Video: Edge Detection Overview

[00:01] [slide 1] In this first video about edge detection, we focus on general ideas behind an approach that's especially relevant to understanding the early stages of processing in human vision. In a separate video, you'll learn how these ideas are implemented in computer vision systems.

Visual processing begins with an image, like this image of coins - you can think of the image as an array of measurements of the brightness that's reflected from surfaces in a scene and measured by a camera or the eye. Important physical changes in the scene give rise to variations in the light intensities in the image. Here, the physical changes include things like, the borders of the coins, the engravings on their surfaces, the highlights around the rims of the coins, and they also cast a very narrow shadow against the background surface. For a tiny part of the image in this red box here, the grid of brightness measurements on the right shows the shade of gray for each picture element or pixel, and the numerical intensity is on a scale from 0 to 255, corresponding to the range from black to white.

These variations in brightness provide the first clue about objects in the scene, so our first goal in processing an image is to detect and describe the important intensity changes taking place in the image. But looking at this snippet here, intensity is changing everywhere, it's hard to find two adjacent pixels with the same brightness, so how do we distinguish the most important changes and pinpoint where they occur? How do we determine whether they're high-contrast, sharp changes like we see along the border of each coin, or small, gradual changes of brightness like we see in the engravings on the faces of the coins? We'll use an approach to detect and describe intensity changes that embodies three basic operations: smoothing of the image intensities, taking a derivative to measure change, and extracting features in the result of this processing.

[02:22] [slide 2] We'll start by looking at these operations in one dimension. Imagine that we take a horizontal slice through a digital image and then show the brightnesses there as a plot, for example, this plot here at the top of the picture, in blue. There are some small fluctuations of intensity that are likely due to noise in the camera's sensors, and we can just ignore those. Changes can also be happening at different scales. In the case of the coins image, at a coarse scale that we'd see if we stood very far away or blurred the image a lot, we'd just see a bunch of light circles on a dark background. At a finer scale, we see the details of the engravings on the coins. In general, if we smooth out the intensity profile by a small amount, as shown in the middle here, we'd get rid of the minor fluctuations due to noise, but still preserve the fine details. If we smoothed it more, we'd only see the changes taking place at a coarse scale. So the purpose of smoothing is to reduce the effects of noise and set the scale at which we want to describe the changes of intensity.

[03:40] [slide 3] So let's say we now have our smoothed intensity profile as shown on the top in this slide, how do we identify where the significant intensity changes are taking place? We still have the challenge that brightness is changing everywhere, from one location to the next. The next step is to measure these changes by taking a derivative. The first derivative of this signal, shown in the middle here, indicates how much the signal changes from one location to the next. So the first derivative, if the intensity change is increasing from left to right, then the first derivative is positive. If intensity is decreasing as we move from left to right, then the first derivative will be negative. Around the middle of these changes, highlighted, for example, with these three black circles, if we follow down the dashed lines here, you see that the first derivative actually hits a peak around the middle of that change. The peak might be positive if the intensity is increasing, or if it’s decreasing, the peak would be negative. We can also measure another derivative, a second derivative of the original signal, and here, at the locations of the intensity changes where there was a peak in the first derivative, the second derivative crosses zero, as shown by the green circles. It might cross zero going from positive to negative values, or negative to positive values, depending on the direction of change in the original intensity profile. It turns out that it’s fairly easy to locate these peaks in the first derivative, or zero-crossings in the second derivative, and we can use these features as a way to pinpoint where the significant intensity changes are occurring in the original intensity profile.

[05:54] [slide 4] Now how do these ideas generalize to a two-dimensional image? We’ll start with an image of this iconic Wellesley lamppost. In two dimensions, we’ll again smooth the image a bit and we’re going to compute a kind of second derivative that mimics the processing that takes place in the human eye. The picture in the upper right corner shows the result of this computation. Like we saw in one dimension, the values here can be both positive or negative, and the very positive values are shown as the bright parts of this result here and the negative values are shown darker in this depiction here. In the bottom left, all the positive values from this result are displayed as white and all the negative values are displayed as black. The transition between the two are the places where it’s changing sign, and those are the zero-crossing points that I mentioned in the case of one dimension. They correspond to roughly where the edges are in the original picture, for example, around the lamppost here and the lamp itself. Some of these contours appear darker than others, for example, on the lamppost and within the lamp, and we’ll see that what I’m portraying there is, the darker contours are very high-contrast, sharp changes taking place in the original image, and the lighter contours, for example, some of the branches in the background, those are lower-contrast edges in the original image.

[07:50] [slide 5] I also said that we can smooth the image by different amounts to capture changes taking place in the image at different scales. I’ll use a different example for this - an image of a Handi-Wipe cloth, shown as a black and white image. At a fine scale, there are intensity changes around the individual holes in the cloth, but at a coarse scale which you might see if you stood far away or blurred the image out, mainly there’s the light and dark stripes. If we smooth the image by a small amount, we can preserve the changes corresponding to the little holes in the cloth here, as shown on the left, and if we smooth by a larger amount, then we lose the holes and we just see the stripes.

[08:42] [slide 6] Now you may be wondering, does this process of locating significant intensity changes have any relevance to human vision? To give you just a taste, we'll look at the Craik-O'Brien-Cornsweet illusion. Now you might say, what illusion? I clearly see two light gray circles on a dark gray background. But that's not what's actually here. Imagine we take the brightnesses along a horizontal strip in the middle of this image and plot them as a graph. I'm showing this below here. To give you a reference, I'll add this red line at the brightness of the dark gray background. On the right, the inside of the circle really is brighter than the background, and that's shown in the profile here. But something different is actually shown on the left. It starts out the dark gray background, but then as we approach the edge of the circle, it actually gets darker here, and then right at the edge, it suddenly jumps up, but then gets darker again as we approach the center. In fact, the center of this circle is that same dark gray as the surrounding region, but we don't see it that way. We're especially sensitive to these sharp changes of intensity at the borders, and nearly blind to gradual changes taking place on either side of the edge in this case. So intensity changes also play a key role in human vision as well.

[slide 7] That's a broad overview of the first stage of visual processing, that we refer to as "edge detection." How do we actually compute this information from a digital image like we saw at the start of this video? The next video jumps into the details of this process and introduces a generic type of computation that we refer to as *convolution* that's useful for the edge detection stage. Later, we'll see that this kind of computation is also performed in biological systems, but in a way that's very different from how it's done in computer vision systems.