

CS 223B Computer Vision

Problem Set 3

Due: Feb. 22nd, 2011

1 Probabilistic Recursion for Tracking

In this problem you will derive a method for tracking a point of interest through a sequence of images. In this problem there exists a true location of an object or interest point that you want to estimate in each frame. We observe this true location through a noisy image. These observations will then be used in conjunction with prior knowledge of the general position where our object or interest point is known to lie. Formalizing this description we can write down the following random variables

$x_k \in \mathbb{R}^d$: the ground truth location of our object or interest point at time k

$X_k = [x_1, \dots, x_k]^T$: time history to step k

$z_k \in \mathbb{R}^c$: our noisy position measurement at time k of x_k

$Z_k = [z_1, \dots, z_k]^T$: history to time step k

To gain our estimate for the position of the object or interest point at time step k we would like to solve for $p(x_k|Z_k)$. In addition, our estimate needs to be computationally efficient, and should be easily updated for when we wish to estimate the position at $k + 1$. Thus we will calculate our estimate using a probabilistic recursion. To be able to compute this recursion it can only depend on distributions which are stored in memory. The distributions we have stored in memory are

$p(z_k|x_k)$: our measurement model

$p(x_k|x_{k-1})$: transition model

$p(x_{k-1}|Z_{k-1})$: the result of our recursion at the previous iteration

- (a) In this part of the problem we will assume that Z_{k-1} and x_k are conditionally independent. Using this assumption derive an expression for $p(x_k|Z_k)$ in terms of $p(z_k|x_k)$, $p(x_k|Z_{k-1})$ and $p(z_k|Z_{k-1})$. This is conditioned solely on our ground truth location at time k and measurements up to and including time $k - 1$. Justify each step of your derivation (it should not be more than a few lines).
- (b) The $p(x_k|Z_{k-1})$ term prevents our expression on being solely dependent on distributions that are stored in memory. We will now assume that x_{k-1} and Z_{k-1} are conditionally independent. Now using this assumption write our recursion as an expression which is solely dependent on distributions we have stored in memory. Justify your answer.

Remark: If both our measurement and transition model are Gaussian distributions, there exists a closed form solution for the recursion. In general, the recursion is approximated using numerical techniques known as Monte Carlo methods.

2 Sifting through SIFT

SIFT has become a household name for anyone involved with computer vision. However, it is often naively or incorrectly applied when a person simply downloads a compiled version off the internet and applies it to their task. The purpose of this problem is to illuminate subtleties of SIFT that can be overlooked, while also giving practical experience to the description and matching methods developed by Lowe et al.

- (a) It is desirable to locate image features that are stable across image instances and noise variation. Interest points are related to distinct changes in neighboring pixel values, but choosing the correct operator was the focus of several papers. For comparison, consider Fig. 1 which contains an input image, a difference of gaussians (DoG) of the input, and a gradient of the input. If you wanted to detect interest points that produced a pronounced signal across image variation and random noise, which operator would you use? Specifically, write down the mathematical formulation for both approaches (DoG and gradient) and identify properties that would support your decision. *hint: think about the what these kernels look like in the frequency domain.*
- (b) A large part of SIFT's immense popularity (10,000+ citations) is its excellent marriage of a keypoint detector and a descriptor. Choosing the keypoints involves many heuristics such as a spatial sampling, removing edge points, and contrast thresholding (described in sections 3 and 4 of the Lowe et al. 2004), whereas the descriptor is calculated in a straightforward manner. In this problem, we will give you a keypoint detector and ask you to create the descriptor. Please download `PS3Prob2.zip` from the course webpage to complete the rest of this problem.
 - (i) The keypoint descriptor of SIFT is created by first computing the gradient magnitude and orientation at each image sample point in a region around the keypoint location, then collecting the results into orientation histograms. Please refer to Section 6.1 Lowe's 2004 Paper for details of implementation. Download `PS3Prob2.zip` from the




Figure 1: Different metrics for stable interest points of a given image




Figure 2: Diagram of keypoint descriptor formation. The blue circle is used to indicate the presence of a gaussian centered at the keypoint.

course webpage and examine `sift_demo.m`. The interest point detector is given to you, and it is up to you to create the 128-length (4x4 array of histograms with 8 orientation bins per histogram) feature vector descriptor for each keypoint. Test the results your of your full SIFT pipeline with the given test image `elmo.jpg` and the given `sift_demo.m` file.

- (ii) Consider the situation where you are trying to detect a rigid object that has a consistent orientation across object instances (Refer to Fig. 1 for a practical example). What change to the orientation assignment of the original SIFT algorithm could improve detection of the specified object? *hint: what invariances does SIFT have and what are the tradeoffs?*




Figure 3: In the example problem of detecting pedestrians with a stationary camera, it is a safe assumption that they will always be of similar orientation.

3 Single Object Recognition Via SIFT

In his 2004 SIFT paper, David Lowe demonstrates impressive object recognition results even in situations of affine variance and occlusion. In this problem, we will explore a similar approach for recognizing and locating a given object from a set of test images. It might be useful to familiarize yourself sections 7.1 and 7.2 of the paper, which can be found in class website under the reading for February 10th. The code and data necessary to solve this problem can be found in `PS3Prob3.zip` on the course webpage.

