

Image Matching and Visual Search

large scale methods and applications

Yannis Avrithis

National Technical University of Athens
Image, Video and Multimedia Systems Laboratory
Image and Video Analysis Team

Athens, April 2012



outline

- 1 local features and bag-of-words
- 2 local feature detection
- 3 visual vocabularies
- 4 spatial matching and re-ranking
- 5 geometry indexing
- 6 feature selection
- 7 clustering of photo collections
- 8 location and landmark recognition
- 9 implementation: ivl library


outline

- 1 local features and bag-of-words
- 2 local feature detection
- 3 visual vocabularies
- 4 spatial matching and re-ranking
- 5 geometry indexing
- 6 feature selection
- 7 clustering of photo collections
- 8 location and landmark recognition
- 9 implementation: ivl library

image matching




image matching



matching local feature points

[Scott and Longuet-Higgins, RSL 1991]



- given two sets of points $a_i, i = 1, \dots, m$ and $b_j, j = 1, \dots, n$ on the same plane, let d_{ij} be the distance between a_i and b_j
- following earlier theories of Ullman and Marr, the problem is to **associate** points a_i and b_j in a **one-to-one correspondence** such that the **sum of squared distances** between corresponding points is **minimized**

a spectral approach

- 1 construct the $m \times n$ proximity matrix G with elements

$$g_{ij} = \exp(-d_{ij}^2/2\sigma^2)$$

- 2 perform singular value decomposition of G

$$G = USV^T$$


where U, V are orthogonal matrices of dimension m, n and S is a non-negative diagonal $m \times n$ matrix

- 3 replace each diagonal element s_{ij} of S by 1 and reconstruct

$$P = UEV^T$$

- 4 finally, associate points a_i and b_j if element p_{ij} of P is the greatest element in its row and its column

a spectral approach




scale, translation

rotation


matching discriminative local features

[Lowe, ICCV 1999]




matching discriminative local features

[Lowe, ICCV 1999]

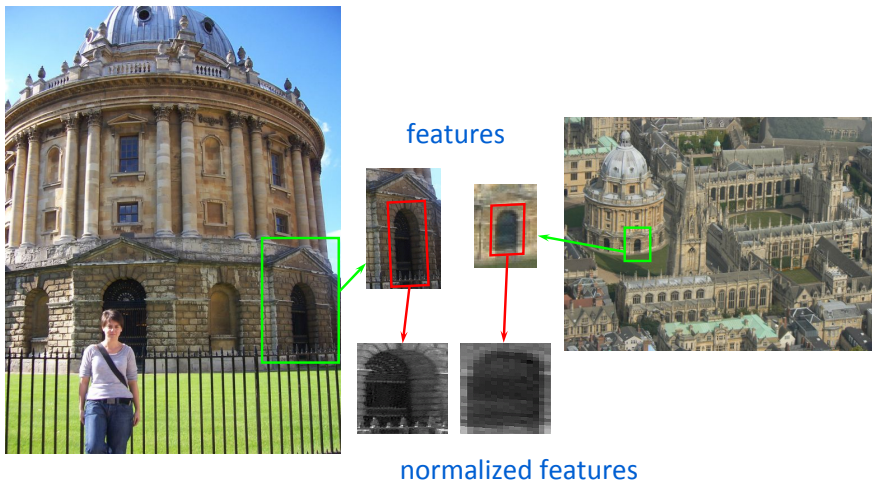


features



matching discriminative local features

[Lowe, ICCV 1999]




forget about geometry: bag-of-words

[Sivic and Zisserman, ICCV 2003]



