Visual Search at Pinterest

Yushi Jing, David Liu, Andrew Zhai, Dmitry Kislyuk, Jiajing Xu, Jeff Donahue, Yunsong Guo, Sarah Tavel

{jing,dliu,andrew,dkislyuk,jiajing,jdonahue,yunsong,sarah}@pinterest.com

Introduction

Pinterest is a visual bookmarking tool that helps users discover and save creative ideas.

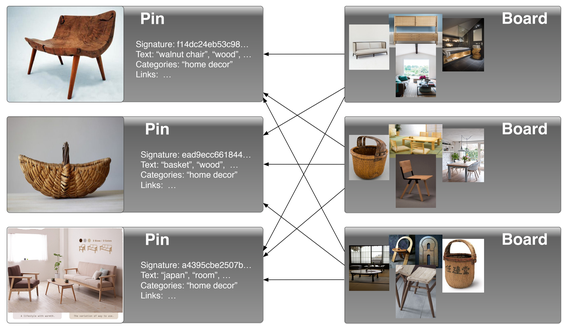

Figure 1: Users save pins to boards

Recommendations are generated using user-curated signals such as pin co-occurrences on boards, text signals, categories, and other metadata.

Figure 2: A pin (left) and its recommendations (right)

Problem: User-curated signals are limited in that unpopular or new Pins may not have these signals as we require user interaction. Instead of depending on user-curated signals, how can we leverage visual signals to immediately recommend Pins to users?

Visual Features

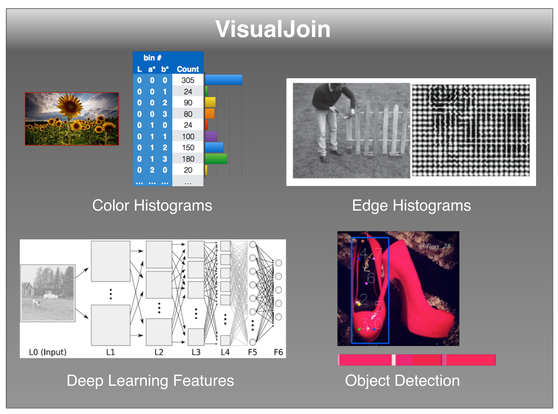

Figure 3: Visual features that we extract incrementally.

Incremental Feature Extraction

We built a feature extraction pipeline that efficiently maintains a complete set of image features as users add new images and as engineers add/modify features.

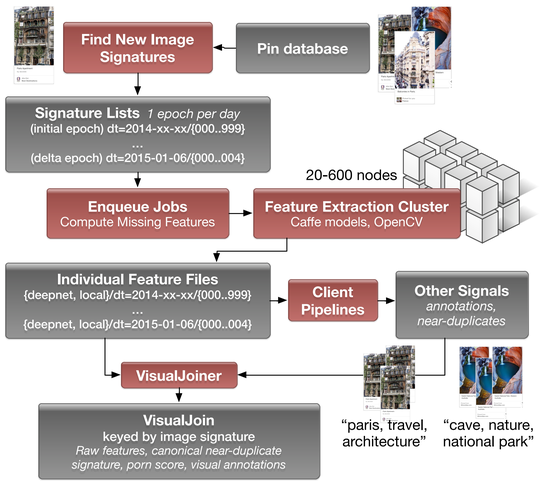

Figure 4: The process used to incrementally compute image features.

Peach System Overview

Peach is a large-scale distributed visual search system that provides real-time k-nearest neighbor lookup and reranking. Built upon open source tools and widely available platforms, such as Caffe, OpenCV, FLANN, Zookeeper, Thrift, and Amazon EC2.

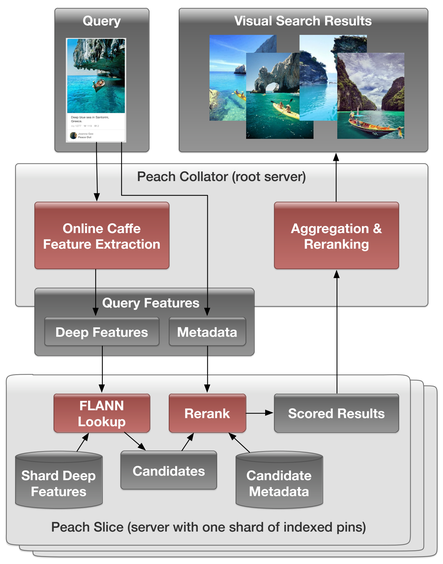

Figure 5: Illustration of how a query is processed in *Peach*.

Feature Evaluation

We evaluated visual search relevance with a dataset generated from text search results.

