

A Brief Introduction to ICM

Jerod Weinman

August 16, 2012

Background

Images get corrupted by many things, from noisy transmissions (e.g., from a Mars rover to NASA) to dirty scanner glass and camera lenses. In such situations, we want to restore the image to its original or proper state as best we can. One method for doing this is called “Iterated Conditional Modes.” [1] We won’t go in to all of the rich theory behind it (c.f. Li [2]), but the name does tell us two things. It’s iterated, so it will probably involve a loop. And by “mode” we mean finding a maximum (or actually a minimum, in our case).¹ In short, we’ll be iteratively optimizing something within a loop, and that something is our best guess of what a corrupted image pixel should be.

In this lab, we are going to be concerned with binary images. That is, images whose pixels are either on or off, black or white. Thus, the only noise that is possible is whether a black pixel (or bit) was flipped to a white one, and vice versa.

When these images are *huge* (very high resolution telescope scans of the sky, or Google-scale satellite images of the earth), distributed parallel processing can make the job much faster.

Image representation

An image is thought of as a two dimensional array, where each pixel is indexed by its row and column. In C code, this is often something like

```
image[i][j]
```

which refers to the element of the image at row i and column j , which we’d call pixel (i, j) . To look at the left neighbor of the pixel, we change the column index by one:

```
image[i][j-1]
```

Now, image folks like to view the world from the top-down. Therefore, rows of the image start at the top, and work their way down. This is in contrast to the usual Cartesian coordinate system you’re using to graphing where “0” on the y axis is at the bottom-left. Remember also that C indexing starts at zero. Making all this concrete with some code,

```
image[0][5]    // the TOP row, and the sixth column
image[M-1][3]  // if there are M rows, this pixel is in the bottom row
image[i-1][j]  // this is the upper neighbor of (i,j)
image[i+1][j+1] // this is the lower-right neighbor of (i,j)
```

Hopefully that’s enough to get you started with the simple data structures we’ll be using.

Oftentimes, it is more convenient to abstract from the two dimensional indexing and think of a pixel as a single index $p = (i, j)$. If \mathbf{x} represents the entire image, then we can go back and forth between the mathematical construct/notation, and the program code notation:

$$\text{image}[i][j] \leftrightarrow x_p. \tag{1}$$

As it turns out, this will make it much easier to describe the ICM algorithm using mathematical notation.

¹The use of “mode” is a bit of a confusing appropriation of a term from statistics (meaning the most frequent observation) that can be interpreted as finding a maximum. It would be better if the M in ICM stood for maximization (or minimization).

ICM motivation

This image restoration algorithm makes two basic assumptions, one based on the contents of images in general, and another based on the noise process, having nothing to do with image properties. The first assumption about images says:

Assumption 1 Neighboring pixels tend to have the same values.

Think about it. Images are made up of regions corresponding to objects with large chunks that tend to have roughly the same color and brightness. In a binary image, there will be a lot of black regions and a lot of white regions. Only the borders between these regions violate the assumption, but we'll just have to hope that isn't a big deal. The power behind this assumption is that if we see a white pixel in a whole sea of black pixels, we might reasonably conclude that it was corrupted by noise, accidentally flipped, and that it too should probably be a black pixel. (The reverse situation would also hold.)

The second assumption about the noise process says:

Assumption 2 Each pixel is corrupted (flipped) independently, and with some probability.

This assumption says (rather straightforwardly) that the corruption, or lack thereof, of any two pixels is not dependent in any way. Thus, if one pixel is flipped by a noisy transmission line, it has no bearing on whether any of its neighbors are flipped. This covers the first part of the assumption. The second part simply says that there is some prior knowledge about how likely it is a pixel is flipped. A 100% probability of flipping is not exactly noise, because *every* pixel would be flipped, producing an inverse image. A 10% probability that a pixel would be flipped is a moderately noisy, but recognizable image. A 50% probability would be an utter disaster, because every pixel would be essentially random. So long as there is a moderate amount of noise, we can create a process using these assumptions that does a pretty good job of fixing up the image.

ICM implementation

Once again, there is some very fascinating theory behind the image models that are derived from the above assumptions. Because this is a systems class, we're going to focus only on the basic algorithm, in order to reiterate that topics like parallelization are important in a wide variety of computational domains.

One important aspect of the ICM algorithm is the "pixel neighborhood." Recall that assumption 1 states that neighboring pixels have the same value. A pixel must therefore ask, "who is my neighbor?" The typical answer to this question is the 8 pixels surrounding $p = (i, j)$. The two dimensional indices of these eight neighbors are shown below.

$(i-1, j-1)$	$(i-1, j)$	$(i-1, j+1)$
$(i, j-1)$	(i, j)	$(i, j+1)$
$(i+1, j-1)$	$(i+1, j)$	$(i+1, j+1)$

Note that pixels on the outside edges of the image will not have all of these neighbors. For instance, the top row of the image will not have the upper row of 3 neighbors. Corner pixels only have 3 neighbors.

We will refer the set of neighbors for a given pixel $p = (i, j)$ as $\mathcal{N}(p)$. This will make it easy to refer to properties of the neighborhood of p . For instance, if $x_p \in \{0, 1\}$ where 0 means a pixel is black, and 1 means a pixel is white, the total number of white pixels surrounding p is given by the sum

$$\sum_{q \in \mathcal{N}(p)} x_q. \quad (2)$$

Now we're finally ready to get to the algorithm. Our goal is to estimate the original, uncorrupted image \mathbf{y} from the observed noisy image \mathbf{x} . To do this, the algorithm iteratively minimizes a *cost* of the reconstruction of each pixel. We'll do this iteratively, so we represent the restored image at iteration k as $\mathbf{y}^{(k)}$. Updates $\mathbf{y}^{(k+1)}$ are repeatedly calculated from the current restored image $\mathbf{y}^{(k)}$.

