashlight Save Propensity” similar to the Related Pins metric used previously). After investigation into the reason for the poor performance, we saw two significant issues, as shown in Figure 11: (a) bad object detections (manifested as either low localization quality or object “hallucination”, left example), and (b) irrelevant retrieval results, which could either be a symptom of the first failure case (right example), or a standalone issue (middle example).

To address these problems, we developed a method of ranking object detections with the confidence of the visual search system, as we found that detector confidence alone is not sufficient to suppress poor detections. In particular we looked into three signals returned by our visual search system: visual Hamming distance of the top result (from

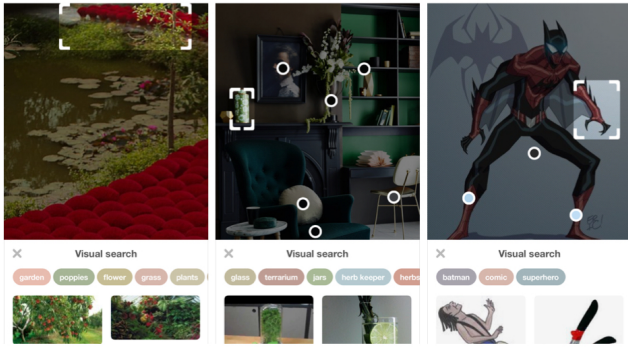


Figure 11: Examples of object detection false positives (shown as selected object), which are successfully suppressed using the category conformity suppression method.




Figure 12: Pinterest Lens: user can take a picture on their phone camera to get objects and ideas related to the photo (e.g. a strawberry image can lead to chocolate strawberry recipes).

our 4096 dimensional binarized convnet features), top annotation score for annotations returned by our visual search system (aggregated tf-idf scored annotations from the visual search results), and category conformity (maximum portion of visual search results that are labeled with the same category). We list our results in Table 5, where we impose a minimum threshold on each signal. Our best variation results in a 4.9% increase in user engagement on Flashlight. We found that category conformity in particular was critical in our improved engagement with Flashlight when object detection was enabled. A low category conformity score indicates irrelevant visual search results which we use as a proxy to suppress both of the error types shown in Figure 11.

8. PINTEREST LENS

Pinterest Lens, as shown in Figure 12, is a new discovery experience accessed via a mobile phone camera. Unlike Flashlight, Lens is not optimized to return visually similar results, but instead is implemented to return a diverse set of engaging results semantically relevant to the query. For example, a photo of blueberries would yield not only visually similar blueberries, the results may also include recipes for various types of blueberry desserts. In this section, we describe the technology used in Lens, and will report metrics in a future publication, as the product launch was very recent at the time of this writing.

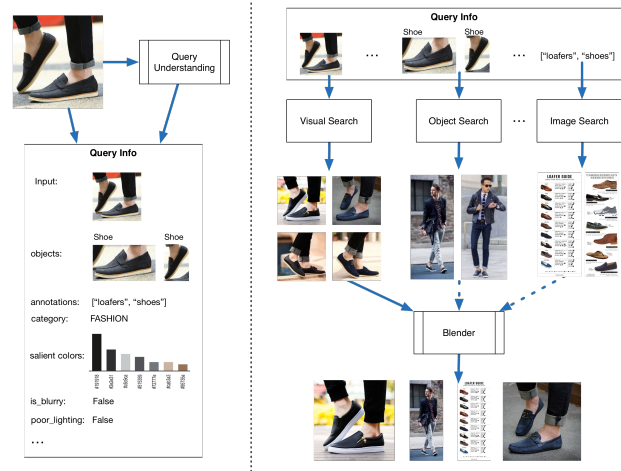


Figure 13: Lens is divided into two components. The query understanding layer (left) first computes visual features. The blending layer (right) then fetches results from multiple content sources.

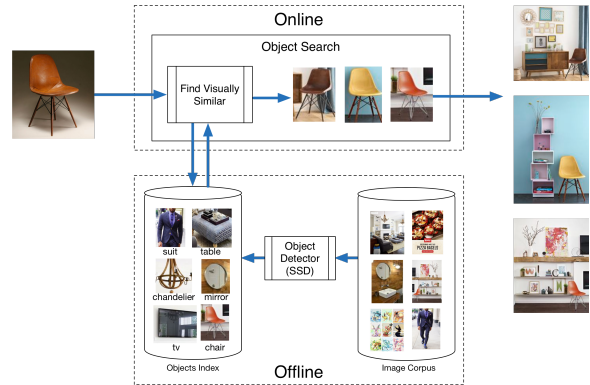


Figure 14: Given a query object, we find visually similar objects. We return the whole images (scenes) corresponding to the objects.

