

Intro to Computer Vision: Quiz 3

Alvin Lin

August 2017 - December 2017

Question 1

We are interested in finding key points in this image in a scale and rotation invariant way. We can do this by taking the Laplacian of Gaussian on the image in order to find regions of high gradient. These regions identify significant edges and corners on the image to use as keypoints. We can obtain the orientation of the keypoint by calculating the orientations of all the pixels around the keypoint and selecting the most frequently occurring orientation (using a histogram to bin all the orientations). To uniquely identify this keypoint, we take a 16x16 window of pixels around the keypoint, broken up into a 4x4 grid. For each 4x4 window in the 4x4 grid, we calculate gradient magnitudes and orientations, which give us 128 features which uniquely identify this keypoint. To make this rotation invariant, we subtract the keypoint orientation obtained earlier. By doing this process across two images, we obtain a series of rotation invariant fingerprints for keypoints, which we can compare for similarity. If clusters of keypoints match across both images, we can conclude those regions of the image are of the same scene.

If you have any questions, comments, or concerns, please contact me at alvin@omgimanerd.tech