vector quantization \rightarrow visual words




query




15

vector quantization \rightarrow visual words




query




15

vector quantization \rightarrow visual words




query


19




15




vector quantization \rightarrow visual words




query



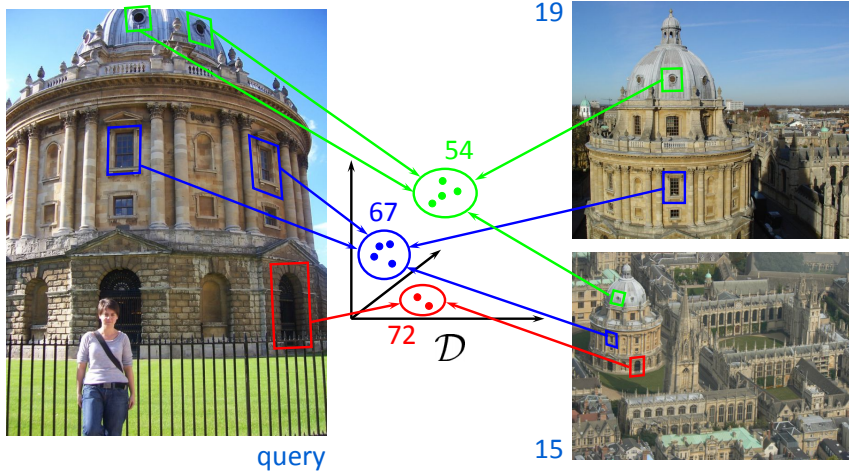
19



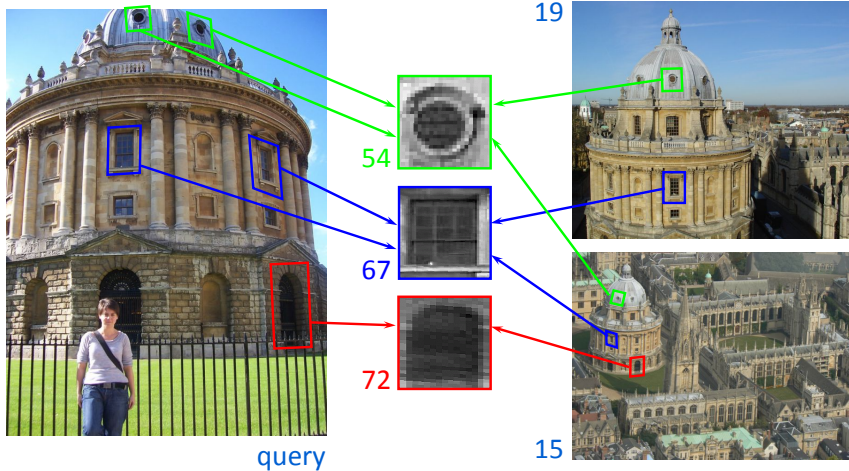
15



vector quantization \rightarrow visual words



vector quantization \rightarrow visual words



inverted file indexing

54	
67	
72	


query

54											
67											
72											


12	13	14	15	16	17	18	19	20	21	22	

images


inverted file indexing



inverted file indexing



inverted file indexing



inverted file indexing

54	
67	
72	

54											
67											
72											


query

ranked
shortlist

	1	3	1	2	1	1				
12	13	14	15	16	17	18	19	20	21	22



images

back to geometry: re-ranking




original images

back to geometry: re-ranking




local features

back to geometry: re-ranking



tentative correspondences


back to geometry: re-ranking



RANSAC inliers

RANSAC

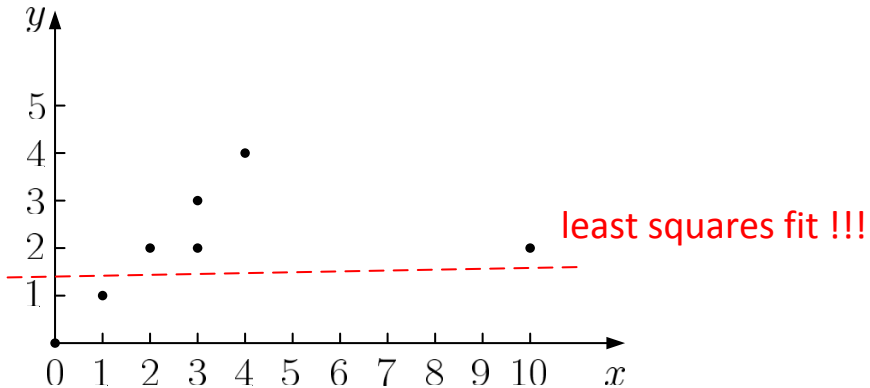
[Fischler and Bolles, CACM 1981]



problem: fit line to data


RANSAC

[Fischler and Bolles, CACM 1981]



RANSAC


[Fischler and Bolles, CACM 1981]



solution: choose 2 random points ...

RANSAC


[Fischler and Bolles, CACM 1981]



... fit line to them ...

RANSAC


[Fischler and Bolles, CACM 1981]



... classify remaining points to inliers ...

RANSAC


[Fischler and Bolles, CACM 1981]



... and outliers

RANSAC


[Fischler and Bolles, CACM 1981]



repeat ...

RANSAC


[Fischler and Bolles, CACM 1981]



... and repeat

RANSAC

[Fischler and Bolles, CACM 1981]



finally: maximum inliers

outline

- 1 local features and bag-of-words
- 2 local feature detection**
- 3 visual vocabularies
- 4 spatial matching and re-ranking
- 5 geometry indexing
- 6 feature selection
- 7 clustering of photo collections
- 8 location and landmark recognition
- 9 implementation: `ivl` library

edge-based feature detection


[Rapantzikos and Avrithis, ECCVW 2010]

- **blob-like regions** starting from single-scale edges
- local maxima of **Euclidean distance transform** expected to lie in region interior or close to ridges
- **greedily merge** maxima guided by edge strength, to reproduce the effect of smoothing in **scale-space evolution**
- regions of **arbitrary shape and scale**, unaffected by spurious or disconnected edges

original image




binary edge map



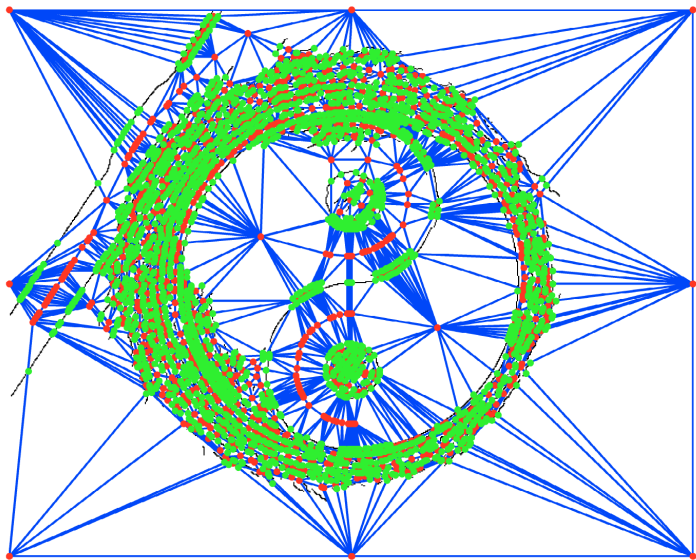
binary distance map




distance map + local maxima




Delaunay triangulation



convex hulls of selected regions



original image + features




weighted α -shapes


[Varytimidis et al., submitted to ECCV 2012]



input



triangulation



α -detector



MSER

medial features


[Avrithis and Rapantzikos, ICCV 2011]

- **additively weighted distance map** directly from image gradient, computed exactly in linear time


$$\mathcal{D}_d(f)(x) = \min_{y \in X} \{d(x, y) + f(y)\}, \quad x \in X$$

- **weighted medial** capturing region structure and topology
- region/boundary **duality** and image **partition**


region/boundary duality




distance propagation



negated distance & medial




partition




backpropagation


original image




weighted distance map and medial



region/boundary duality




original image + features



fragmentation factor



binary input



point labels





image partition


$$\phi(\kappa) = \frac{1}{a(\kappa)} \sum_{e \in E(\kappa)} w^2(x(e))$$

- simple selection criterion: is a region **well-enclosed by boundaries**?
- **arbitrary shape and scale**, without explicit scale-space construction

the challenge of shape




the challenge of shape




the challenge of scale




the challenge of scale




the challenge of scale



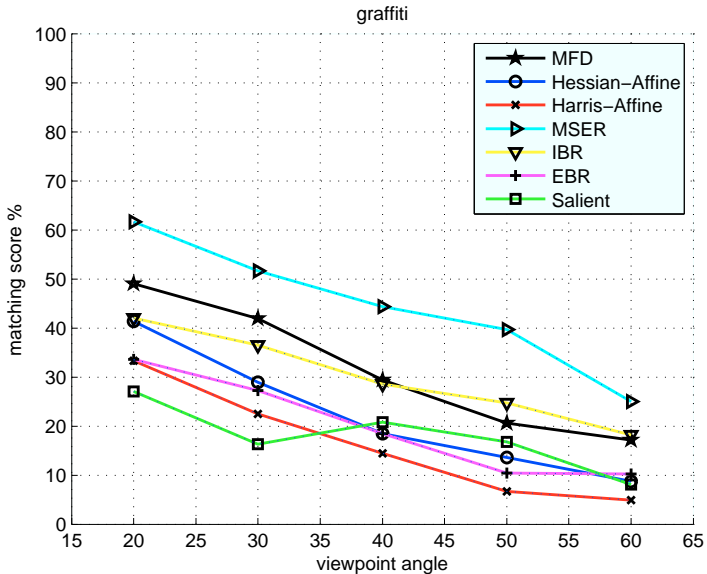
viewpoint: graffiti scene



viewpoint: graffiti scene



viewpoint: graffiti scene




scale + rotation: boat scene



scale + rotation: boat scene



scale + rotation: boat scene




blur: bikes scene



blur: bikes scene




blur: bikes scene




texture + blur: trees scene



texture + blur: trees scene



texture + blur: trees scene




viewpoint: wall scene



viewpoint: wall scene



viewpoint: wall scene



application to segmentation

[Avrithis and Leonardos, unpublished 2012]



application to segmentation

[Avrithis and Leonardos, unpublished 2012]