The two assumptions above are reflected in two components of an additive cost. First, due to assumption 2 there is a penalty if the reconstruction flips a bit from the observation. This value corresponds to the probability of a bit being flipped from noise. A *large* penalty for each flipped bit in the restoration corresponds to a *small* probability of corruption. A second penalty term, due to assumption 1, is paid for every pair of neighboring pixels that are different.

Once again, this promotes smoothness of the image. The larger the penalty, the smoother and more uniform we are asking the reconstruction to be. These two penalty terms, which we will call α and β , interact to determine how the image is reconstructed.

Now, on to the heart of it. Given the current reconstruction $\mathbf{y}^{(k)}$, we can define the cost of giving a particular pixel p either value $y_p \in \{0, 1\}$

$$C(y_p) = \alpha(1 - \delta(y_p, x_p)) + \beta \sum_{q \in \mathcal{N}(p)} (1 - \delta(y_p, y_q^{(k)})) \quad (3)$$

The function $\delta(a, b)$ is known as the “Kronecker delta function” and has the simple definition

$$\delta(a, b) = \begin{cases} 1 & \text{if } a = b \\ 0 & \text{if } a \neq b. \end{cases} \quad (4)$$

The meaning of the cost $C(y_p)$ is given intuitively above. The first term corresponds to assumption 2 and says “add α if the reconstructed pixel and the observed (original) pixel don’t match.” The second term corresponds to assumption 1 and says “add β for every neighbor in the current reconstruction that doesn’t match y_p .” We then assign an updated reconstruction by choosing each pixel value as the one that minimizes the cost

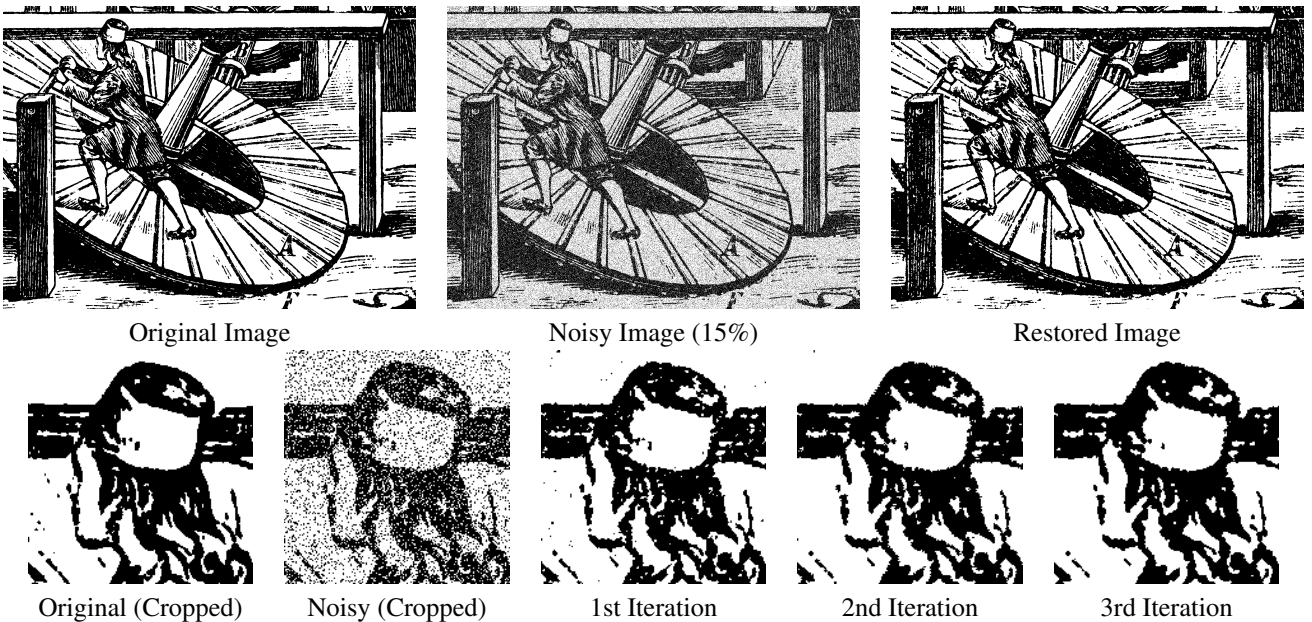
$$y_p^{(k+1)} \leftarrow \arg \min_{y_p \in \{0, 1\}} C(y_p). \quad (5)$$

Note that all of the reconstructed pixels are updated synchronously. In other words, we must use values from the previous reconstruction $\mathbf{y}^{(k)}$ for calculating every updated reconstructed pixel in $\mathbf{y}^{(k+1)}$. The algorithm converges when the updates cease to change the reconstruction,

$$\mathbf{y}^{(k+1)} = \mathbf{y}^{(k)}. \quad (6)$$

Sometimes, though, the algorithm tends to oscillate between a few states,² so an iteration limit on k is needed.


Below is an example original image, a corrupted version of the image, and a few steps of the iterations for restoring it with $\alpha = 2$ and $\beta = 1$.



²This oscillation can be easily detected, but you won’t be asked to do that.

References

- [1] J. Besag. On the statistical analysis of dirty pictures (with discussions), 1986.
- [2] S. Z. Li. *Markov Random Field Modeling in Image Analysis*. Computer Science Workbench. Springer-Verlag, Tokyo, second edition, 2001.

Original treadmill image is in the public domain. Other text and images are copyright ©2008, 2012 Jerod Weinman.  This work is licensed under a Creative Commons Attribution-Noncommercial-Share Alike 3.0 United States License.