Topic 11:

Feature Detection & Image Matching

- Introduction to the image matching problem
- Image matching using SIFT features
- The SIFT feature detector
- The SIFT descriptor





Goal: Identify "features" or patches in image I that appear in another image, I'





Indicates a correspondence between location (x, y) in image I and location (x', y') in image I'





Q:

Is it possible to solve this problem by direct template matching between two images?





Image



Q:

Is it possible to solve this problem by direct template matching between two images?







Q:

A:

Is it possible to solve this problem by direct template matching between two images?



Yes, but it would be impossibly inefficient (ie. must search over all possible pairs of patches)

Feature-Based Image Matching



Feature-Based Image Matching



Errors in Feature-Based Image Matching





In general, some/many of these correspondences may be incorrect.

Two types of error:

1. False positive matches

algorithm returns a correspondence between 2 locations where none exists

2. False negative matches

algorithm fails to detect a correspondence between two instances of the same feature/patch.



Errors in Feature-Based Image Matching





GOAL: minimize false positive and false negatives across a wide range of imaging conditions.

1. False positive matches

algorithm returns a correspondence between 2 locations where none exists

2. False negative matches

algorithm fails to detect a correspondence between two instances of the same feature/patch.



Evaluating a Feature Detector's Performance



Feature Matching & Transformation Invariance





To be most useful, the feature detector & matching algorithm must be insensitive to a wide range of image transformations.

"Transformed" source images







A feature detector is called **invariant** to a certain image transformation if it can reliably detect features in a transformed version of the source image

Melba elba

OF THE HEDGE FUND WORL

"Transformed" source images





A feature detector is called **invariant** to a certain image transformation if it can reliably detect features in a transformed version of the source image

"Transformed" source images



Brightness transformation



"Transformed" source images



A feature detector is called **invariant** to a certain image transformation if it can reliably detect features in a transformed version of the source image Distortion due to change in viewpoint





"Transformed" source images

A feature detector is called **invariant** to a certain image transformation if it can reliably detect features in a transformed version of the source image



Distortion due to change in viewpoint & magnification (ie. scale)

Topic 11:

Feature Detection & Image Matching

- Introduction to the image matching problem
- Image matching using SIFT features
- The SIFT feature detector
- The SIFT descriptor

SIFT: Scale Invariant Feature Transform

- Developed by David Lowe in 1999
- One of the most powerful representations for feature detection and matching
- Widely used in applications that range from robotics, to image retrieval & recognition, image stitching, video analysis.







Scale + Viewpoint





The SIFT Feature Detection Algorithm

Goal: Represent an image I as a collection of SIFT features that can be identified reliably in other images where the same (or similar) objects are present.

Input: Image I Output: A set of k SIFT keypoints {p₁, p₂, ..., p_k} & feature vectors {f₁, f₂, ..., f_k}.

"Source" image I



"Query" image I



The SIFT Feature Detection Algorithm

Goal: Represent an image I as a collection of SIFT features that can be identified reliably in other images where the same (or similar) objects are present.

Input: Image I Output: A set of k SIFT keypoints {p₁, p₂, ..., p_k} & feature vectors {f₁, f₂, ..., f_k}.

"Source" image I



"Query" image I



The SIFT Keypoints

Keypoint: A location (x, y) in the source image, with an associated orientation & scale, that is "visually distinct" from its surroundings.

Input: Image I Output: A set of *k* SIFT

keypoints





Detected keypoints



The SIFT Feature Vectors

Feature vector (of a keypoint p_i): A vector of fixed length that represents the image patch centred at p_i .

Input: Image I Output: A set of k A set of k feature vectors keypoints $p_i = (x_i, y_i, \rho_i, \theta_i)$ $f_i =$ location describes the patch orientation scale centred at pi

Detected keypoints



Representing an Image Using SIFT Features: Steps

Source Image I



vectors



The number k of detected keypoints depends on I

The dimension of each feature vector is the same for all images.

Matching 2 Images Using SIFT Features



1. Identify keypoints 2. Build feature vectors

Matching 2 Images Using SIFT Features



3. Match vectors

Step 1: Compute a Set of Keypoints

Source image I



1. Identify keypoints



Goals:

- Identify distinctive image locations
- Assign scale & orientation to each keypoint
- Should be able to detect some keypoint in images that vary in magnification, brightness, etc.

Step 1: Compute a Set of Keypoints

Keypoint: A location (x, y) in the source image, with an associated orientation & scale, that is "visually distinct" from its surroundings.

