Image Gradients

Class Notes for CMSC 426, Fall 2005

David Jacobs

Introduction

The gradient of an image measures how it is changing. It provides two pieces of information. The magnitude of the gradient tells us how quickly the image is changing, while the direction of the gradient tells us the direction in which the image is changing most rapidly. To illustrate this, think of an image as like a terrain, in which at each point we are given a height, rather than an intensity. For any point in the terrain, the direction of the gradient would be the direction uphill. The magnitude of the gradient would tell us how rapidly our height increases when we take a very small step uphill.

Because the gradient has a direction and a magnitude, it is natural to encode this information in a vector. The length of this vector provides the magnitude of the gradient, while its direction gives the gradient direction. Because the gradient may be different at every location, we represent it with a different vector at every image location.

Computing the Gradient

We'll begin by showing how to compute the gradient at an image location. Then we'll show that the thing that we compute actually does encode the gradient direction and magnitude. We form the gradient vector by combining the partial derivative of the image in the *x* direction and the *y* direction. We can write this as:

$$\nabla I = \left(\frac{\partial I}{\partial x}, \frac{\partial I}{\partial y} \right)$$

To explain this, let's review partial derivatives. When we take the partial derivative of *I* with respect to *x* we are determining how rapidly the image intensity changes as *x* changes. For a continuous function, *I(x,y)* we could write this as:

$$\frac{\partial I(x,y)}{\partial x} = \lim_{\Delta x \to 0} \frac{I(x+\Delta x, y) - I(x,y)}{\Delta x}$$

In the discrete case, we can only take differences at one pixel intervals. So we can take the difference between *I(x,y)* and the pixel before it, or the pixel after it. Or, as we discussed in the class notes on correlation and convolution, in the section on *Taking Derivatives with Correlation* we can treat the pixels before and after *I(x,y)* symmetrically, and compute:

$$\frac{\partial I(x,y)}{\partial x} \approx \frac{I(x+1, y) - I(x-1, y)}{2}$$

By similar reasoning, we can also compute:

$$\frac{\partial I(x,y)}{\partial y} \approx \frac{I(x,y+1)-I(x,y+1)}{2}$$

Combining these with the first expression in this section gives us a complete method for computing the image gradient.

Properties of the Gradient

As we have computed it, the gradient seems to be a vector that describes how quickly the image changes when we move in either the *x* or the *y* direction. But what happens when we move in other directions? When we use derivatives to answer a question like this, it means that we are considering how the image changes with a *very small* movement in position. We will see that the gradient allows us to answer this question.

Let's consider how the image changes as we move from position *(x,y)* by a small amount, Δ, in an arbitrary direction θ, which will take us to the position *(x+ Δ cos* θ*, y +* Δ *sin* θ*).* At *(x,y)* the image intensity is *I(x,y)*. First, from calculus, we know that if Δ is small then we will have:

$$I(x+\Delta\cos\theta, y) \approx I(x,y) + \Delta\cos\theta\frac{\partial I(x,y)}{\partial x}$$

That is, as we move a little in the *x* direction, the change in intensity is the distance we've moved times the derivative of the intensity with respect to the *x* direction. This is exactly what the derivative is for, to tell us how intensity changes as we move a little. (We could also say we've taken the Taylor series expansion, and dropped higher order terms). By the same reasoning, we have:

$$\begin{aligned} I(x+\Delta\cos\theta, y+\Delta\sin\theta) &\approx I(x+\Delta\cos\theta, y) + \Delta\sin\theta\frac{\partial I(x,y)}{\partial y} \\ &\approx I(x,y) + \Delta\cos\theta\frac{\partial I(x,y)}{\partial x} + \Delta\sin\theta\frac{\partial I(x,y)}{\partial y} \end{aligned}$$

This derivation assumes that the derivative of *I* doesn't change much if we move a small amount in the image.

We can rewrite this formula. First, we define:

$$v \equiv (\Delta\cos\theta, \Delta\sin\theta)$$

v is a vector that encodes how we've moved in the image. Then we can rewrite the previous formula as:

$$I(x+\Delta\cos\theta, y+\Delta\sin\theta) - I(x,y) \approx \langle v, \nabla I\rangle$$

This tells us that the change in intensity as we move by a small amount in the image can be found by taking the inner product between the gradient and a vector that describes the movement.

This way of looking at things allows us to answer some questions, using the important formula for the inner product:

$$\langle v, w \rangle = \|v\| \|w\| \cos \alpha$$

where α is the angle between the two vectors, *v* and *w*.

For example, we would like to know what direction we can move in so that the image changes as much as possible, that is, what direction is uphill? We know that for different choices of *v*, $<v, \nabla I>$ is proportional to the cosine of the angle between *v* and ∇I, and will have its maximum value when this angle is 0, when *v* points in the direction of ∇I. So the direction of ∇I is the direction in which the image changes most. Also, we know that if *v* is a unit vector in the direction of ∇I, then the amount the image intensity will change will be: $\|\nabla I\|$, because $\|v\|=1$ and $\cos \alpha = 1$. Therefore the magnitude of the gradient tells us the amount that the image intensity is changing when we move in the direction of greatest change.

A second interesting point is that if *v* is perpendicular to ∇I then the inner product between the two is zero. This means that the image intensity doesn't change at all when we move in that direction.

Conclusions

There are two things we're trying to convey with this note. The first is a method of computing the image gradient. This is done by computing the image derivative in the *x* and the *y* directions, and then combining these into a vector. The second point is that the direction of this vector is the direction in which the image is changing most rapidly, and the magnitude of this vector is the rapidity with which the image is changing in this direction. It is not obvious (at least to me) that these two ideas would be connected. It is interesting that by taking image derivatives in two, essentially arbitrary directions, we wind up with a vector that tells us where and how quickly the image changes, in such a direct way.

The image gradient is important in boundary detection because images often change most quickly at the boundary between objects. In class we will see how to use the gradient to find edges, which we define as places where 1) the magnitude of the image gradient is high; and 2) it is higher than the gradient magnitude at other locations that are in the direction of the gradient.