Image Matching using Local Symmetry Features

Daniel Cabrini Hauagge and Noah Snavely
{hauagge, snavely}@cs.cornell.edu

Motivation

Symmetry is a powerful feature in the structure of our world, evident in the shape and appearance of many natural and man-made scenes. For computer vision applications, analysis of symmetry is attractive for a number of reasons: symmetries are potentially a stable and robust feature of an object, yet, when considered at all scales and locations, are also potentially quite descriptive. Observe, for instance, the images below. It is clear that they all depict the Eiffel Tower, even though there is dramatic change in appearance. One can observe though that many of the local symmetries are preserved, like the vertical bilateral symmetry or the smaller bilateral symmetries of the metallic structure.

Contributions

- A simple and flexible definition of local symmetries.
- A scale space definition based on local symmetries.
- Two multi-scale feature detectors based on local symmetries.
- A new descriptor based on local symmetries.
- A challenging new dataset with a collection of images of man made structures that are rich in symmetries and vary widely in appearance.

Definitions

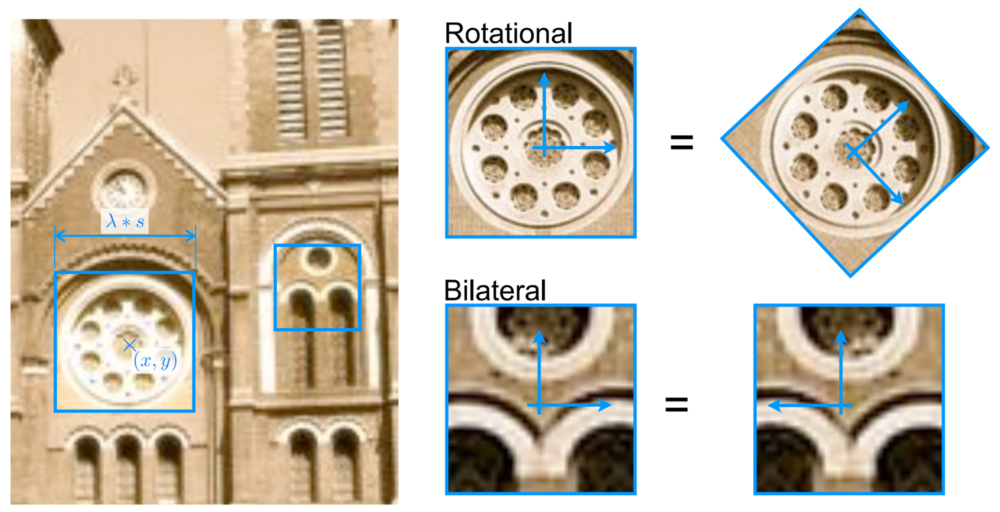

For a given location *(x,y)* and scale *s* we want to determine if a window of size λ*s around *(x,y)* exhibits rotational or bilateral symmetry.

