

Imaging Science

Lost in the Supermarket: Decoding Blurry Barcodes

By *Todd Wittman*

Linear barcodes—the ubiquitous alternating black and white stripes whose relative widths encode information—are used in shipping, tracking, and identification and come in more than 250 varieties. The most familiar is probably the UPC barcode used in supermarkets; based on their grocery-shopping experience, most readers probably consider the reading of barcodes a trivial problem.

Indeed, when I presented a poster on this topic in May at the SIAM IS04 conference in Salt Lake City, the most common initial response was: “Huh. I didn’t think this was a problem.” Certainly, the problem as it appears in the checkout lane has been solved. The standard countertop laser scanner is equipped with mirrors that obtain signals from multiple angles. Handheld laser scanners are by necessity less powerful and less accurate—they are not capable of gathering multiple signals, and the optical devices cannot be as large and unwieldy as those of their countertop counterparts. Nevertheless, barcodes have found many applications outside the grocery store.

Shipping companies like UPS and the US Postal Service have developed their own coding schemes for package identification and tracking. Many states issue driver’s licenses with two-dimensional barcodes that encode information about the driver and a rough sketch of the photograph. Some hospitals place barcodes on every drug container, syringe, and even patient. Entomologists glue tiny barcodes to the backs of flying insects and strategically place scanners around hives to track insect flight patterns and mating habits. Some libraries and warehouses are experimenting with inventory robots, placing barcodes not only on the items but also on the walls, where they give the robots encoded directions for navigating obstacles. As the use of barcodes becomes more sophisticated, there is a need for smaller and more accurate barcode scanners, as well as superior decoding algorithms.

In 1999, Ruben Muniz, Luis Junco, and Adolpho Otero described a striking example of the current limits of barcode decoding [1]. They looked into the large pharmaceutical warehouses that distribute drugs to local drugstores in most metropolitan areas; a pharmacy places a request to the warehouse by sending an order form with three barcodes, representing the patient, medicine, and price. The authors visited their local warehouse in the Spanish province of Asturias, which receives about 1.5 million orders every month. The order forms are usually printed on low-resolution dot-matrix printers; doctors’ sprawling signatures frequently cover at least part of the barcodes. The automated scanner was able to decode only 50% of the orders, forcing human intervention 750,000 times a month. Using a Hough transform, the authors were able to improve the decoding rate to 85% of orders; even so, hundreds of thousands of unreadable order forms remained.

Rather than trying to improve the scanner’s optics, Fadil Santosa (my thesis adviser) suggested an alternative approach: Look for ways to improve the scanner’s deblurring algorithms. I must admit that I was skeptical when he first proposed this research problem. Foreshadowing the response to my poster at IS04, I asked, “Is this really a problem?” Fortunately, his research instincts are far superior to mine, and my work on barcodes has uncovered several interesting avenues of research in signal and image processing. (This work also caught the attention of the R&D department at Symbol Technologies, the world’s largest manufacturer of handheld barcode scanners. I am writing this article for *SIAM News* from Symbol’s New York office, where, as a research intern, I am working on mathematical techniques for improving the resolution of imaging scanners. If a billion-dollar industry is willing to give money to a math grad student, it must be for a genuine research problem.)

The ideal barcode signal u is a one-dimensional 0–1 step signal; the signal u_0 received by a scanner is a noisy, blurred version of u . The goal is to recover u , given u_0 . We model the blurring process by

$$u_0 = G_{a,w} * u + n, \tag{1}$$

where n is the noise and $G_{a,w}$ is a Gaussian of amplitude a and width scaled by a factor w :

$$G_{a,w}(x) = a \exp\left(-\left(x/w\right)^2\right). \tag{2}$$

Because the parameters a and w are uncorrelated, G does not necessarily have unit area. The results of the convolution are similar in appearance to the signals actually encountered (see Figure 1).

