CS 280 Computer Vision
Spring 2018 Efros, Malik, Yu Lecture 10

1 Lecture

Edges, Templates, Textures

Gradient magnitude indicates where the changes happen; it doesn't care about orientation.

There's a tradeoff between smoothing and good edge localization. We want to smooth, because it removes noise that would otherwise confound the derivatives, but smoothing will also blur the edges. This is where the *Canny edge detector* comes in. The basic idea: we first threshold the norm of the gradient and remove the weaker transitions (with smaller gradients). Then we want to turn the thick regions of the gradient into curves. For this we can use non-maximum suppression. We retain only the pixels that are maxima along the gradient direction by choosing *along the gradient direction* only the pixels which are greater in gradient magnitude than their neighbors. If not retained, a pixel will have its magnitude set to 0.

There is one problem, though. Some pixels along an edge, which we want, might not survive thresholding and non-maximum suppression. So the second trick of the Canny edge detector is *hysteresis thresholding*: we will first use a high threshold to start edge curves (high gradients will survive this non-maximum suppression). Then we use a low threshold to continue them, those strong edges we identified initially (we will, in effect, *grow* the edge).

In this way, weak pixels need to be attached to a strong pixel in order to be considered. And we end up with less noise, at least as isolated pixels that would have survived ONLY a low threshold.

The edges we're interested tend to be continuous, and this is what hysteresis thresholding captures!

In summary, the Canny edge detector undergoes the following sequence of steps:

1. Filter the image with the derivative of a Gaussian.
2. Find the magnitude of the gradient.
3. Find the orientation of the gradient, which will give us the gradient direction at each pixel so we can...
4. Perform non-maximum suppression, and thereby thin wide "ridges" down to a single pixel width.
5. Perform thresholded linking: use a high threshold to start edges and a low threshold to continue them.

Smoothing vs. Derivative Filters

The values of a smoothing filter (e.g. a Gaussian) should sum to 1 so that constant regions are not affected by the filter. The values of a derivative filter (e.g. derivative of a Gaussian) should sum to 0 so that there is no response in constant regions.

A derivative filter will yield high absolute values when applied to points of high contrast.

Template Matching

Many methods for template matching exist. The *zero-mean filter* (just the template with its mean subtracted out) is the fastest option but a subpar matcher. SSD is the next fastest, but is sensitive to overall intensity. NCC is slow but achieves superior results, as it is invariant to local average intensity and contrast. *Note: using the template as-is (as a filter) never works. It just ends up being a smoothing filter.*

Object Recognition

In object recognition, there's a distinction between instance recognition and category recognition. In *instance recognition*, the task is to find a particular object (e.g. find a particular chair again). Often, simple template matching does okay at this.

In *category recognition*, the task is to find all objects of a certain type (e.g. find all chairs). Templates don't work for this, because a template describes an instance which doesn't describe the "essence" of the entire category on its own. We need to focus on things that might be invariant across the entire category.

Texture

Texture is characterized by a spatially repeating pattern. Many natural phenomena manifest as textures, like rocks or piles of fruit.

How do we compare textures and decide whether they're made of the same "stuff?" The pixel values might be different, but humans can still tell if textures come from the same "stuff." This is evidenced by Julesz's studies of random dot stereograms and texton discrimination. As Julesz found, human vision is sensitive to certain types of differences between elements (even if subtle) but not others.