Key Observation

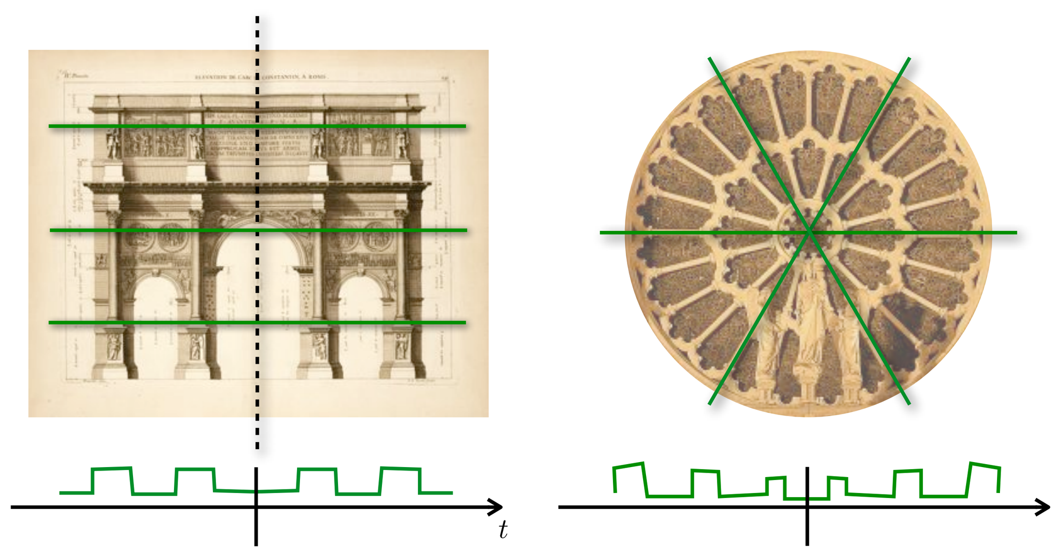

Slices through a symmetric image (in green) can be seen as even functions (plot in the bottom). Detecting symmetries can be posed as determining whether these slices are all approx. even functions.

Scoring Local Symmetries

To compute the Symmetry Distance (SD) for a given window we need three components: a weight mask (*e.g.* gaussian), a remapping function $g_{sp}(x)$ that computes the symmetrically opposing point to x for symmetry type s (*e.g.* rotational, horizontal, or vertical), and a pairwise distance $d(x,y)$ that compares the image at locations x and y (*e.g.* absolute difference between image intensities).

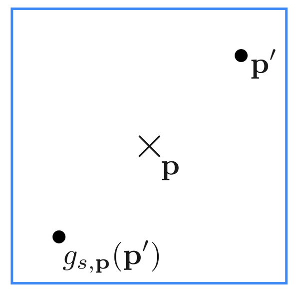

$$SD(\mathbf{p}) = \sum_{\mathbf{p}'} \underbrace{w(\|\mathbf{p}' - \mathbf{p}\|)}_{\text{weight mask}} \underbrace{d(\mathbf{p}', g_{s,\mathbf{p}}(\mathbf{p}'))}_{\text{pairwise distance function}}.$$

remapping function: $g_{s,\mathbf{p}}(\mathbf{p}')$

Pairwise Distance Function

- SYM-I: absolute difference in intensities.
- SYM-G: the dot product between (appropriately reflected) histograms of oriented gradients for each point.

The following figure shows our symmetry distance measure for three different symmetry types for SYM-G.