Input: Image I Output: A set of *k* SIFT

keypoints





Detected keypoints

length

scale

denotes

Topic 11:

Feature Detection & Image Matching

- Introduction to the image matching problem
- Image matching using SIFT features
- The SIFT feature detector
- The SIFT descriptor

Computing SIFT Keypoints: Basic Steps



Orientation assign

Extremum pruning

Location refinement

Step 1a: Construct a Gauss-like Pyramid

Step 1a: Compute a pyramid of Gauss-filtered images organized into **octaves** of s + 1 images

each image is smoothed by a factor of k more than the image below it





Gaussian

Step 1a: Construct a Gauss-like Pyramid

SIFT terminology An **octave** is a set of Gauss-convolved images, I_1 , ..., I_s representing a doubling of the scale parameter σ between I_1 and I_s .

each image is smoothed by a factor of k more than the image below it

each octave contains s + 1 images

$$I_{s} = I * G_{k^{s}\sigma}$$

$$\vdots$$

$$I_{2} = I * G_{k(k\sigma)}$$

$$I_{1} = I * G_{k\sigma}$$

$$I_{0} = I * G_{\sigma}$$

Gaussian

Step 1a: Construct a Gauss-like Pyramid

Step 1a: Compute a pyramid of Gauss-filtered images organized into octaves of s + 1

each

s + 1



Computing SIFT Keypoints: Basic Steps



Orientation assign

Extremum pruning

Location refinement
Step 1b: Compute Pyramid of DOG Images

Step 1b: Compute a pyramid of DOG-filtered images

$$D(x, y, \rho) = I(x, y) * (G(x, y, k\rho) - G(xy, \rho))$$

for $\rho = \sigma, k\sigma, k^2\sigma, \dots, k^{S-1}\sigma$

$$I_{s} = I * G_{k^{s}\sigma}$$

$$I_{2} = I * G_{k(k\sigma)}$$

$$I_{1} = I * G_{k\sigma}$$

$$I_{0} = I * G_{\sigma}$$

$$Gaussian$$

$$D(x, y, k^{s}\sigma)$$

$$D(x, k^{s}\sigma)$$

 $D(x, y, k^{S}\sigma)$

Step 1b: Compute Pyramid of DOG Images

Step 1b: Compute a pyramid of DOG-filtered images



Reminder: Difference-of-Gaussian Filtering

 $I * G_{\rho}$



Reminder: Difference-of-Gaussian Filtering

 $I * G_{k\rho}$



Reminder: Difference-of-Gaussian Filtering

Difference of two Gaussiansmoothed versions of I:

$$I * G_{k\rho} - I * G_{\rho} = I * (G_{k\rho} - G_{\rho})$$

(just the difference between two Gaussian masks)



But we know that $G_{k\rho} - G_{\rho} = k\rho(k\rho - \rho) \nabla^2 G_{\rho}$

 $D(x, y, \rho) = \left[\nabla^2 (I * G_{\rho}) \right] \rho^2 k(k-1)$

Computing SIFT Keypoints: Basic Steps



Orientation assign

Extremum pruning

Location refinement

Step 1c: Detecting DOG Extrema

Step 1c (Extremum detection): Find all pixels that correspond to extrema of D(x, y, ρ)

Extrema of the image Laplacian are readily distinguishable from their surroundings. There are usually just a few thousand in each pyramid.



The Difference-Of-Gaussians (DOG) Filter

Finding local extrema in a single image $D(x, y, \rho)$



minimum at

 (x, y) if
 D(x,y,ρ) < all
 neighbours



maximum if
 D(x,y,ρ) > all
 neighbours



Step 1c: Detecting DOG Extrema

Step 1c (Extremum detection): Find all pixels that correspond to extrema of $D(x, y, \rho)$



Step 1c: Detecting DOG Extrema

Step 1c (Extremum detection): Find all pixels that correspond to extrema of $D(x, y, \rho)$



Step 1c: Detecting DOG Extrema: Algorithm

Step 1c (Extremum detection): Find all pixels that correspond to extrema of D(x, y, ρ)

For each (x, y, ρ), check whether D(x, y, ρ) is greater than (or smaller than) all of its neighbours in current scale and adjacent scales above & below.



Step 1c: SIFT Keypoints = DOG Extrema

Step 1c (Extremum detection): Find all pixels that correspond to extrema of D(x, y, ρ)

An extremum detected at D(x, y, ρ) defines the keypoint (x, y, ρ).



Computing SIFT Keypoints: Basic Steps



Orientation assign

Extremum pruning

Location refinement

Step 1d: Refining Location of Extrema

Step 1d (Extremum localization) Refine the location of detected extrema through a quadratic least-squares fit.

- The original SIFT method/paper just uses the pixel location of the extrema as the location of the key point
 - Revised method finds the subpixel location interpolated location of the extrema using 2nd order Taylor expansion of *D* at (*x*, *y*, *ρ*)
 - This improves the results when matching significantly



Step 1d: Refining Location of Extrema



2. Take derivatives with respect to $\Delta \vec{x}$

$$\frac{\partial D}{\partial (\Delta \vec{x})} = \left(\frac{\partial D}{\partial \vec{x}}\right)^T + \left(\frac{\partial^2 D}{\partial \vec{x}^2}\right) (\Delta \vec{x})$$