The overall Lens architecture is separated into two logical components as shown in Figure 13. The first component is the query understanding layer where we returned a variety of visual and semantic features for the input image such as annotations, objects, and salient colors. The second component is result blending as Lens results come from multiple sources. Visually similar results come from Flashlight, semantically similar results are computed using the annotations to call Pinterest image search, and contextual results (e.g. living room designs containing the chair in the photo taken by the user's camera) come from our object search system explained below. Not all sources are triggered per request. The blender will dynamically change blending ratios and content sources based on our derived features from the query understanding layer. For instance, image search would not be triggered if our annotations are low confidence.

Object Search

Object search is a visual search system where instead of indexing only whole images as per traditional systems, we also

index objects. One use case of such a system is to retrieve results which *contain* a queried object. For example, in Figure 3, if a user takes a picture of a Roman numeral clock, they may be interested in furniture which complements this clock or living room designs containing this clock.

To build this system, we use our SSD object detector, as described in Section 5, to extract objects from billions of images on Pinterest for indexing. This results in a corpus with more than a billion objects. SSD is also run during query time on the given input image as object search is only triggered if the query image contains an object. Given a query object, we compute visually similar objects and then map these objects back to their whole images (scenes) to return to the user. Figure 14 describes an end-to-end view of the pipeline. The idea of matching query with objects was previously explored in [21] [6]; to the best of our knowledge this is the first production system to serve this class of recommendations.

9. CONCLUSIONS

This paper presents an overview of our visual discovery engine powering various visual discovery experiences at Pinterest, and shares the rationales behind our technical and product decisions such as the use of binarized features, object detection, and interactive user interfaces. By sharing our experiences, we hope visual search becomes more widely incorporated into today’s commercial applications.

10. ACKNOWLEDGEMENTS

Visual discovery is a collaborative effort at Pinterest. We’d like to thank Maesen Churchill, Jamie Favazza, Naveen Gavini, Yiming Jen, Eric Kim, David Liu, Vishwa Patel, Albert Pereta, Steven Ramkumar, Mike Repass, Eric Sung, Sarah Tavel, Michelle Vu, Kelei Xu, Jiajing Xu, Zhefei Yu, Cindy Zhang, and Zhiyuan Zhang for the collaboration on the product launch, Hui Xu, Vanja Josifovski and Evan Sharp for the support.

Additionally, we’d like to thank David Tsai, Stephen Holiday, Zhipeng Wu and Yupeng Liao for contributing to the visual search architecture at VisualGraph prior to the acquisition by Pinterest.

11. REFERENCES

- [1] Introducing the future of visual discovery on pinterest. <https://engineering.pinterest.com/blog/introducing-future-visual-discovery-pinterest>. Published: 2017-02-08.
- [2] Our crazy fun new visual search tool. <https://blog.pinterest.com/en/our-crazy-fun-new-visual-search-tool>. Published: 2015-11-08.
- [3] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. J. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Józefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. G. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. A. Tucker, V. Vanhoucke, V. Vasudevan, F. B. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *CoRR*, abs/1603.04467, 2016.
- [4] P. Agrawal, R. Girshick, and J. Malik. Analyzing the performance of multilayer neural networks for object recognition. 2014.
- [5] K. Aizawa and M. Ogawa. Foodlog: Multimedia tool for healthcare applications. *IEEE MultiMedia*, 22(2):4–8, 2015.
- [6] S. Bell and K. Bala. Learning visual similarity for product design with convolutional neural networks. *ACM Trans. Graph.*, 34(4):98:1–98:10, July 2015.
- [7] L. Bertelli, T. Yu, D. Vu, and B. Gokturk. Kernelized structural svm learning for supervised object segmentation. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 2153–2160. IEEE, 2011.
- [8] R. Datta, D. Joshi, J. Li, and J. Wang. Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Survey*, 40(2):5:1–5:60, May 2008.
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [10] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2011 (VOC2011) Results. <http://www.pascal-network.org/challenges/VOC/voc2011/workshop/index.html>.
- [11] P. F. Felzenszwalb, R. B. Girshick, and D. A. McAllester. Cascade object detection with deformable part models. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2241–2248, 2010.
- [12] A. Frome, Y. Singer, and J. Malik. Image Retrieval and Classification Using Local Distance Functions. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 417–424. MIT Press, Cambridge, MA, 2007.
- [13] Y. Gao, O. Beijbom, N. Zhang, and T. Darrell. Compact bilinear pooling. *arXiv preprint arXiv:1511.06062*, 2015.
- [14] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. *arXiv preprint arXiv:1311.2524*, 2013.
- [15] I. J. Goodfellow, Y. Bulatov, J. Ibarz, S. Arnaud, and V. D. Shet. Multi-digit number recognition from street view imagery using deep convolutional neural networks. *CoRR*, abs/1312.6082, 2013.
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
- [17] M. Jaderberg, K. Simonyan, A. Vedaldi, and A. Zisserman. Reading text in the wild with convolutional neural networks. *CoRR*, abs/1412.1842, 2014.
- [18] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [19] Y. Jing and S. Baluja. Visualrank: Applying pagerank to large-scale image search. *IEEE Transactions on*