	BSDS500						
	Covering			PRI		VI	
	ODS	OIS	Best	ODS	OIS	ODS	OIS
Human	0.72	0.72	–	0.88	0.88	1.17	1.17
gPb-owt-ucm ₃	0.59	0.65	0.74	0.83	0.86	1.69	1.48
gPb-mad-ucm _g	0.58	0.64	0.74	0.83	0.86	1.62	1.39
gPb-mad-esm _c	0.55	0.62	0.71	0.82	0.86	1.83	1.51
Mean Shift [15]	0.54	0.58	0.66	0.79	0.81	1.85	1.64
Felz-Hutt [18]	0.52	0.57	0.69	0.80	0.82	2.21	1.87
gPb-mad-sfm	0.52	0.56	0.62	0.79	0.82	1.83	1.70
gPb-mad-esm _a	0.51	0.54	0.60	0.79	0.80	1.86	1.82
Canny-owt-ucm [3]	0.49	0.55	0.66	0.79	0.83	2.19	1.89
myCanny-mad-ucm _g	0.48	0.55	0.65	0.79	0.83	2.10	1.77
gPb-mad-ucm _φ	0.46	0.54	0.63	0.77	0.80	2.07	1.81
NCuts [16]	0.45	0.53	0.67	0.78	0.80	2.23	1.89
gCanny-mad-esm _c	0.45	0.53	0.63	0.78	0.83	2.31	1.91
gCanny-mad-sfm	0.42	0.49	0.58	0.77	0.80	2.18	1.95
gCanny-mad-esm _a	0.40	0.47	0.53	0.76	0.77	2.41	2.27
gCanny-mad-ucm _φ	0.35	0.42	0.50	0.74	0.77	2.43	2.29
Quad-Tree	0.32	0.37	0.46	0.73	0.74	2.46	2.32

outline

- 1 local features and bag-of-words
- 2 local feature detection
- 3 visual vocabularies**
- 4 spatial matching and re-ranking
- 5 geometry indexing
- 6 feature selection
- 7 clustering of photo collections
- 8 location and landmark recognition
- 9 implementation: `ivl` library

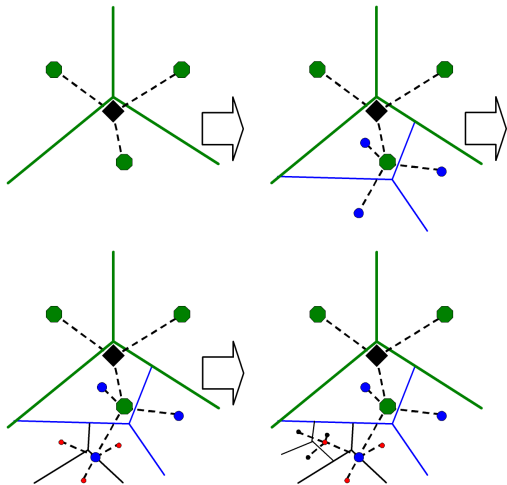
hierarchical k -means

[Nister and Stewenius, CVPR 2006]

- large scale clustering: $n = 10^7$ data points into $k = 10^6$ clusters, in $d = 10^2$ dimensions!
- complexity of k -means per iteration is $O(ndk)$: not practical!
- build a vocabulary tree by hierarchical k -means
- e.g. with a branching factor of $b = 10$, one needs only $l = \log_b k = 6$ levels, of complexity $O(ndb)$ each
- the same tree is used for nearest neighbor search and scoring

hierarchical k -means

[Nister and Stewenius, CVPR 2006]




approximate k -means

[Philbin et al., CVPR 2007]

- in k -means, most computation is spent on searching for **nearest neighbors** between points and cluster centers
- replace exact search by an **approximate nearest neighbor** (ANN) search, implemented by randomized k -d trees
- now a single level of complexity $O(ndt)$ is needed, where t is a **fixed** number of tests, e.g. $t = 100$
- **more flexible** than hierarchical k -means!


approximate k -means

[Philbin et al., CVPR 2007]




exact nearest neighbors: k -d tree

[Bentley, ACM 1975]



exact nearest neighbors: k -d tree

[Bentley, ACM 1975]



approximate NN: randomized k -d trees

[Silpa-Anan and Hartley, CVPR 2008]

indexing


- build m different k -d trees, each with a different structure
- use a **random** e.g. splitting plane, rotation, or projection for each tree

search


- **parallel search** among all m trees, with a limit of t nodes in total
- traverse all trees once, then use a **shared priority queue**

FLANN implementation


[Muja and Lowe, VISAPP 2009]




(a)



(b)



(c)



(d)

Gaussian mixtures

- each cluster j represented by **component** p_j with

$$p_j(\cdot) = \pi_j \mathcal{N}(\cdot | \boldsymbol{\mu}_j, \sigma_j \mathbf{I}),$$

modeling its population π_j , position $\boldsymbol{\mu}_j$ and scale σ_j

- **responsibility** of component p_j for data point \mathbf{x}_i

$$\gamma_{ij} = \frac{p_j(\mathbf{x}_i)}{\sum_{\ell} p_{\ell}(\mathbf{x}_i)}$$

- **maximum likelihood** estimates of parameters $\pi_j, \boldsymbol{\mu}_j, \sigma_j$ obtained as weighted averages over data, with responsibilities as weights
- iteratively compute responsibilities and parameters by **expectation maximization** (EM)

approximate Gaussian mixtures

[Avrithis and Kalantidis, submitted to ECCV 2012]

incremental search


- keep **all** t nearest neighbors found for each data point, not just the best
- use them **across** iterations, limiting the effort spent in new search
- limit responsibilities to this approximate nearest neighbor set: complexity is still $O(ndt)$

dynamic estimation of k

- start with **all data points** as components
- **purge** overlapping clusters and **expand** remaining ones at each iteration


approximate Gaussian mixtures—2d example

iteration=0, clusters=50




approximate Gaussian mixtures—2d example

iteration=1, clusters=15




approximate Gaussian mixtures—2d example


iteration=2, clusters=10




approximate Gaussian mixtures—2d example



approximate Gaussian mixtures—learning



approximate Gaussian mixtures—distractors




outline

- 1 local features and bag-of-words
- 2 local feature detection
- 3 visual vocabularies
- 4 spatial matching and re-ranking**
- 5 geometry indexing
- 6 feature selection
- 7 clustering of photo collections
- 8 location and landmark recognition
- 9 implementation: `ivl` library

local patches


- each local feature is associated with an image patch L , which also represents an affine transform
- the **rectified** patch \mathcal{R}_0 is transformed to the patch via L
- the patch is rectified back to \mathcal{R}_0 via L^{-1}



fast spatial matching (FSM)

[Philbin et al., CVPR 2007]

- single patch correspondence $L \leftrightarrow R$
- the transformation from one patch to the other is RL^{-1}
- each correspondence provides a transformation hypothesis
- hypotheses are now $O(n)$; we can try them all for inliers
- overall complexity is $O(n^2)$




relaxed spatial matching


[Tolias and Avrithis, ICCV 2011]

