Features

All is Vanity, by C. Allan Gilbert, 1873-1929

Readings

- M. Brown et al. Multi-Image Matching using Multi-Scale Oriented Patches, CVPR 2005

Today's lecture

- Feature detection
- Feature matching
- Applications

Invariant local features

Find features that are invariant to transformations

- geometric invariance: translation, rotation, scale
- photometric invariance: brightness, exposure, …

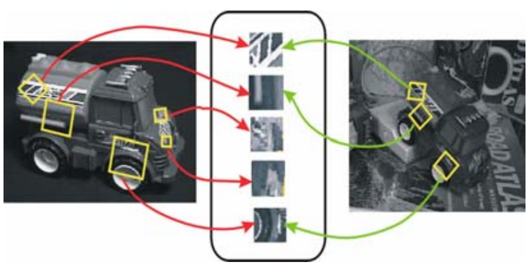

Feature Descriptors

Advantages of local features

Locality

- features are local, so robust to occlusion and clutter

Distinctiveness:

- can differentiate a large database of objects

Quantity

- hundreds or thousands in a single image

Efficiency

- real-time performance achievable

Generality

- exploit different types of features in different situations

More motivation…

Feature points are used for:

- Image alignment (e.g., mosaics)
- 3D reconstruction
- Motion tracking
- Object recognition
- Indexing and database retrieval
- Robot navigation
- … other

What makes a good feature?

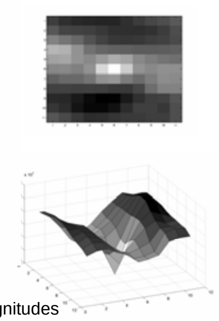

$$A^T A = \sum \nabla I (\nabla I)^T$$

– gradients are different, large magnitudes
– large $_1$, large $_2$

Feature detection

Want a “feature detection” function

- gives large values only for image patches that are good features
- How might you define f in terms of $_1$, $_2$?

The Harris operator

Want a “feature detection” function

- gives large values only for image patches that are good features
- How might you define f in terms of $_1$, $_2$?

$$f = \frac{\lambda_1 \lambda_2}{\lambda_1 + \lambda_2} \quad \text{(harmonic mean)}$$

$$= \frac{det(A^T A)}{trace(A^T A)}$$

- Called the “Harris Corner Detector” or “Harris Operator”
- Lots of other detectors, this is one of the most popular

The Algorithm:

- Find points with large response (f > threshold)
- Take the points of local maxima of f

Input images

Compute f

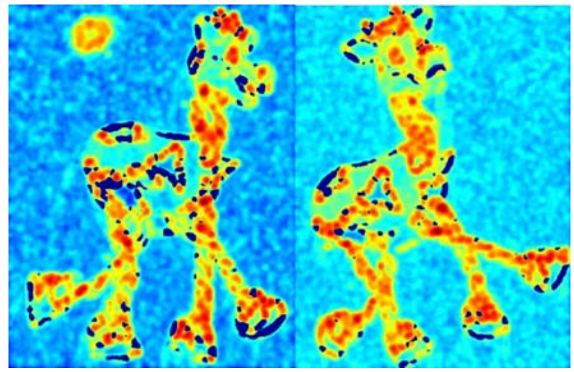

Threshold (f > value)

Find local maxima of f

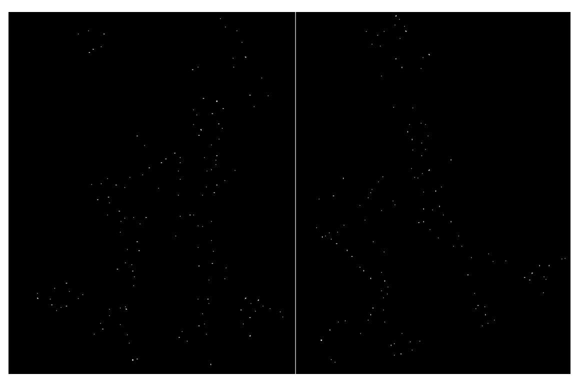

Harris features (in red)

Invariance

Suppose you rotate the image by some angle

- Will you still pick up the same features?

What if you translate the image instead?

Change in brightness?

Scale?

Scale invariant detection

Suppose you're looking for corners

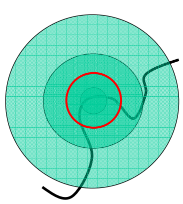

Key idea: find scale that gives local maximum of f

- f is a local maximum in both position and scale

Feature descriptors

We know how to detect good points

Next question: **How to match them?**

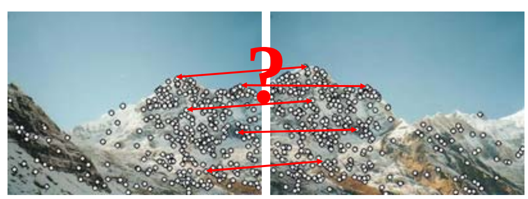

Lots of possibilities (this is a popular research area)

