CS 510, Spring 2019, First Half of Semester Review
March 11, 2019

The following is not an exhaustive list of what has been covered, and to be very explicit, the First Midterm may cover material discussed in class not mentioned below. That said, hopefully the following will be helpful as you review the broad topics and techniques we have covered so far in CS510.

1. The semester began with an overview of the human vision system in human visual processing. From this overview, there are a handful of take-home messages you should be comfortable discussing. These include the fovea, saccadic eye movements, the placement of primary visual cortex, and the concept of a retina topic mapping.
2. Illusory contours, as evident in such famous examples as the Kanizsa Triangle, teach us much about the logic of human visual perception. Be ready to recognize this concept and address is implications relative to physics and psychology.
3. The second lecture dealt with visual attention. From that lecture, you should be very comfortable the distinction between overt and covert tension. You should also be able to differentiate between the three theories of covert attention presented in that lecture. Finally, in terms of overall significance, one of the most important take-home messages from that lecture concerns human visual cognition and the concept of "inattentional blindness".
4. This semester we have dealt with many pragmatic issues associated with using open CV. For example, you're comfortable with code that can create an image that is small and may be properly thought of as a window into a larger image.
5. Applying geometric transformations to images is very simple in open CV. You can comfortably work with code that rotates, translates, scales and even warps images. You can take it a step further and actually deal with a non-affine transformation that performs the equivalent of a perspective projection change in view for a planar object, i.e. an image.
6. In CS5 10, you're very unlikely to be asked about the Fourier transform in the context of single and double integrals. Yes, their continuous domain definition is best expressed in such a manner, but instead you're more likely to be asked questions about before and after representations of 1D and 2D discrete signals. However, it is important to understand the relationship between phase and the use of complex numbers when representing a Fourier transform.
7. When operating upon images it is very common to apply linear operators to images. For example, principal components analysis falls into this category. So does the discrete Fourier transform, and this is an aspect of linear transformations that you now understand and can explain to others. If this seems vague, be ready to describe the basis vectors that go along with the Fast Fourier Transform.
8. It may seem like a nitpick, but be clear on exactly how many complex numbers it takes to fully encode the discrete Fourier transform of a 64 x 64 pixel image.
9. A variety of examples were presented the semester showing how to manipulate an image through changes to its representation in the frequency domain. This included "removing noise" when the noise happened to have a very distinctive periodicity. Be comfortable recognizing some of these very handy uses of the frequency domain.
10. The Nyquist Rate is a very important concept touched upon in lecture. Review this concept and be sure you understand it along with the related concept of aliasing.
11. In edge detection the distinction between and abstract continuous image surface and a physically realizable discrete image come to the forefront. You are now comfortable with both representations and in particular the mapping between the clean definition of a gradient on an analytical function and the approximations that must arise when working with real images.
12. The two concepts of edge detection and smoothing go hand-in-hand throughout the history of computer vision. Be ready to explain the Sobel edge operator as the result of combining a difference operator and a smoothing operator. If the word operator is too vague, replace it with "a one-by-three pixel mask".
13. At least one full lecture was devoted to different ways to capture the similarity between two images being a bit more specific, much of that discussion applied to any pair of vectors. The presentation included different distance

norms and also correlation. These mathematical definitions are so utterly essential to pattern recognition in general and computer vision in particular that you should have them committed to memory.

14. An entire lecture was devoted to template matching. It is easy today to overlook the practical value and simplicity of template matching. Be clear on what is meant by template matching as well as the kinds of problems for which it is a good solution as well as the great many problems where can be expected to fail.
15. Since you used the template matching function built into open CV, you can recognize visually, and also relate back to fundamental mathematical definitions, the five distinct ways of measuring similarity between a template and an image. To clarify, not all of these distinguished when viewed in a screenshot, but some can and you should give thought to which pairs of alternative similarity measures do appear notably different when displayed.
16. Our lab at CSU is justifiably proud of the work of David Bolme. We spent an entire lecture talking about the ASEF and MOSSE algorithms as well as the tracker that was built on top of the basic correlation filter construction algorithm. You can recognize and interpret every significant step in both of these algorithms with perhaps one exception. That exception being the precise least-squares approach taken in MOSSE relative to the more brute force averaging used by ASEF.
17. Principal components analysis plays an utterly fundamental role in characterizing data, reducing the dimensionality of data, and even in recognition including the dawn of practical human face recognition algorithms. There are at least two very distinct ways this course has sought to develop an intuition as well as a practical understanding of PCA. Make sure you are comfortable with both.
18. Some researcher claim, mistakenly, that when using PCA the important information is always conveyed by the dimensions with the largest Eigenvalues. Be able to speak to this misconception in a very clear and practical way.
19. The Hough transform is taught in this course as a powerful example of a voting scheme that yields relative robust output. As this idea is so fundamental you should be comfortable with all aspects of the Hough Transform.
20. The Canny Edge Detector is a truly seminal bit of work in computer vision and you should be able to recite from memory the key contributions made by John Canny in his development of this algorithm.