- do not seek for inliers
- rather, look for hypotheses that agree with each other
- how? build a hierarchical partition of 4d transformation space and count hypotheses that fall in the same bin
- inspired by Hough voting—hence Hough pyramid matching (HPM)
- for ℓ levels (e.g. $\ell = 5$), complexity drops from $O(n^2)$ to $O(n\ell)$!


toy example—Hough pyramid



Level 0












Level 1




Level 2

toy example—correspondences, strengths

	p	q	strength
c_1			$(2 + \frac{1}{2}2 + \frac{1}{4}2)w(c_1)$
c_2			$(2 + \frac{1}{2}2 + \frac{1}{4}2)w(c_2)$
c_3			$(2 + \frac{1}{2}2 + \frac{1}{4}2)w(c_3)$
c_4			$(1 + \frac{1}{2}3 + \frac{1}{4}2)w(c_4)$
c_5			$(1 + \frac{1}{2}3 + \frac{1}{4}2)w(c_5)$
c_6			0
c_7			0
c_8			$\frac{1}{4}6w(c_8)$
c_9			$\frac{1}{4}6w(c_9)$


toy example—affinity matrix



relaxed spatial matching ...


- is **invariant** to similarity transformations
- is **flexible**, allowing non-rigid motion and multiple matching surfaces or objects
- imposes **one-to-one** mapping

relaxed spatial matching—examples




fast spatial matching

relaxed spatial matching—examples



relaxed spatial matching

relaxed spatial matching—examples




relaxed spatial matching

relaxed spatial matching—examples




fast spatial matching

relaxed spatial matching—examples




relaxed spatial matching

relaxed spatial matching—examples



fast spatial matching

relaxed spatial matching—examples




relaxed spatial matching

relaxed spatial matching—examples



fast spatial matching

relaxed spatial matching—examples




relaxed spatial matching

world cities dataset

- 927 annotated images
- 17 groups of photos, each from a landmark scene in Barcelona
- 5 queries from each group


relaxed spatial matching—distractors



relaxed spatial matching ...

- is non-iterative, and **linear** in the number of correspondences
- in a given query time, can re-rank **one order of magnitude** more images than the state of the art
- needs **less than one millisecond** to match a pair of images, on average

relaxed spatial matching—timing



this work is becoming part of...



<http://opencv.willowgarage.com/>

outline

- 1 local features and bag-of-words
- 2 local feature detection
- 3 visual vocabularies
- 4 spatial matching and re-ranking
- 5 geometry indexing**
- 6 feature selection
- 7 clustering of photo collections
- 8 location and landmark recognition
- 9 implementation: ivl library






weak geometric consistency (WGC)

[Jegou et al., ECCV 2008]

- when an image undergoes rotation or scaling, the **orientation** and **scale** of local features is consistently modified
- quantize orientation and scale **differences** between feature pairs
- maintain **several scores** for each image, one for each difference bin
- this is not enough to recover a full transformation, but does improve ranking

weak geometric consistency (WGC)

[Jegou et al., ECCV 2008]




feature map hashing

[Avrithis et al., ACM-MM 2010]

- estimate image alignment via **single correspondence**
- for each feature, construct a **feature map** encoding normalized positions and appearance of all remaining features
- represent an image by a collection of such feature maps
- RANSAC-like matching is reduced to a number of **set intersections**


feature maps—example

- well aligned feature sets are likely to have maps with a high degree of overlap




feature maps—example

- well aligned feature sets are likely to have maps with a high degree of overlap



feature map similarity


$$S_F(P, Q) = \max_{v \in V(P, Q)} \max_{\substack{\hat{x} \in H_v(P) \\ \hat{y} \in H_v(Q)}} f_P^T(\hat{x}) f_Q(\hat{y})$$



feature map similarity

$$S_F(P, Q) = \max_{v \in V(P, Q)} \max_{\substack{\hat{x} \in H_v(P) \\ \hat{y} \in H_v(Q)}} f_P^T(\hat{x}) f_Q(\hat{y})$$

feature map of image P wrt origin \hat{x}




feature map similarity

$$S_F(P, Q) = \max_{v \in V(P, Q)} \max_{\substack{\hat{x} \in H_v(P) \\ \hat{y} \in H_v(Q)}} f_P^T(\hat{x}) f_Q(\hat{y})$$

feature map of image P wrt origin \hat{x}

feature map of image Q wrt origin \hat{y}




feature map similarity

over all origins mapped to visual word v

$$S_F(P, Q) = \max_{v \in V(P, Q)} \max_{\substack{\hat{x} \in H_v(P) \\ \hat{y} \in H_v(Q)}} f_P^T(\hat{x}) f_Q(\hat{y})$$

feature map of image P wrt origin \hat{x}

feature map of image Q wrt origin \hat{y}



feature map similarity


over all visual words that P, Q have in common

over all origins mapped to visual word v


$$S_F(P, Q) = \max_{v \in V(P, Q)} \max_{\substack{\hat{x} \in H_v(P) \\ \hat{y} \in H_v(Q)}} f_P^T(\hat{x}) f_Q(\hat{y})$$

feature map of image P wrt origin \hat{x}

feature map of image Q wrt origin \hat{y}




feature map similarity—example



fast spatial matching (35 inliers)

feature map similarity—example



feature map similarity (32 inliers)

towards indexing

with min-wise independent permutations [Broder, CCS 2000]

- FMS is a fast way of matching 2 images, but still not enough for indexing
- a feature map is an extremely **sparse histogram**; bin count typically takes values in $\{0, 1\}$
- each feature map f is represented by a **set** \bar{f} of non-empty bins
- then, use **min-wise independent permutations** a.k.a. **min-hashing** as an equivalent to **random sampling**

an example

[Chum et al. 2007]

<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	$\{a, b, c\}$	$\{b, c, d\}$	$\{a, e, f\}$
permutations						hash values		
3	6	2	5	4	1	2	2	1
1	2	6	3	5	4	1	2	1
3	2	1	6	4	5	1	1	3
4	3	5	6	1	2	3	3	1

an example

[Chum et al. 2007]

<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	$\{a, b, c\}$	$\{b, c, d\}$	$\{a, e, f\}$
permutations						hash values		
3	6	2	5	4	1	2	2	1
1	2	6	3	5	4	1	2	1
3	2	1	6	4	5	1	1	3
4	3	5	6	1	2	3	3	1

an example

[Chum et al. 2007]

<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	$\{a, b, c\}$	$\{b, c, d\}$	$\{a, e, f\}$
permutations						hash values		
3	6	2	5	4	1	2	2	1
1	2	6	3	5	4	1	2	1
3	2	1	6	4	5	1	1	3
4	3	5	6	1	2	3	3	1

an example

[Chum et al. 2007]


<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	$\{a, b, c\}$	$\{b, c, d\}$	$\{a, e, f\}$
permutations						hash values		
3	6	2	5	4	1	2	2	1
1	2	6	3	5	4	1	2	1
3	2	1	6	4	5	1	1	3
4	3	5	6	1	2	3	3	1

an example

[Chum et al. 2007]


