

Automatic Image Alignment, Part 2




© Mike Nese

CS180: Intro to Comp. Vision and Comp. Photo
Efros & Kanazawa, UC Berkeley, Fall 2025


*with a lot of slides stolen from
Steve Seitz and Rick Szeliski*

Feature Detectors




Feature Descriptors

Schmid & Mohr 1997, Lowe 1999, Baumberg 2000, Tuytelaars & Van Gool 2000, Mikolajczyk & Schmid 2001, Brown & Lowe 2002, Matas et. al. 2002, Schaffalitzky & Zisserman 2002



Invariant Local Features

Image content is transformed into local feature coordinates that are invariant to translation, rotation, scale, and other imaging parameters



Features Descriptors

Applications

Feature points are used for:

- Image alignment (homography, fundamental matrix)
- 3D reconstruction
- Motion tracking
- Object recognition
- Indexing and database retrieval
- Robot navigation
- ... other

Today's lecture

- 1 Feature detector
 - scale invariant Harris corners
- 1 Feature descriptor
 - patches, oriented patches


Reading:

Multi-image Matching using Multi-scale image patches, CVPR 2005

Feature Detector – Harris Corner

Harris corner detector


C.Harris, M.Stephens. "A Combined Corner and Edge Detector". 1988




The Basic Idea

We should easily recognize the point by looking through a small window


Shifting a window in *any direction* should give a *large change* in intensity




Harris Detector: Basic Idea



“flat” region:
no change in
all directions



“edge”:
no change along
the edge direction





“corner”:
significant change
in all directions

Corner Detection: Mathematics

Change in appearance of window W for the shift $[u, v]$:

$$E(u, v) = \sum_{(x, y) \in W} [I(x + u, y + v) - I(x, y)]^2$$




Corner Detection: Mathematics


Change in appearance of window W for the shift $[u, v]$:

$$E(u, v) = \sum_{(x, y) \in W} [I(x + u, y + v) - I(x, y)]^2$$

$I(x, y)$



$E(u, v)$




Corner Detection: Mathematics

Change in appearance of window W for the shift $[u, v]$:

$$E(u, v) = \sum_{(x, y) \in W} [I(x + u, y + v) - I(x, y)]^2$$

We want to find out how this function behaves for small shifts

$E(u, v)$



Corner Detection: Mathematics

- First-order Taylor approximation for small motions $[u, v]$:

$$\begin{aligned} I(x+u, y+v) &= I(x, y) + I_x u + I_y v + \text{higher order terms} \\ &\approx I(x, y) + I_x u + I_y v \\ &= I(x, y) + \begin{bmatrix} I_x & I_y \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} \end{aligned}$$

- Let's plug this into

$$E(u, v) = \sum_{(x, y) \in W} [I(x+u, y+v) - I(x, y)]^2$$

Corner Detection: Mathematics

$$\begin{aligned} E(u, v) &= \sum_{(x, y) \in W} [I(x + u, y + v) - I(x, y)]^2 \\ &\approx \sum_{(x, y) \in W} [I(x, y) + \begin{bmatrix} I_x & I_y \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} - I(x, y)]^2 \\ &= \sum_{(x, y) \in W} \left(\begin{bmatrix} I_x & I_y \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} \right)^2 \\ &= \sum_{(x, y) \in W} \begin{bmatrix} u & v \end{bmatrix} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} \end{aligned}$$

Corner Detection: Mathematics

The quadratic approximation simplifies to

$$E(u, v) \approx \begin{bmatrix} u & v \end{bmatrix} M \begin{bmatrix} u \\ v \end{bmatrix}$$

where M is a *second moment matrix* computed from image derivatives:



$$M = \sum_{(x,y) \in W} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$$

Interpreting the second moment matrix

- The surface $E(u,v)$ is locally approximated by a quadratic form. Let's try to understand its shape.
- Specifically, in which directions does it have the smallest/greatest change?

$$E(u,v) \approx [u \ v] M \begin{bmatrix} u \\ v \end{bmatrix}$$

$$M = \sum_{(x,y) \in W} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$$



Interpreting the second moment matrix

First, consider the axis-aligned case
(gradients are either horizontal or vertical)


$$M = \sum_{(x,y) \in W} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} = \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix}$$

If either a or b is close to 0, then this is **not** a corner,
so look for locations where both are large.

Interpreting the second moment matrix

Consider a horizontal “slice” of $E(u, v)$: $[u \ v] M \begin{bmatrix} u \\ v \end{bmatrix} = \text{const}$


This is the equation of an ellipse.



Visualization of second moment matrices



Visualization of second moment matrices




Interpreting the second moment matrix

Consider a horizontal “slice” of $E(u, v)$: $[u \ v] M \begin{bmatrix} u \\ v \end{bmatrix} = \text{const}$

This is the equation of an ellipse.