- simple: match square windows around the point
- state of the art (e.g.,): SIFT
  - David Lowe, UBC http://www.cs.ubc.ca/~lowe/keypoints/

Rotation invariance for feature descriptors

Find dominant orientation of the image patch

- From the motion lecture, this is given by the eigenvector of A^TA corresponding to the larger eigenvalue
- Rotate the patch according to this angle

MOPS [Brown, Szeliski, Winder, CVPR'2005]

Detections at multiple scales

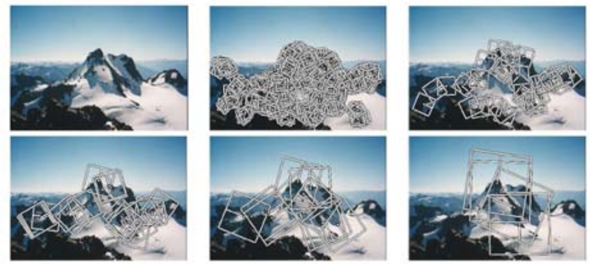

Figure 1. Multi-scale Oriented Patches (MOPS) extracted at five pyramid levels from one of the Matier images. The boxes show the feature orientation and the region from which the descriptor vector is sampled.

Matching features

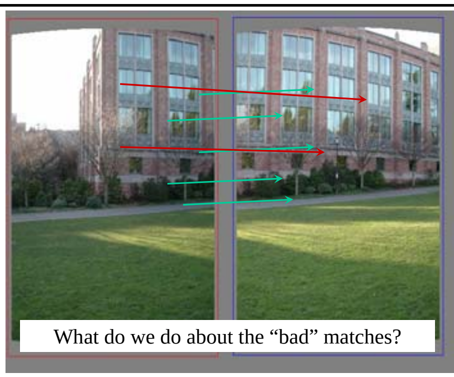

RAndom SAmple Consensus

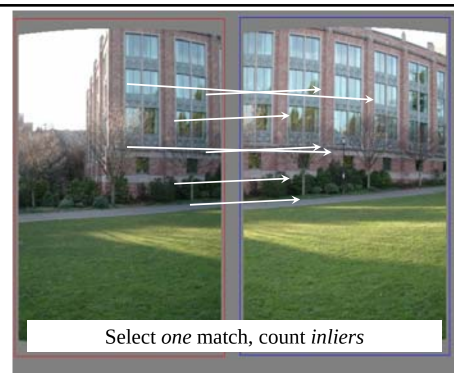

RAndom SAmple Consensus

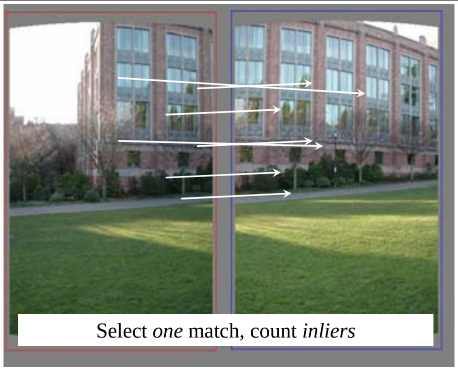

Least squares fit

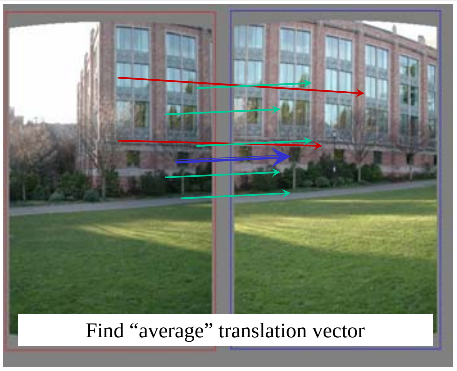

RANSAC (RANdom SAmpling Consensus)

Popular approach for robust model fitting with outliers

RANSAC loop:

1. Select K feature matches (at random)
2. Fit model (e.g., homography) based on these features
3. Count *inliers*:
   - number of other features that fit the model to within some specified threshold
4. The model with the largest number of inliers wins
5. Re-fit the model based on all of these inliers

More info:
http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL_COPIES/FISHER/RANSAC/

RANSAC

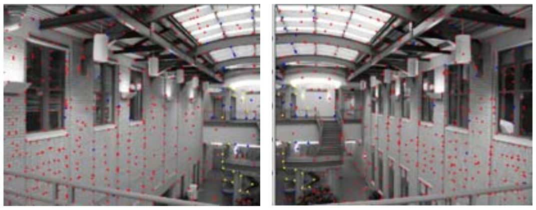

Lots of applications

Features are used for:

- Image alignment (e.g., mosaics)
- 3D reconstruction
- Motion tracking
- Object recognition
- Indexing and database retrieval
- Robot navigation
- … other

Autostitch (Brown and Lowe)

Fully automatic panorama generation

- Input: set of images
- Output: panorama(s)