<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	$\{a, b, c\}$	$\{b, c, d\}$	$\{a, e, f\}$
permutations						hash values		
3	6	2	5	4	1	2	2	1
1	2	6	3	5	4	1	2	1
3	2	1	6	4	5	1	1	3
4	3	5	6	1	2	3	3	1

matching maps




multiple matching pairs of feature maps

matching maps




multiple matching pairs of feature maps

matching maps




multiple matching pairs of feature maps

matching maps




multiple matching pairs of feature maps

matching maps



multiple matching pairs of feature maps

matching maps



multiple matching pairs of feature maps

retrieval

indexing



- construct inverted file of triplets (\hat{v}, w, π) (origin, hash value, permutation)
- memory requirements $10\times$ a typical baseline system

query

- retrieve images by triplets (\hat{v}, w, π) of query image
- re-estimate transformation parameters using LO-RANSAC
- re-ranking is an order of magnitude faster than FastSM, because an initial estimate is already available

European cities dataset 50K (EC50K)

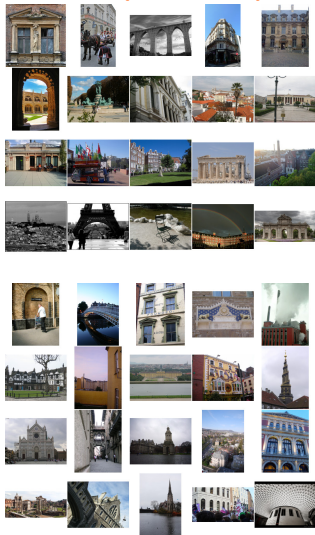

- 778 annotated images
- 20 groups of photos
- 5 queries from each group



publicly available: <http://image.ntua.gr/iva/datasets/ec50k>


European cities dataset 50K (EC50K)

- 778 annotated images
- 20 groups of photos
- 5 queries from each group
- 50,000 distractor images



publicly available: <http://image.ntua.gr/iva/datasets/ec50k>

feature map hashing—results EC50K



outline

- 1 local features and bag-of-words
- 2 local feature detection
- 3 visual vocabularies
- 4 spatial matching and re-ranking
- 5 geometry indexing
- 6 feature selection**
- 7 clustering of photo collections
- 8 location and landmark recognition
- 9 implementation: ivl library


selecting useful features

[Turcot and Lowe, ICCV 2009]

- **index space** is now the bottleneck in going to large scale, not speed
- select features by matching across **multiple views** of the same object or scene
- for each image in the database, find similar views, perform spatial matching, and **select** features appearing as inliers

selecting useful features

[Turcot and Lowe, ICCV 2009]



large scale geometry indexing


[Tolias et al., submitted to CVIU, 2012]

- feature map hashing implies **random selection**
- instead, select **robust** features, again by matching across similar views in dataset
- individual selection criteria for **origins** and **inlier features**
- dramatic reduction in index size

feature selection



results EC1M




feature selection by symmetry

[Tolias et al., submitted to ACM-MM 2012]

- feature selection so far relies on **multiple views**
- how about **unique views** of an object or scene?
- in fact, most images in a dataset are unique
- exploit **self-similarities**, **repeating patterns** and **symmetries**

matching scheme




direct matching




flipped matching




selected features



precision vs distractors



precision vs memory



outline

- 1 local features and bag-of-words
- 2 local feature detection
- 3 visual vocabularies
- 4 spatial matching and re-ranking
- 5 geometry indexing
- 6 feature selection
- 7 clustering of photo collections**
- 8 location and landmark recognition
- 9 implementation: `ivl` library

community photo collections

clustering / landmark recognition

- focus on popular subsets
- applications: browsing, 3D reconstruction




[Crandall *et al.*, ICCV 2009]

community photo collections

retrieval / location recognition

- include **all** images, has not yet scaled enough
- applications: automatic geo-tagging, camera pose estimation



Estimated Location Similar Image Incorrectly geo-tagged Unavailable

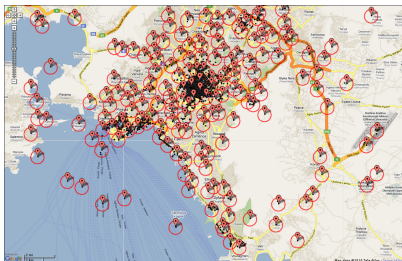


Suggested tags: Sint Antoniesbreestraat, Zwanenburgwal, Amsterdam
Frequent user tags: Anthoniesluis, sluijswacht, krom, stare, Skirt

view clustering

[Avrithis et al., ACM-MM 2010]

- identify images that potentially depict views of the **same scene**
- **geo clustering**: according to location
- **visual clustering**: according to visual similarity



- use **sub-linear indexing** in the clustering process

kernel vector quantization (KVQ)

[Tipping and Schölkopf, AIS 2001]

properties

- codebook vectors are points of the original dataset:
 $Q(D) \subseteq D$
- distortion upper bounded by r :
for all $x \in Q(D)$


$$\max_{y \in C(x)} d(x, y) < r$$

- the cluster collection

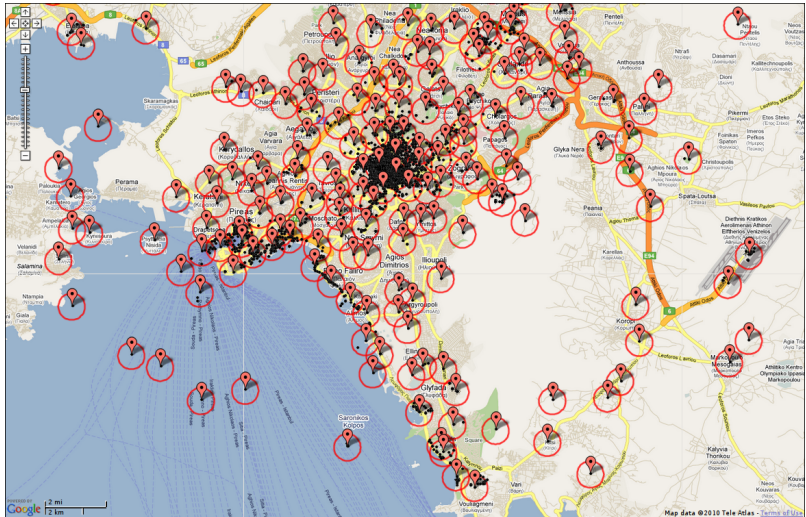
$$\mathcal{C}(D) = \{C(x) : x \in Q(D)\}$$