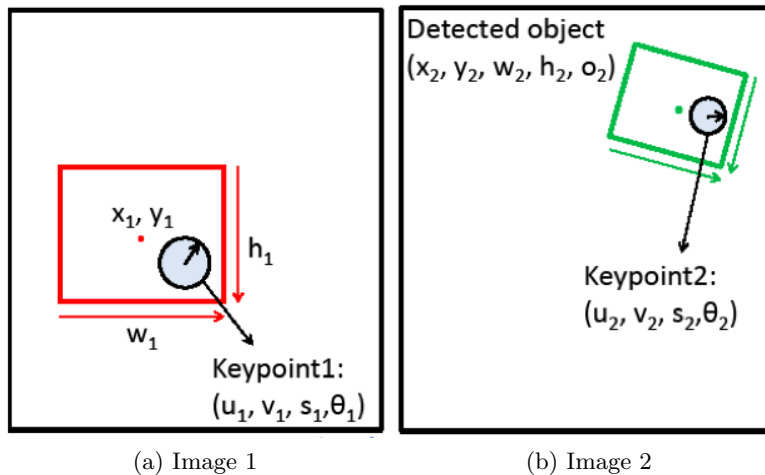


Figure 4: Sample Output, showing training image, keypoint correspondences and resulting bounding box

- (a) Given a keypoint descriptor g from one image and a set of keypoint descriptors from a second image $f_1 \dots f_n$, write the algorithm and equations to determine which keypoint in $f_1 \dots f_n$ (if any) matches g . The stub function `matchKeypoints.m` is given to you: implement this matching algorithm in MATLAB and test its performance using the `matchObject.m` skeleton code. Load the data in `PS3_Prob3.mat` and run the system with the following line:

```
>>matchObject(stopim{1}, sift_desc{1}, keypt{1}, obj_bbox, stopim{3}, ...
    sift_desc{3}, keypt{3});
```

Note that the SIFT keypoints and descriptors are given to you in `PS3_Prob3.mat` file. Your result should match the sample output in Fig. 4. Be sure to turn in your code and a sample image similar to Fig. 4



(a) Image 1

(b) Image 2

Figure 5: Two sample images for part (b)

- (b) Suppose that you have matched a keypoint in the object region of the first image to a keypoint in a second image (see Figure 5 above). The keypoint in the first image is associated with a bounding box of a particular width, height and center position relative to the location of the keypoint. Show how to compute the predicted center position, width,

height, and relative orientation $(x_2, y_2, w_2, h_2, o_2)$ of the object in Image 2 assuming that the relation holds across rotation, translation and scale. For example, you can define o_2 in terms of θ_1 and θ_2 and w_2 in terms of w_1, s_1 and s_2 . Once you have defined the five features of the new bounding box in terms of the two keypoint features and the original bounding box, describe how you would utilize the Hough transform to determine the best bounding box in Image 2 given n correspondences.

- (c) Implement the function `getObjectRegion.m` to recover the position, scale, and orientation of the objects (via calculating the bounding boxes) in the test images. You can use a coarse Hough transform by setting the number of bins for each dimension equal to 4. If you are not able to localize the objects (this could happen in two of the test images), explain what makes these cases difficult.

4 Single Object Matching Via Shape Context

Depending on a given problem, some feature descriptors may be better suited than others for object recognition. In contrast to other detectors studied in class, Shape Context [Belongie et al 2002] measures similarity between shapes and exploits it for classification. The methodology also calculates a transformation that would maximally align two potential matches.




Figure 6: Visualization of polar-coordinate binning

- (a) One natural application of Shape Context is to match handwritten digits. In this problem we have provided the data, interest point detection, and matching code for you.
 - (i) Write a function `compute_shape_context.m` to generate the shape context feature to be used when matching. Specifically your function should take in the minimum radius, maximum radius, number of bins for radius, number of bins for angle, the input data from the edge detection, the mean radius and outliers. Your code should output the mean radius and the shape context feature for each data point. Write your code in `compute_shape_context.m`. In this file there is a description of the input and output data structures. This code will take in data from the interest point detector and output data to be used in the matching process for shape context. Run the `test_shape_context.m` script to verify your implementation. Turn in your code for `compute_shape_context.m`.

- (ii) Run the `compare_digits.m` script and turn in the error for each match. The error for each match we calculate the error for you and print it out to a figure after running the `compare_digits.m` script.




Figure 7: Example data from handwritten digit dataset

- (b) Now we consider different images and matching scenarios other than handwriting to test how shape context performs.
 - (i) We will use preprocessed versions of the images in Fig 8 to test your shape context descriptor. Run the `extend_shape_context.m` script. Turn in the resulting error values and indicate what properties of shape context yield this performance.




Figure 8: Other test images for shape context

- (ii) Considering the invariances of Shape Context when would you expect Shape Context descriptor to have poor performance? Give specific circumstances citing the invariances of Shape Context, and use the previous football result as evidence. Given that Shape Context performs well for the digit data, when else would we expect shape context to perform well.