Uses **SIFT** (Scale-Invariant Feature Transform) to find/align images

http://www.cs.ubc.ca/~mbrown/autostitch/autostitch.html

Microsoft version

- part of the Digital Image Pro and Digital Image Suite

1. Solve for homography

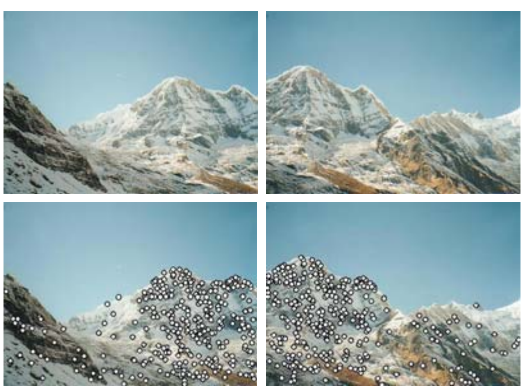

1. Solve for homography

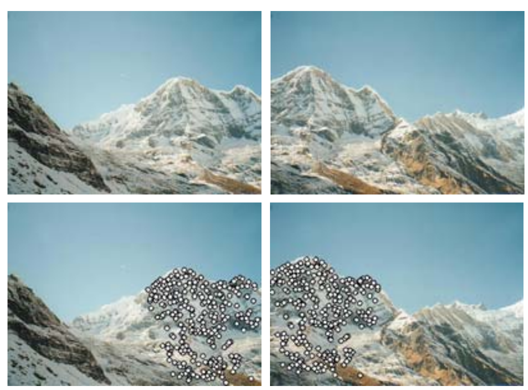

1. Solve for homography

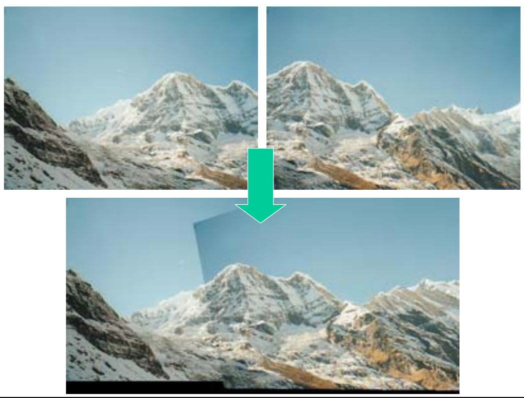

2. Find connected sets of images

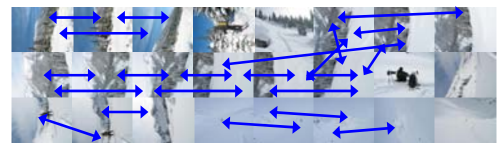

2. Find connected sets of images

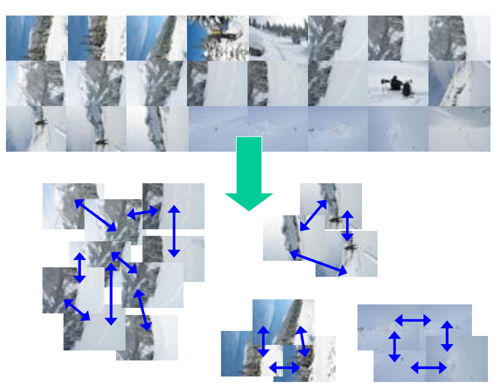

2. Find connected sets of images

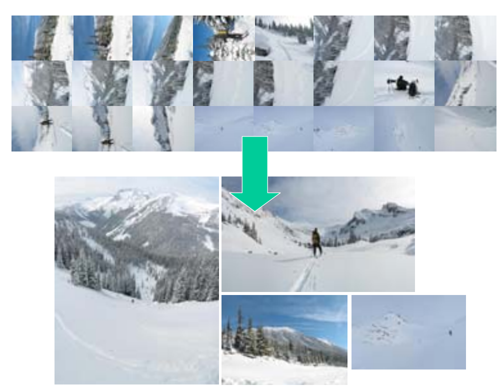

Object recognition (David Lowe)

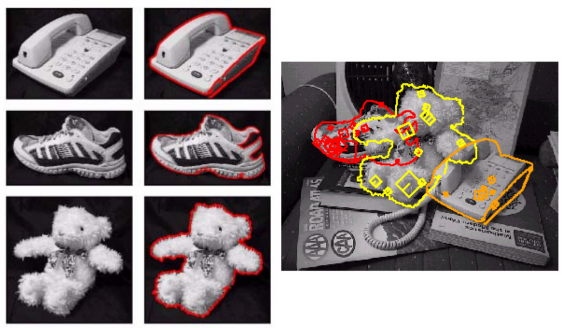

Sony Aibo

SIFT usage:

- Recognize charging station
- Communicate with visual cards
- Teach object recognition

The office of the past (Kim et al.)

http://grail.cs.washington.edu/projects/office/

3D scene recovery