is a cover for D

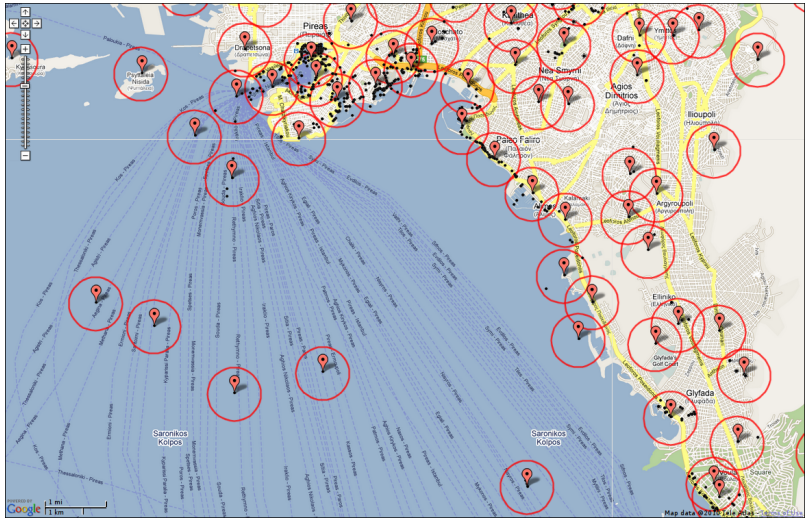
- clusters are overlapping



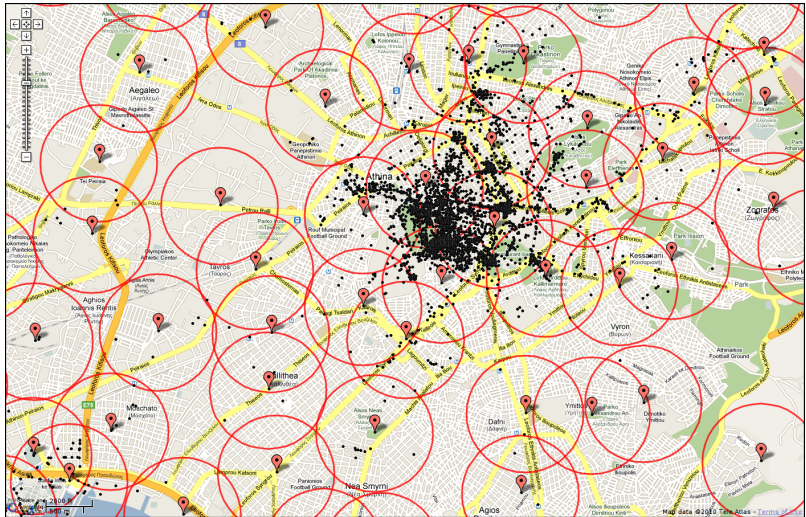
geo-clustering—example




geo-clustering—example



geo-clustering—example




geo-clustering—example



visual clustering

visual similarity measure


- $I(F_p, F_q)$: number of **inliers** between visual feature sets F_p, F_q of photos p, q respectively



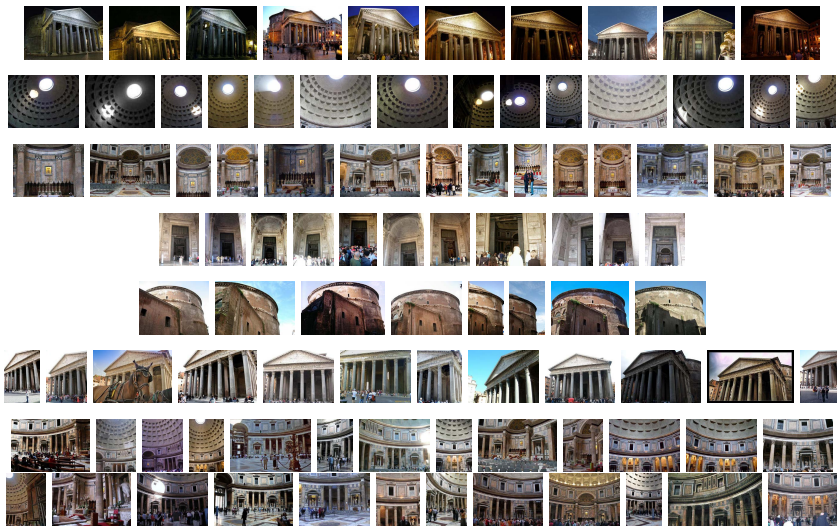
visual clustering—example

1,146 geo-tagged Flickr images of Pantheon, Rome

- 258 resulting visual clusters
- 30 images at each visual cluster on average
- an image belongs to 4 visual clusters on average



visual clustering—example



scene maps

[Avrithis et al., ACM-MM 2010]

- the image associated to the center of a view cluster shares at least one **rigid object** with all other images in the cluster
- treat this image as a **reference** for the cluster and **align** all other images to it
- initial estimates available from the view clustering stage—only local optimization needed
- construct a 2D **scene map** by grouping similar local features
- extend index, retrieval, and spatial matching for scene maps

view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images




view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—input images




view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images




view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images



view cluster alignment—example

Palau Nacional, Montjuic, Barcelona—aligned images




scene map construction—example

visual cluster containing 30 images of Palau Nacional, Montjuic




scene map construction—example



before vector quantization

scene map construction—example



after vector quantization

scene map indexing

index construction

- scene maps and images have the same representation—sets of features
- index all scene maps by visual word in an inverted file

query

- re-rank using the single correspondence assumption [Philbin et al. 2007]
- whenever a scene map $S(p)$ is found relevant, all images $q \in C_v(p)$ are retrieved as well

European cities 1M dataset (EC1M)

- 1,081 images in Barcelona, annotated into 35 groups
- 5 queries from each group
- all geo-tagged Flickr images



17 landmark groups




18 non-landmark groups

publicly available: <http://image.ntua.gr/iva/datasets/ec1m/>

European cities 1M dataset (EC1M)

- 908,859 **distractor** images from 21 European cities, excluding Barcelona
- most depict urban scenery like the ground-truth



publicly available: <http://image.ntua.gr/iva/datasets/ec1m/>

mining statistics—scene maps

- 1M images, 58 hours, single machine (8GB RAM), landmarks and non-landmarks



mining statistics—related work

- [Chum *et al.*, PAMI 2010] **web-scale clustering**: 5M images, 28 hours, single machine (64GB RAM), **popular subsets only**
- [Agarwal *et al.*, ICCV 2009] **building Rome in a day**: 150K images, 24 hours, 500 cores
- [Frahm *et al.*, ECCV 2010] **building Rome in a cloudless day**: 3M images, 24 hours, GPU
- [Heath *et al.*, CVPR 2010] **image webs**: 200K images, 4,5 hours, 500 cores


retrieval comparisons

- **baseline**: bag-of-words with **fast spatial matching** [Philbin et al. 2007]
- **QE1**: iterative **query expansion**, re-query using the retrieved images and merge results, 3 times iteratively
- **QE2**: create a scene map using the initial query's result and re-query once
- both QE schemes similar to **total recall** [Chum *et al.*, 2007]

query timing

Method	time	mAP
Baseline BoW	1.03s	0.642
QE1	20.30s	0.813
QE2	2.51s	0.686
Scene maps	1.29s	0.824

retrieval statistics



outline

- 1 local features and bag-of-words
- 2 local feature detection
- 3 visual vocabularies
- 4 spatial matching and re-ranking
- 5 geometry indexing
- 6 feature selection
- 7 clustering of photo collections
- 8 location and landmark recognition**
- 9 implementation: ivl library

location and landmark recognition

[Y. Kalantidis et al., MTAP 2011]