Diagonalization of M : $M = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$

The axis lengths of the ellipse are determined by the eigenvalues and the orientation is determined by R



Interpreting the eigenvalues

Classification of image points using eigenvalues of M :



Harris Detector: Mathematics

Measure of corner response:

$$R = \frac{\det M}{\text{Trace } M}$$

$$\det M = \lambda_1 \lambda_2$$


$$\text{trace } M = \lambda_1 + \lambda_2$$

Harris detector: Steps

1. Compute Gaussian derivatives at each pixel
2. For each sliding window in image, compute second moment matrix M in a Gaussian window around each pixel
3. Compute corner response function R
4. Threshold R
5. Find local maxima of response function (nonmaximum suppression)


C.Harris and M.Stephens. [“A Combined Corner and Edge Detector.”](#)
Proceedings of the 4th Alvey Vision Conference: pages 147—151, 1988.

Harris Detector: Workflow




Harris Detector: Workflow

Compute corner response R



Harris Detector: Workflow

Find points with large corner response: $R > \text{threshold}$



Harris Detector: Workflow

Take only the points of local maxima of R




Harris Detector: Workflow



Harris Detector: Some Properties

Rotation invariance





Ellipse rotates but its shape (i.e. eigenvalues) remains the same

Corner response R is invariant to image rotation

Harris Detector: Some Properties


Partial invariance to *affine intensity* change

- ✓ Only derivatives are used => invariance to intensity shift $I \rightarrow I + b$
- ✓ Intensity scale: $I \rightarrow a I$




Harris Detector: Some Properties

But: non-invariant to *image scale*!




All points will be
classified as **edges**



Corner !

Scale Invariant Detection


Consider regions (e.g. circles) of different sizes around a point
Regions of corresponding sizes will look the same in both images



Scale Invariant Detection


The problem: how do we choose corresponding circles *independently* in each image?

Choose the scale of the “best” corner



Feature selection


Distribute points evenly over the image




Adaptive Non-maximal Suppression

Desired: Fixed # of features per image

- Want evenly distributed spatially...
- Sort points by non-maximal suppression radius
[Brown, Szeliski, Winder, CVPR'05]



(a) Strongest 250



(b) Strongest 500



(c) ANMS 250, $r = 24$




(d) ANMS 500, $r = 16$

Feature descriptors

We know how to detect points

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



Point descriptor should be:

1. Invariant

2. Distinctive

Feature Descriptor – MOPS

Multi-Scale Oriented Patches

Interest points

- Multi-scale Harris corners
- Orientation from blurred gradient
- Geometrically invariant to rotation

Descriptor vector

- Bias/gain normalized sampling of local patch (8x8)
- Photometrically invariant to affine changes in intensity

[Brown, Szeliski, Winder, CVPR'2005]

Detect Features, setup Frame

Orientation = blurred gradient

Rotation Invariant Frame

- Scale-space position (x, y, s) + orientation (θ)



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.

MOPS descriptor vector

8x8 oriented patch





- Sampled at 5 x scale

Bias/gain normalisation: $I' = (I - \mu)/\sigma$



Automatic Feature Matching





Feature matching



Feature matching

- Pick best match!
 - For every patch in image 1, find the most similar patch (e.g. by L2 distance).
 - Called “nearest neighbor” in machine learning
- Can do various speed ups:
 - Hashing
 - compute a short descriptor from each feature vector, or hash longer descriptors (randomly)
 - Fast Nearest neighbor techniques
 - *kd*-trees and their variants
 - Clustering / Vector quantization
 - So called “visual words”

What about outliers?




Feature-space outlier rejection

Let's not match all features, but only these that have “similar enough” matches?

How can we do it?

- $\text{SSD}(\text{patch1}, \text{patch2}) < \text{threshold}$
- How to set threshold?



Feature-space outlier rejection: symmetry

Let's not match all features, but only these that have
“similar enough” matches?


How can we do it?

- Symmetry: x 's NN is y , and y 's NN is x



Feature-space outlier rejection: Lowe's trick

A better way [Lowe, 1999]:

- 1-NN: L2 of the closest match
- 2-NN: L2 of the second-closest match
- Look at how much better 1-NN is than 2-NN, e.g. 1-NN/2-NN
- That is, is our best match so much better than the rest?




Feature-space outlier rejection




Can we now compute H from the blue points?

- No! Still too many outliers...
- What can we do?


Matching features




Random Sample Consensus



Random Sample Consensus



Least squares fit



RANSAC for estimating homography

RANSAC loop:

- 
1. Select four feature pairs (at random)
 2. Compute homography H (exact)
 3. Compute *inliers* where $\text{dist}(p_i', \mathbf{H} p_i) < \varepsilon$
 4. Keep largest set of inliers
 5. Re-compute least-squares H estimate on all of the inliers

RANSAC