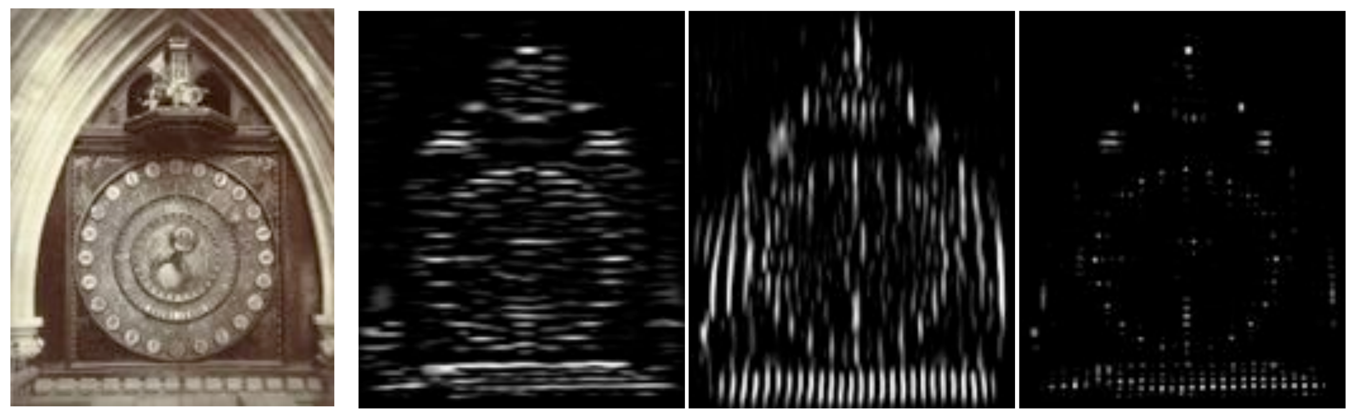

Horizontal | Vertical | Product of V and H

Local Symmetry Features

Detectors

In order to detect features at multiple scales we fix the window size and vary the image size. Local maxima with respect to scale and location are then used as our interest points. The following figure shows detection for SIFT and the two detectors we define. Note how the symmetry-based detectors more reliably fire on features such as the ring of circles around the clock's circumference.

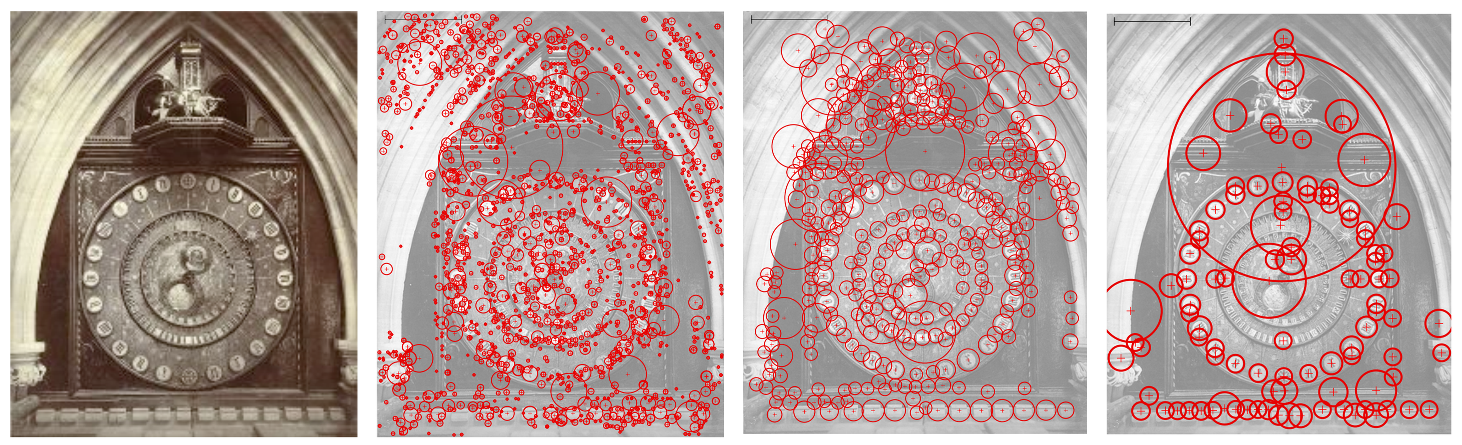

SIFT | SYM-I | SYM-G

An important difference between our two detectors is that SYM-I uses rotational symmetries to detect local symmetries while SYM-G uses the product of the horizontal and vertical symmetries.

Descriptor

Similar to encoding proposed by [Shechtman & Irani 2007], imposes a log polar grid around the interest point and records max symmetry score within each cell resulting in one descriptor per symmetry type. Repeating the process for different symmetry types and concatenating the resulting descriptors results in our final descriptor. We compute this descriptor for SYM-I only and compute horizontal, vertical, and rotational symmetries.