- assume that a subset of similar photos are correctly [geo-tagged](#), and not too far apart
- recognize the [location](#) where the query photo is taken, as the centroid of the most populated spatial (geo) cluster
- cross-validate locations and text (title, tags) of similar images with [Geonames](#) entries and geo-referenced [Wikipedia](#) articles
- link to known [landmarks](#) or [points of interest](#)

location recognition—examples



landmark recognition—examples




Suggested tags: Park Güell, Barcelona
Frequent user tags: Best of, me, Palau Güell



Suggested tags: La Barceloneta, Barcelona
Frequent user tags: honeymoon, wedding, straÙe



Suggested tags: Triumphal arch, Arc de Triomf, Barcelona
Frequent user tags: Sant Pere, Santa Caterina i La Ribera, macba, Passeig de Lluís Companys, lluis companys, Sant Beltra



Suggested tags: FC Barcelona Museum, Camp Nou, Barcelona
Frequent user tags: champions league, vfb, vfb stuttgart, Zoo de Barcelona, Camp Nou



Suggested tags: Montjuïc circuit, Museu Nacional d'Art de Catalunya, Barcelona
Frequent user tags: Montjuïc, castellers, Travelling Pooh, architecture, mnac



Suggested tags: Sagrada Família, Sagrada Família, Barcelona
Frequent user tags: gaudi, Sagrada Família, sagrada, familia, expiatorio

<http://viral.image.ntua.gr>

query



results



Estimated Location Similar Image Incorrectly geo-tagged Unavailable



Suggested tags: Buxton Memorial Fountain, Victoria Tower Gardens, London

Frequent user tags: Victoria Tower Gardens, Buxton Memorial Fountain, Winchester Palace, Architecture, Victorian gothic

Similar Images



Similarity: 0.619
Details Original ●●



Similarity: 0.491
Details Original ●●



Similarity: 0.397
Details Original ●●



Similarity: 0.385
Details Original ●●

similar of similar



Original



Original



Original



Original



Original



Original



Original



Original

similar of similar



Original ●●



Original ●●




Original ●●



Original ●●



Original ●●



Original ●●



Original ●●



Original ●●

similar of similar



Original ●●



Original ●●



Original ●●



Original ●●



Original ●●



Original ●●



Original ●●



Original ●●

suggested tags



Suggested tags: Buxton Memorial Fountain, Victoria Tower Gardens, London

Frequent user tags: Victoria Tower Gardens, Buxton Memorial Fountain, Winchester Palace, Architecture, Victorian gothic

related wikipedia articles



WIKIPEDIA
The Free Encyclopedia

Main page

Contents

Featured content

Current events

Random article

Interaction

About Wikipedia

Community portal

Recent changes

Contact Wikipedia

Donate to Wikipedia

Help

Toolbox

What links here

Related changes

Upload file

Special pages

Permanent link

Cite this page

Print/export

New features Log in / create account

Article **Discussion**

Read **Edit** View history

Buxton Memorial Fountain

From Wikipedia, the free encyclopedia

The **Buxton Memorial Fountain** is a memorial and *drinking fountain* in *London*, the *United Kingdom*, that commemorates the *emancipation of slaves* in the *British Empire* in 1834.

It was commissioned by *Charles Buxton* MP, and was dedicated to his father *Thomas Fowell Buxton* along with *William Wilberforce*, *Thomas Clarkson*, *Thomas Babington Macaulay*, *Henry Brougham* and *Stephen Lushington*, all of whom were involved in the abolition. It was designed by Gothic architect *Samuel Sanders Teulon* (1812–1873) in 1865 coincidentally with the passing of the *Thirteenth Amendment to the United States Constitution*, which effectively ended the western slave-trade.^[1]

It was originally constructed in *Parliament Square*, erected at a cost of £1,200. As part of the postwar redesign of the square it was removed in 1949 and not reinstated in its present position in *Victoria Tower Gardens* until 1957.^[2] There were eight decorative figures of British rulers on it, but four were stolen in 1960 and four in 1971. They were replaced by fibreglass figures in 1980. By 2005 these were missing, and the fountain was no longer working. Between autumn 2006 and February 2007 restoration works were carried out. The restored fountain was unveiled on 27 March 2007 as part of the commemoration of the 200th anniversary of the act to abolish the slave trade.^[3]

A memorial plaque commemorating the 150th anniversary of the *Anti-Slavery Society* was added in 1989.

Description

[\[edit\]](#)

The base is octagonal, about twelve feet in diameter, having open arches on the eight sides, supported on clustered shafts of polished Devonshire marble around a large central shaft, with four massive granite basins. Surmounting the pinnacles at the angles of the octagon are eight figures of bronze, representing the different rulers of England; the *Britons* represented by *Caractacus*, the *Romans* by *Constantine*, the *Danes* by *Canute*, the *Saxons* by *Alfred*, the *Normans* by *William the Conqueror*, and so on, ending with *Queen Victoria*. The fountain bears an inscription to the effect that it is "intended as a memorial of those members of Parliament who, with Mr. *Wilberforce*, advocated the abolition of the British slave-trade, achieved in 1807, and of those members of Parliament who, with Sir T.



The Buxton Memorial Fountain, designed by Samuel Sanders Teulon, celebrating the emancipation of slaves in the British Empire in 1834, in Victoria Tower Gardens, *Mitbank*, *London*.

related wikipedia articles



WIKIPEDIA
The Free Encyclopedia

Main page

Contents

Featured content

Current events

Random article

Donate

Interaction

About Wikipedia

Community portal

Recent changes

Contact Wikipedia

Help

Toolbox

What links here

Related changes

Upload file

Special pages

Permanent link

Cite this page

Print/export

New features Log in / create account

Article **Discussion**

Read **Edit** View history

Search

Victoria Tower Gardens

From Wikipedia, the free encyclopedia

Coordinates: 51°29′49.0″N 0°7′30.0″W﻿ / ﻿51.49694°N 0.12500°W﻿ / 51.49694; -0.12500

Victoria Tower Gardens is a public [park](#) along the north bank of the [River Thames](#) in [London](#). As its name suggests, it is adjacent to the [Victoria Tower](#), the south-western corner of the [Palace of Westminster](#). The park, which extends southwards from the Palace to [Lambeth Bridge](#), sandwiched between [Millbank](#) and the river, also forms part of the [Thames Embankment](#).

Contents [hide]

- Features
- Transport
- History
- External links
- References

Features

The park features:

- A reproduction of the [sculpture *The Burgbers of Calais*](#) by [Auguste Rodin](#), purchased by the [British Government](#) in 1911 and positioned in the Gardens in 1915.
- A 1930 statue of the [suffragette Emmeline Pankhurst](#), by [A.G. Walker](#).
- The [Buxton Memorial Fountain](#) – originally constructed in [Parliament Square](#), this was removed in 1940 and placed in its present position in 1957. It was commissioned by [Charles Buxton](#) MP to commemorate the [emancipation of slaves](#) in 1834, dedicated to his father [Thomas Fowell Buxton](#), and designed by [Gothic architect Samuel Sanders Teulon](#) (1812–1873) in 1865.
- A stone wall with two modern-style goats with kids – situated at the southern end of the Gardens.

