Dataset Fingerprints: Exploring Image Collections Through Data Mining

Konstantinos Rematas[1], Basura Fernando[1], Frank Dellaert[2] Tinne Tuytelaars[1]
[1]KU Leuven, ESAT-PSI, iMinds. [2]Georgia Tech.

Figure 1: Given an image collection, our system extracts patterns of discriminative mid level features and uses the connection between them to enable structure specific browsing.

As the amount of visual data increases, so does the need for summarization tools that can be used to explore large image collections and to quickly get familiar with their content. In this paper, we propose *dataset fingerprints*, a new and powerful method based on data mining that extracts meaningful patterns from a set of images. The discovered patterns are compositions of discriminative mid-level features that co-occur in several images. The main contributions of our work are:

- a fully unsupervised visual pattern discovery method,
- discovered patterns cover large parts of the images,often corresponding to full objects or meaningful parts,
- different patterns are connected based on co-occurrence, allowing a user to "browse" the images from one pattern to the next.

The traditional way of browsing such image datasets is based on manual labels (e.g. the "explore" button on the COCO web-page [3] and the "example images" in VOC [2]). In this paper we propose a novel way of summarizing and browsing an image collection, based on data mining. There are sets of images that *see* the same underlying structures, visual patterns and constellations. We express the problem of finding these structures as an *unsupervised* data mining problem: our image collection is the database, the images are the transactions and mid-level features are the items. We then look for *patterns*, i.e. combinations of mid-level features that appear in several images.

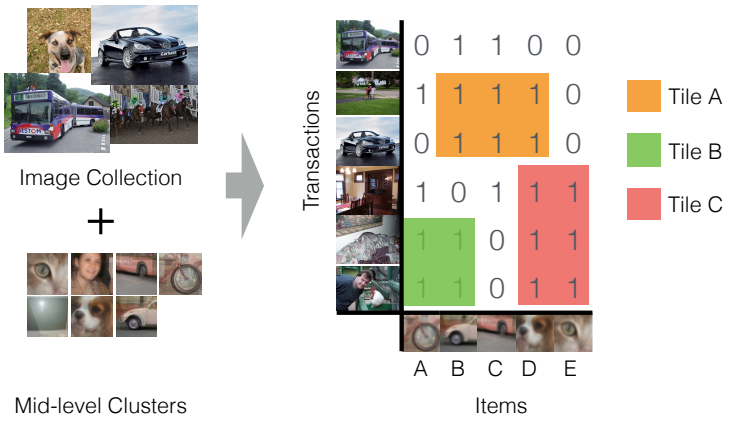

We start our dataset summarization by extracting random patches similar to [1, 4]. This set is used as seeds for the mid-level clusters. If a cluster is activated for a given image, then we add that item i to the transaction t of that image. We repeat this procedure over all images to create the set of transactions T . Afterwards, we use LCM [5] closed itemset mining algorithm to find an initial set of closed patterns (tiles).

This is an extended abstract. The full paper is available at the Computer Vision Foundation webpage.

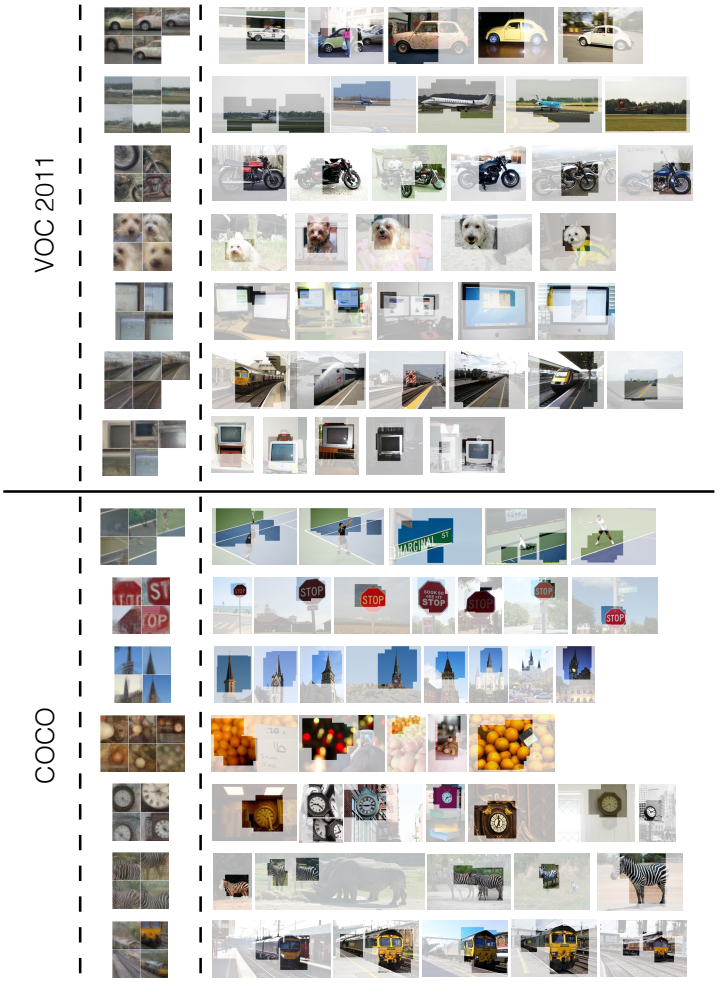

Our objective is to find the set of tiles that together cover the entire dataset with minimum redundancy. Ideally, we would like to cover the entire dataset with a few non-overlapping tiles. To do this we need to make sure that any of two tiles that are selected should not have common transactions *or* they should not have common items.

In conclusion, we propose *dataset fingerprints*, a new way of exploring image collections, based on interesting patterns discovered therein. The patterns can be linked together in a graph-structure, allowing a user to browse the images and gradually explore the archive.

[1] Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Context as supervisory signal: Discovering objects with predictable context. In *ECCV*, 2014.

[2] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010.

[3] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. *ECCV*, 2014.

[4] Saurabh Singh, Abhinav Gupta, and Alexei A Efros. Unsupervised discovery of mid-level discriminative patches. In *ECCV*. Springer, 2012.

[5] Takeaki Uno, Tatsuya Asai, Yuzo Uchida, and Hiroki Arimura. An efficient algorithm for enumerating closed patterns in transaction databases. In *Discovery Science*, 2004.

http://homes.esat.kuleuven.be/~krematas/DatasetFingerprints/