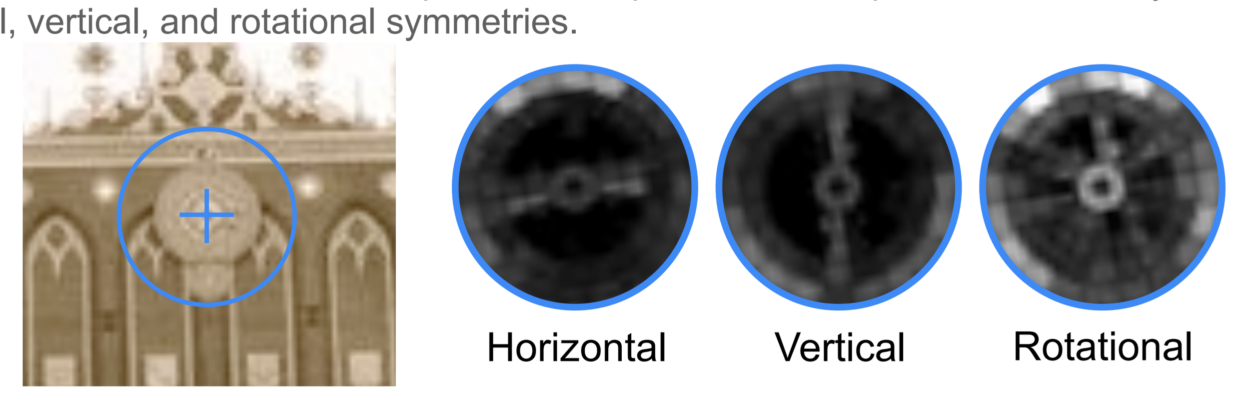

Horizontal | Vertical | Rotational

Dataset

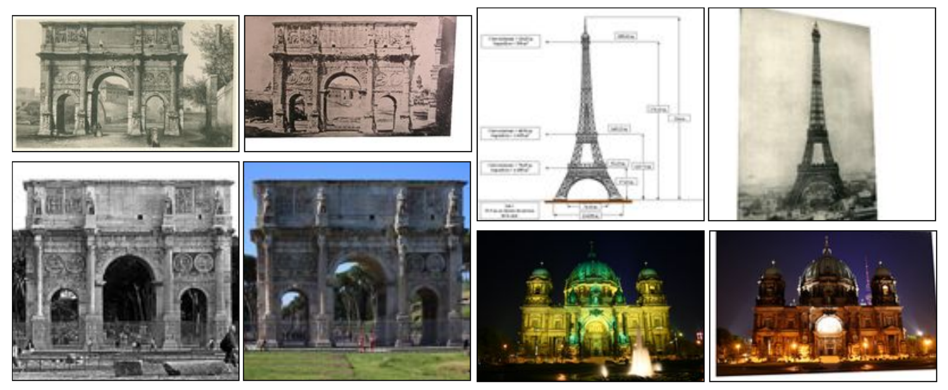

We present a new dataset consisting of 46 image pairs that exhibit large variation due to illumination, rendering style (e.g., paintings and drawings), and capturing technology (e.g., old vs. new pictures).

Results

Detector Evaluation

The following plots show the fraction of matched detections for the top *n* keypoints for one pair of images. Keypoints are ordered by either detector score or scale so that the detections with strongest response or largest support region come first. The table on the right shows aggregate results for the full dataset.

	Scale		Score	
	100	200	100	200
MSER	0.087	0.103	--	--
SIFT (DoG)	0.144	0.153	0.050	0.078
SYM-I	0.135	0.184	0.173	0.206
SYM-G	0.173	0.228	0.227	0.281