- Pattern Analysis and Machine Intelligence (T-PAMI)*, 30(11):1877–1890, 2008.
- [20] Y. Jing, M. Covell, D. Tsai, and J. M. Rehg. Learning query-specific distance functions for large-scale web image search. *IEEE Transactions on Multimedia*, 15:2022–2034, 2013.
- [21] Y. Jing, D. Liu, D. Kislyuk, A. Zhai, J. Xu, and J. Donahue. Visual search at pinterest. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining (SIGKDD)*.
- [22] Y. Jing, H. Rowley, J. Wang, D. Tsai, C. Rosenberg, and M. Covell. Google image swirl: a large-scale content-based image visualization system. In *Proceedings of the 21st International Conference on World Wide Web*, pages 539–540. ACM, 2012.
- [23] M. H. Kiapour, X. Han, S. Lazebnik, A. C. Berg, and T. L. Berg. Where to buy it: matching street clothing photos in online shops. In *International Conference on Computer Vision*, 2015.
- [24] A. Krizhevsky, S. Ilya, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1105. 2012.
- [25] D. C. Liu, S. Rogers, R. Shiau, K. Ma, Z. Zhong, D. Kislyuk, J. Liu, and Y. Jing. Related pins at pinterest, the evolution of a real-world recommender system. In *Proceedings of the International Conference on World Wide Web (WWW)*, 2017.
- [26] S. Liu, Z. Song, M. Wang, C. Xu, H. Lu, and S. Yan. Street-to-shop: Cross-scenario clothing retrieval via parts alignment and auxiliary set. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [27] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. Ssd: Single shot multibox detector. 2016. To appear.
- [28] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [29] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. *arXiv preprint arXiv:1411.4038*, 2014.
- [30] H. Müller, W. Müller, D. M. Squire, S. Marchand-Maillet, and T. Pun. Performance evaluation in content-based image retrieval: Overview and proposals. *Pattern Recognition Letter*, 22(5):593–601, 2001.
- [31] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [32] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Neural Information Processing Systems (NIPS)*, 2015.
- [33] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [34] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 806–813, 2014.
- [35] A. Shrivastava, A. Gupta, and R. Girshick. Training region-based object detectors with online hard example mining. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [36] E. Simo-Serra and H. Ishikawa. Fashion style in 128 floats: Joint ranking and classification using weak data for feature extraction. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [37] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [38] H. O. Song, Y. Xiang, S. Jegelka, and S. Savarese. Deep metric learning via lifted structured feature embedding. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [39] C. Szegedy, S. Ioffe, and V. Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. *CoRR*, abs/1602.07261, 2016.
- [40] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *arXiv preprint arXiv:1409.4842*, 2014.
- [41] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1701–1708, 2014.
- [42] J. Wang, Y. Song, T. Leung, C. Rosenberg, J. Wang, J. Philbin, B. Chen, and Y. Wu. Learning fine-grained image similarity with deep ranking. In *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition*, CVPR '14, pages 1386–1393, Washington, DC, USA, 2014. IEEE Computer Society.
- [43] K. Yamaguchi, M. H. Kiapour, L. E. Ortiz, and T. L. Berg. Retrieving similar styles to parse clothing. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(5):1028–1040, 2015.