3. Extremum \Leftrightarrow derivative is zero \Rightarrow solve for $\frac{\partial D}{\partial \Lambda \overrightarrow{x}} = 0$

$$(\Delta \vec{x}) = \left(\frac{\partial^2 D}{\partial \vec{x}^2}\right)^{-1} \cdot \left(\frac{\partial D}{\partial \vec{x}}\right)$$

Computing SIFT Keypoints: Basic Steps



Orientation assign

Extremum pruning

Location refinement









Computing SIFT Keypoints: Basic Steps



Orientation assign

Extremum pruning

Location refinement

Step 1f: Keypoint Orientation Assignment

Assigning an orientation θ_i to keypoint (x'_i, y'_i, ρ'_i) :



- A. Compute smoothed image
- B. Compute gradient magnitude & orientation in neighbourhood of (x'_i, y'_i) in $I * G_{\rho'_i}$
- C. Compute histogram of orientations
- D. Assigned orientation
 - θ_i = highest peak in histogram

Step 1f: Keypoint Orientation Assignment

Computing Histogram of Orientations



- Orientations divided into 36 bins (one every 10 degrees).
- Pixel (x, y) contributes to the bin corresponding to the gradient orientation θ at (x, y).
- Contribution to the bin is equal to $|\nabla I(x, y)| \cdot G_{1.5\rho'_i}(d)$
- Total bin weight = sum of contributions from all pixels

Computing SIFT Keypoints: Basic Steps



Orientation assign

Extremum pruning

Location refinement

Topic 11:

Feature Detection & Image Matching

- Introduction to the image matching problem
- Image matching using SIFT features
- The SIFT feature detector
- The SIFT descriptor

The SIFT Keypoints

Keypoint: A location (x, y) in the source image, with an associated orientation & scale, that is "visually distinct" from its surroundings.

Input: Image I Output: A set of *k* SIFT

keypoints





Detected keypoints



The SIFT Feature Vectors

Feature vector (of a keypoint p_i): A vector of fixed length that represents the image patch centred at p_i .

Input: Image I Output: A set of k A set of k feature vectors keypoints $p_i = (x_i, y_i, \rho_i, \theta_i)$ $f_i =$ location describes the patch orientation scale centred at pi

Detected keypoints





SIFT Descriptor



1. Compute gradients in 16 x 16 pixel patch of image $I^* G_{\sigma_i}$ centred at (x_i, y_i)

Gaussian-smoothed image at scale of the keypoint



1. Compute gradients in 16 x 16 pixel patch of image $I^* G_{\sigma_i}$ centred at (x_i, y_i) 2. Compute gradient orientation relative to keypoint orientation

$$\theta(x, y) = \tan^{-1} \left[\frac{\partial (I * G_{\sigma_i})}{\partial y} \middle/ \frac{\partial (I * G_{\sigma_i})}{\partial x} \right] - \theta_i$$

SIFT Descriptor





- 1. Compute gradients
- 2. Compute relative gradient orientations
- 3. Compute orientation histogram of each 4x4 pixel block
 Histogram contains 8 bins, each covering 45°

The 4x4 Orientation Histogram



Weight of (x, y):

$$w = G_{\frac{\sigma_i}{2}} \left(\left\| (x - x_i, y - y_i) \right\| \right)$$

⇒ pixels closer to keypoint centre have higher weight Total contribution of (x, y) to the orientation histograms: $C(x, y) = w \cdot ||\nabla I * G_{\sigma_i}(x, y)||$

gradient magnitude of smoothed image



Contribution spread across 2 closest orientations & 3 closest histograms

⇒ no abrupt changes in histogram if keypoint centre displaced by 3 - 4 pixels Total contribution of (x, y) to the orientation histograms: $C(x, y) = w \cdot \|\nabla I * G_{\sigma_i}(x, y)\|$



orientations & 3 closest histograms

⇒ no abrupt changes in histogram if keypoint centre displaced by 3 - 4 pixels orientation 1: $\frac{\theta_1}{\theta_1 + \theta_2}$ orientation 2: $\frac{\theta_2}{\theta_1 + \theta_2}$



- 1. Compute gradients
- 2. Compute relative gradient orientations
- 3. Define an "accumulator" variable for each of the 8

orientations in each of the 16 histograms (128 total).

4. For each pixel, calculate the pixel's contribution to each accumulator variable.

Converting SIFT Descriptors to 128-dim Vectors

Post processing: 1. Normalize f_i :

$$f_i \to \frac{f_i}{\|f_i\|}$$

⇒ gives invariance to linear lighting variations across images, ie. when matching image I and image al + b (because f_i will be the **same** in both images)

2. Clamp f_i :

Clamp all elements of f_i at 0.2

⇒ gives less weight to very large gradient magnitudes

SIFT Descriptor



3. Re-normalize

Matching 2 Images Using SIFT Features



1. Identify keypoints 2. Build feature vectors
Matching 2 Images Using SIFT Features



Matching 2 Images Using SIFT Features





Intuition for matching algorithm:

match established only if it is deemed reliable, ie. if there is only one very similar feature in image l'



1. Identify keypoints 2. Build feature vectors