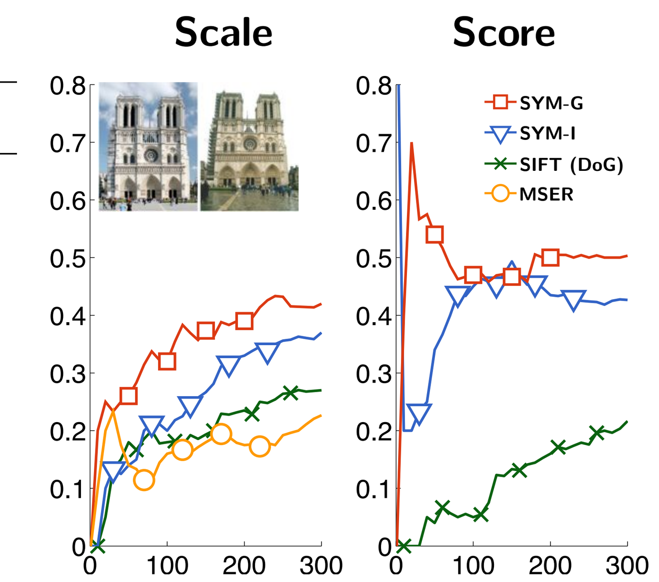

Descriptor Evaluation

The following plots show the precision-recall curves for a single pair of images from the dataset for various sets of feature points. Grid feature points are perfect detections, which allows us to measure descriptor performance isolated from the detector. The table on the lower left corresponds to averages over the entire dataset.

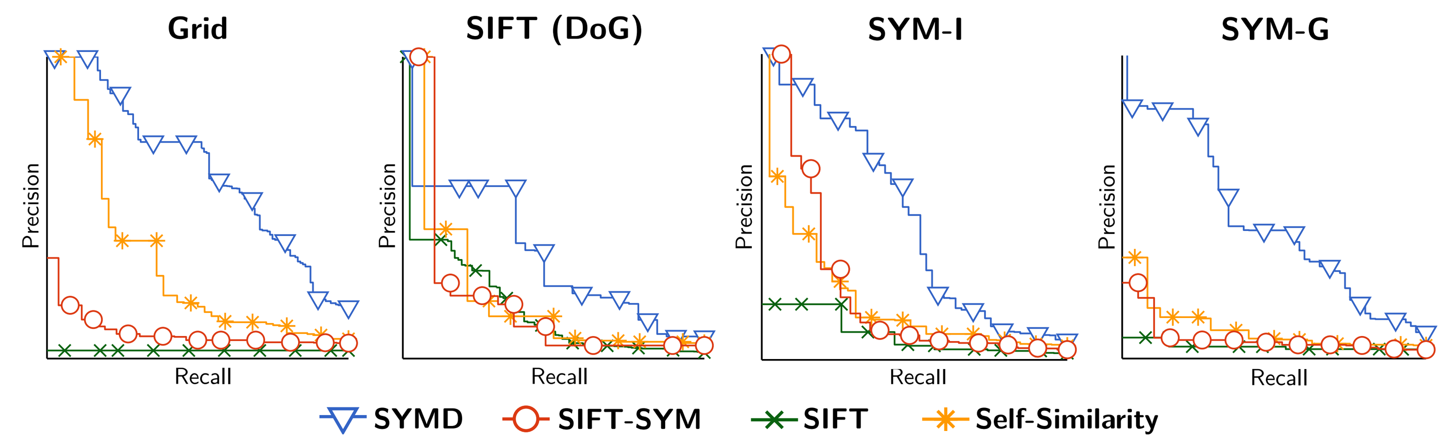

Mean Average Precision for our Dataset

	GRID	SIFT (DoG)	SYM-I	SYM-G
SELF-SIMILARITY	0.29	0.14	0.12	0.16
SIFT	0.49	0.21	0.28	0.25
SYMD	0.41	0.22	0.20	0.25
SIFT-SYMD	**0.58**	**0.28**	**0.35**	**0.36**

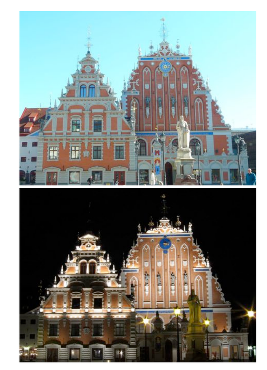

Mean average precision for different combinations of detector and descriptor. Each column corresponds to a different detector, and each row to a different descriptor.

For the full paper, dataset, and code visit the project website at
www.cs.cornell.edu/projects/symfeat