Scalable Triangulation-based Logo Detection

Yannis Kalantidis[1], Lluis Garcia Pueyo[2], Michele Trevisiol[3], Roelof van Zwol[2] and Yannis Avrithis[1]

[1]National Technical University of Athens, [2]Yahoo! Research Santa Clara, [3]Yahoo! Research Barcelona

1. Logo Detection in Natural Scenes

- **Problem:** Given an large annotated database of brand logos and one query image, detect if one or more of the logo instances appear in the query
- Broader category than near-duplicates (many different forms, or *variants*)

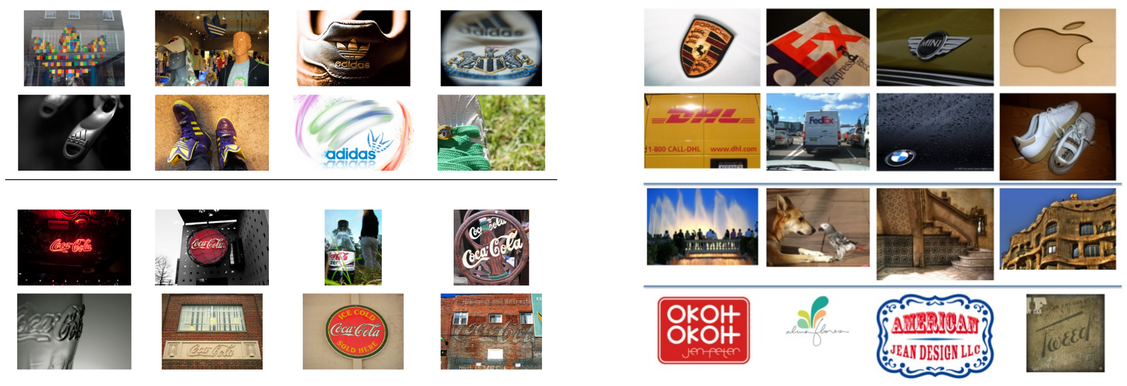

2. Novelty - Contribution

- Novel representation incorporating both visual appearance and local geometry
- *Scalability* in the number of classes: Querying a database of thousands of logo classes typically takes milliseconds
- Relatively few logo *instances* per class needed
- Highly discriminative signatures $\rightarrow$ sparse inverted index

3. Overview of the Approach

- Group *local features* into *triplets* by *multi-scale triangulation*
- Extract a binary *signature* from each triangle
- Represent each logo class by the *union* of signatures over all instances of the class (*generative* model)
- Index using a simple *inverted file* structure
- Extract signatures from the query image and rank classes according to the inverted index response

4. Multi-scale Delaunay Triangulation

Triangle with vertices of scale s_1

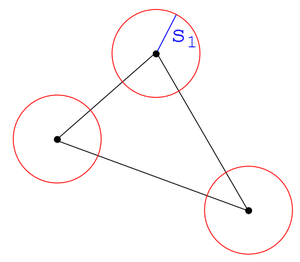

Outlier feature of scale s_2 does not affect triangulation

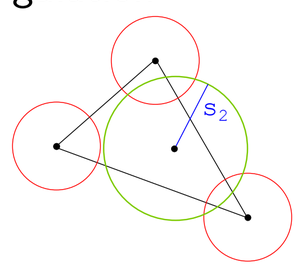

Triangle with vertices of scale s_3

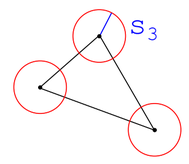

All local features — Multiple scale space levels

5. Triangle representation

- **Signature**: triple of the three visual word labels in lexicographically ascending order

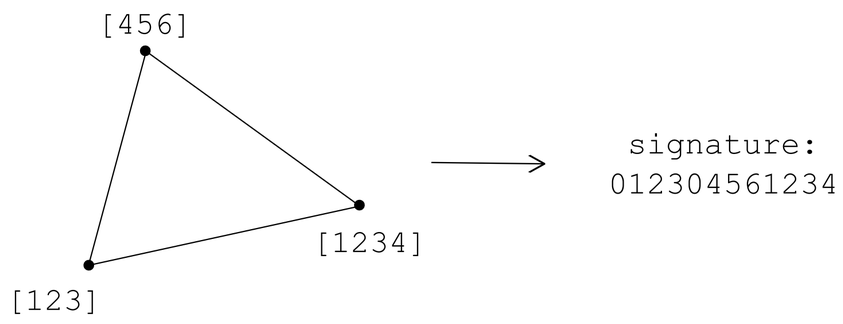

6. Class Models and Recognition

- Represent each logo class by the *union* of signatures of all triangles extracted from all instances of the class, along with their frequency of appearance (*generative* model)
- *Bag-of-signatures* model for each class
- Inverted index with *classes* as index atoms and *tf-idf* weighting

7. Dataset and Evaluation

- **Dataset:** 27 annotated brands (1080 images) and more than $4K$ distractor logo classes; publicly available at:

  `http://image.ntua.gr/iva/datasets/flickr_logos/`
- Accuracy against BoW for all $4K$ classes

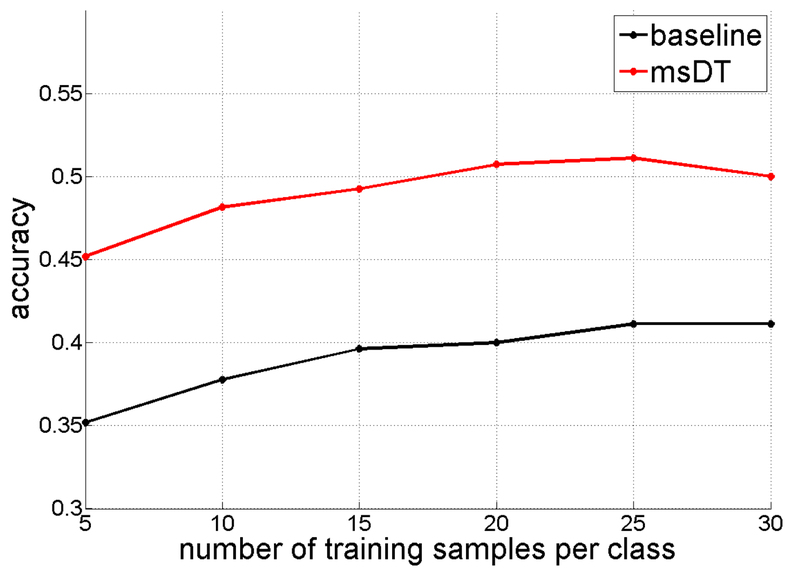

Acknowledgements

This work was supported by the European Commission under contract FP7-215453 WeKnowIt

Contact: {ykalant, iavr}@image.ntua.gr, {lluis, trevi, roelof}@yahoo-inc.com