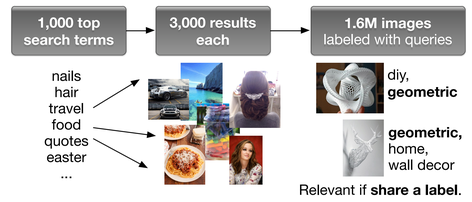

Table 1: Precision of top retrieval results using "Generic" (pre-trained for ILSVRC) and "Fine-Tuned" (on Pinterest annotation classification) models.

Model	p@5	p@10	latency (CPU)
Generic AlexNet FC6	0.051	0.040	193ms
Pinterest AlexNet FC6	0.234	0.210	234ms
Generic GoogLeNet	0.223	0.202	1207ms
Generic VGG-16	0.302	0.269	642ms

Application 1: Related Pins

We use Peach to generate recommendations for the 6% of traffic that have no recommendations, achieving a 2% increase in re-pin engagement on Related Pins.

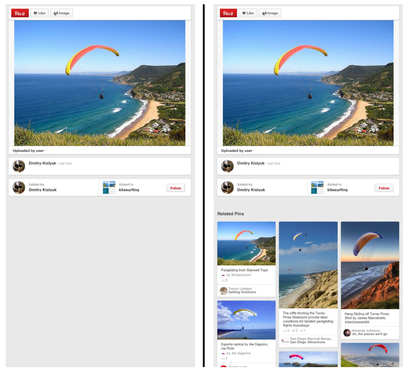

Figure 6: Incorporating Live Related Pins from Peach on a brand-new uploaded photo.

When we reranked *all* recommendations using visual feature similarity, we increased re-pin and click-through engagement on Related Pins by **10%**.

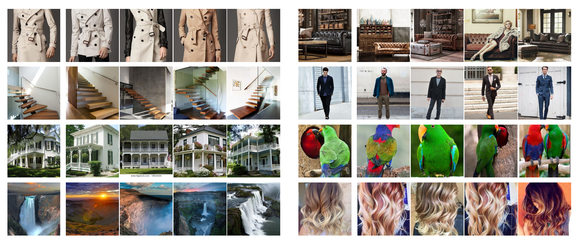

Figure 7: More visual search results used for Related Pins. In each row, the query image is shown, followed by the top results.

Application 2: Similar Looks

Using an offline object detection and localization pipeline, we built *Similar Looks*, which automatically tagged 80 million fashion objects on Pinterest.

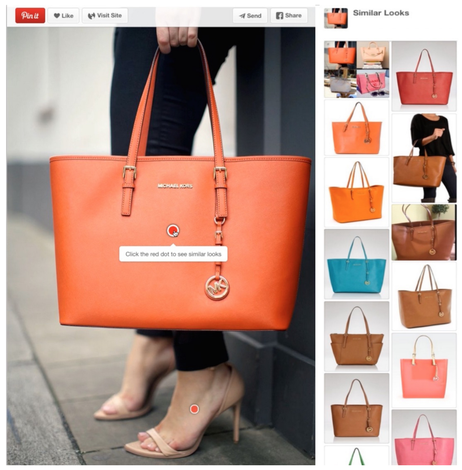

Figure 8: Similar Looks example: "object tags" are placed on each detected object; clicking a tag would show visual search results for that object.

Similar Looks uses a two-step pipeline: matching text annotations on images, followed by object detection using a fast implementation of deformable parts-based models. The text filtering dramatically reduced computational costs and lowered the false positive rate.

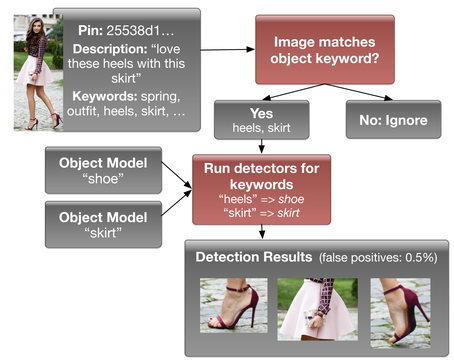

Figure 9: Two-step object detection dramatically reduces false positive rate.

When we launched the "red dot" UI shown in Figure 8, 12% of users who encountered a red dot clicked through, although overall engagement with Pins decreased. We chose to instead blend Similar Looks results into Related Pins, which increased re-pin engagement by 5%.

Visual Search at Pinterest

Yushi Jing, David Liu, Andrew Zhai, Dmitry Kislyuk, Jiajing Xu, Jeff Donahue, Yunsong Guo, Sarah Tavel

{jing,dliu,andrew,dkislyuk,jiajing,jdonahue,yunsong,sarah}@pinterest.com

Visual Search at Pinterest

Yushi Jing, David Liu, Andrew Zhai, Dmitry Kislyuk, Jiajing Xu, Jeff Donahue, Yunsong Guo, Sarah Tavel

{jing,dliu,andrew,dkislyuk,jiajing,jdonahue,yunsong,sarah}@pinterest.com