Each of the parameters in (1) describes a different noise process. The additive noise n can result from electrical noise in the scanner or from source defects, such as stray marks on the paper. The amplitude a is controlled by the intensity of the scanner and the ambient light. The blur factor w models the size of the laser spot relative to the size of the




Figure 1. Top, synthetic image of a UPC barcode encoding 0123456789. Middle, the ideal signal u . Bottom, blurred signal u_0 with $a = .05$, $w = .012$, and a moderate amount of additive noise.

barcode. In particular, w increases as the scanner moves away from the barcode. In reconstructing u , we assume that the Gaussian width is at most the size of the thinnest bar.

Standard commercial decoding techniques are based on classic edge detectors, such as those that determine zero crossings of the second derivative. Such techniques use only local information and would have difficulty decoding the three bars at position 550 in the image at the bottom of Figure 1. Our global approach to solving (1) is to minimize the total-variation (TV) energy

$$E[u|u_0] = \int_{\mathbb{R}} |u_0 - G_{a,w} * u|^2 dx + a \int_{\mathbb{R}} |u'| dx. \quad (3)$$

The first term measures the fidelity or “truth” of the blurring function. The second term monitors the total variation of our reconstructed signal u . The parameter a balances the two terms and needs to be set experimentally. It is well documented that in two dimensions, the TV norm results in “blocky” or “blotchy” images. This is because the TV term controls the jumps, but not the smoothness, of u . This staircasing effect is aesthetically displeasing in two dimensions, but it is a desirable side effect in our one-dimensional problem.

In general, the blind deconvolution problem (1) is famously ill-posed. The problem is central to image processing, as evidenced by the many speakers at IS04 who focused on it, often suggesting a TV approach similar to (3). Fortunately, our knowledge of the point-spread function G and the desired u contains a mathematical ray of sunshine. Selim Esedoglu proved recently that under certain conditions, our problem is well-posed [2]. Specifically, if we assume that $w \geq 0$, a lies in a compact set, and $u = 0$ outside the barcode region, then the infimum of (3) is attained.

Using a Newton’s method approach, we experimented with minimizing (3). We had great success in recovering u with any initial guess, but it was more difficult to recover the correct parameters a and w . For correct results, we found that the initializations for a and w must be very close to the actual values. It may seem strange that the minimization can accurately recover a long signal and yet struggles with two numbers. The explanation is that the energy (3) is convex in u , but not in a and w . Some techniques are available for estimating the size of the blur kernel G using peak locations [3].

For the example shown in Figure 2, the correct u is obtained in about 6 minutes in MATLAB with a 2.4-GHz Celeron processor. This is too slow to be practical, especially in the 10-items-or-less checkout lane. But the approach is promising and could prove useful for applications in which the signal has high levels of blur and speed is not an important factor. To decode more complex barcode schemes, such as two-dimensional barcodes or the UPS bull’s-eye barcode, there is a movement in the industry to replace the optical lasers of scanners with digital cameras. Future research may involve adapting this method for one-dimensional signals to two-dimensional barcode images.




Figure 2. Minimizing the total-variation energy of the noisy signal in Figure 1 produces the solid line. This gives the correct barcode, after thresholding at 0.5. The ideal signal is overlaid as a dotted line.

References

- [1] R. Muniz, L. Junco, and A. Otero, *A robust software barcode reader using the Hough transform*, Proceedings of IEEE International Conference on Information, Intelligence, and Systems, 1999, 313–319.
- [2] S. Esedoglu, *Blind deconvolution of barcode signals*, Inverse Prob., 20 (2004), 121–135.
- [3] E. Joseph and T. Pavlidis, *Bar code waveform recognition using peak locations*, IEEE Trans. Pattern Anal. Machine Intell., 16:6 (1994), 630–640.

Todd Wittman (wittman@math.umn.edu) is a graduate student in mathematics at the University of Minnesota, where he is working under the supervision of Fadil Santosa.