Preface

Precision landmark location in digital images is one of those interesting problems in which industrial practice seems to outstrip the results available in the scholarly literature. Approaching this problem, we remarked that the best-performing commercial close-range photogrammetry systems have specified accuracies of a part per 100,000, which translates to 10–20 millipixels of uncertainty in locating features for measurement. At the same time, articles we identified in the academic literature didn't seem to give any firm answer about when this level of performance is possible or how it can be achieved.

We came to the problem of precision landmark location by a process that must be familiar to many: out of a desire to calibrate an image-metrology test bench using landmarks in well-known locations. It is straightforward to perform the calibration, minimizing errors in a least squares sense; but we were also interested to know what fraction of the residual errors should be attributed to the determination of landmark locations in the images, and whether this error source could be reduced.

To address these questions, we knew we had to go beyond the consideration of binary images. At the limit of sensitivity, it is clear that all of the grayscale information must be used. Likewise, we knew we had to go beyond an idealized model of image formation that considers the pixels to be point-wise samples of the image; since we were looking for changes to the image that arise with landmark displacements of a small fraction of a pixel width, the details of the photosensitive area within each pixel were going to play a role. Finally, rather than focusing on the performance of a particular algorithm, it seemed better to pursue the Cramér–Rao lower bound, and thereby determine an algorithm-independent answer to the question. There was, after all, a sheer curiosity to know whether part-per-100,000 accuracy is really possible for close-range photogrammetry with digital images.

In a convergence of necessity and opportunity, it turns out that considering a richer model of the image formation process is itself instrumental in making the calculation of the uncertainty tractable. The smoothing introduced into the image by diffraction and other sources has been neglected in some past investigations, which have idealized the digital image as point-wise samples from a discontinuous description of the image. Far from being an unwanted complication, the smoothing makes the calculations feasible. If the image were discontinuous, we would have been obliged to represent it with a smoothed approximation in order to calculate the Cramér–Rao lower bound.

With the Cramér–Rao bound in hand, it is possible to determine the gap between the performance of well-known algorithms for landmark location and the theoretical limit, and to devise new algorithms that perform near the limit.

In response to the question of whether part-per-100,000 measurement is possible from digital images, the reader is invited to turn the page and join us in exploring the limits to precision landmark location.

http://www.springer.com/978-1-84628-912-5

Precision Landmark Location for Machine Vision and Photogrammetry
Finding and Achieving the Maximum Possible Accuracy
Gutierrez, J.A.; Armstrong, B.S.R.
2008, XI, 162 p. With online files/update., Hardcover
ISBN: 978-1-84628-912-5