Transport



Victoria Tower Gardens, 2005, with the [Buxton Memorial Fountain](#) at the front and the [Palace of Westminster](#) in the background

[edit]

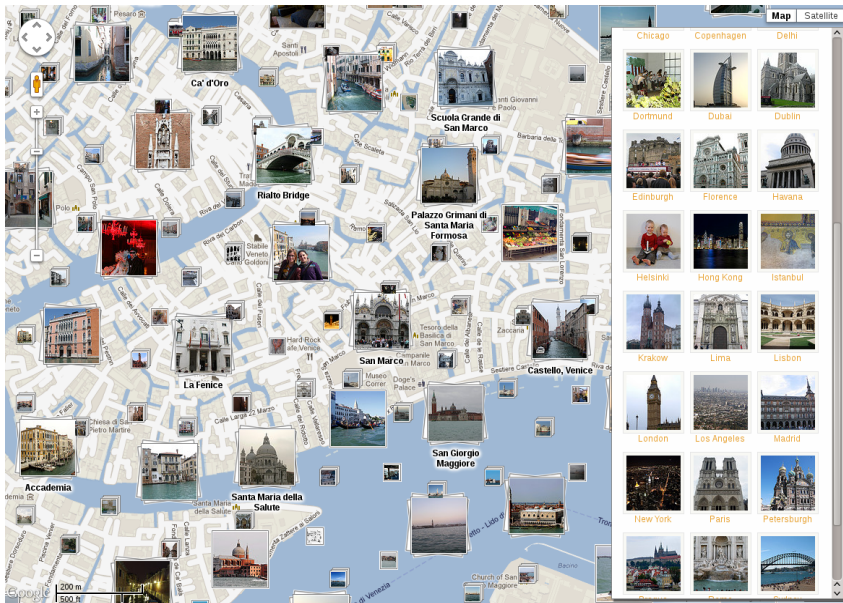
[edit]

VIRaL explore

The interface displays a map of Venice, Italy, with numerous image thumbnails overlaid on various locations. The thumbnails show a wide variety of subjects, including architectural landmarks like St. Mark's Basilica and the Campanile di San Marco, natural scenes like the Grand Canal, and cultural elements like the Venetian Arsenal. A sidebar on the right side of the map lists 27 cities, each with a small thumbnail image representing that city. The cities listed are: Chicago, Copenhagen, Delhi, Dortmund, Dubai, Dublin, Edinburgh, Florence, Havana, Helsinki, Hong Kong, Istanbul, Krakow, Lima, Lisbon, London, Los Angeles, Madrid, New York, Paris, and Petersburg. The map includes standard navigation controls such as a compass, a person icon, and zoom in/out buttons. A scale bar at the bottom left indicates 500 meters and 2000 feet. The Google logo is visible in the bottom left corner. The top right corner features a 'Map' button and a 'Satellite' button. The bottom of the screen has a series of navigation icons for back, forward, and search.

Chicago	Copenhagen	Delhi

VIRaL explore



VIRaL routes

Map Satellite

Identified landmarks

Ca' Pesaro

Frequent user tags

palazzo, italia - venecia, grand canal

User images

Similar images

Viewing Venice by ykaiant.

[Change photo set](#)

outline

- 1 local features and bag-of-words
- 2 local feature detection
- 3 visual vocabularies
- 4 spatial matching and re-ranking
- 5 geometry indexing
- 6 feature selection
- 7 clustering of photo collections
- 8 location and landmark recognition
- 9 implementation: ivl library**

recall feature point matching

- 1 construct the $m \times n$ proximity matrix G with elements

$$g_{ij} = \exp(-d_{ij}^2/2\sigma^2)$$

- 2 perform singular value decomposition of G

$$G = USV^T$$

where U, V are orthogonal matrices of dimension m, n and S is a non-negative diagonal $m \times n$ matrix

- 3 replace each diagonal element s_{ij} of S by 1 and reconstruct

$$P = UEV^T$$

- 4 finally, associate points a_i and b_j if element p_{ij} of P is the greatest element in its row and its column

Matlab code

```
function [m1, m2] =  
match(          x1,          y1,  
          x2,          y2, F s)  
  
[Ax1, Ax2] = meshgrid (x1, x2);  
[Ay1, Ay2] = meshgrid (y1, y2);  
  
D = sqrt((Ax1 - Ax2) .^ 2 + (Ay1 - Ay2) .^ 2);  
G = exp(-D .^ 2 ./ (2 * s .^ 2));  
  
[U, S, V] = svd (G);  
E = S > 0;  
P = U * E * V' ;  
  
[tmp, c] = max (P, [], 2);  
[tmp, r] = max (P, [], 1);  
  
match = r(c) == (1 : length(c) );  
m1 = find(match);  
m2 = c(match)';
```


ivl C++ code

```
template<class F>   ret<array<F>, array<F> >
match(const array<F>& x1, const array<F>& y1,
      const array<F>& x2, const array<F>& y2, F s)
{
    array_2d<F> Ax1, Ax2, Ay1, Ay2, U, S, V, tmp;
    _(Ax1, Ax2) = meshgrid++(x1, x2);
    _(Ay1, Ay2) = meshgrid++(y1, y2);

    array_2d<F> D = sqrt((Ax1 - Ax2) ->* 2 + (Ay1 - Ay2) ->* 2);
    array_2d<F> G = exp(-D ->* 2 / (2 * _[s] ->* 2));

    _(U, S, V) = svd++(G);
    array_2d<F> E = S > 0;
    array_2d<F> P = U ()* E ()* V(!_);

    array<int> c, r;
    _(tmp, c) = max++(P, _ , 2);
    _(tmp, r) = max++(P, _ , 1);

    array<bool> match = r[c] == rng(0, c.length() - 1);
    return _(find(match),
            c[match]);
}
```

ivl library

[Kontosis and Avrithis, expected 2012]

- C++ **template** library, compatible to STL
- supports most types, syntax and built-in operations of **Matlab** language
- fully **optimized**: minimal overhead/temporaries/copying; all array expressions boil down to a single for loop
- uses multiple CPU cores
- integrated with basic image functionalities of **OpenCV**
- integrated with most common **LAPACK** routines

plans

- integration with **QT** to support visualization
- **CUDA** massively parallel implementation on GPU

Credits




Spyros Leonardos




Yannis Kalantidis



Giorgos Tolias



Christos Varitimidis



Kimon Kontosis



Marios Phinikettos



Kostas Rapantzikos



Yannis Avrithis

project pages

<http://image.ntua.gr/iva/research>

VIRaL

<http://viral.image.ntua.gr>

datasets

<http://image.ntua.gr/iva/datasets>

thank you!