Displaying Confidence Images

James G. Nagy* and Dianne P. O'Leary†

July 19, 2000

Abstract

Algorithms for computing images result in an estimate of an image. The image may result from deblurring a measured image, from deconvolving a set of measurements, or from computing an image by modeling physical processes such as the weather. These computations provide an estimated value for each pixel in the image. What is lacking, however, is an estimate of the statistical confidence that we can have in those pixel values or in the features they form. In this work we discuss novel ways to display confidence information, using an algorithm called Twinkle, in order to give the viewer valuable visual insight into uncertainties. The technique is useful whether the confidence information is in the form of a confidence interval or a distribution of possible values. We demonstrate how to display confidence information in a variety of applications: weather forecasts, intensity of a star, and rating a potential tumor in a diagnostic image.

Key Words: image restoration, confidence intervals

1 Introduction

We compute images in a variety of ways:

- We may apply a deblurring algorithm in order to compensate for imperfect camera lenses.
- Filtering may be applied to an image in order to reduce the effects of noise in the pixel measurements.
- A process such as an inverse Radon transform may be used to reconstruct an image from its projections.
- A complex fluid dynamics simulation may result in a flow pattern that is displayed as an image.

*Department of Mathematics, Emory University, Atlanta, Georgia 30322. nagy@mathcs.emory.edu

†Department of Computer Science, University of Maryland, College Park, Maryland. oleary@cs.umd.edu (301) 405-2678; fax (301) 405-6707. This work was supported by the National Science Foundation under Grant CCR 97-32022.

In all of these cases, the result of the computation is an estimate of an image – in particular, an estimate of the value of each pixel in that image. What is lacking, however, is an estimate of the *statistical confidence* that we can have in those pixel values or in the features they form.

In many of these computations, this statistical information can be computed; see, for example, [2]. Even if it is available, however, it must be displayed to the viewer in a visually useful manner. In this work we focus on displaying such confidence information in a novel way as a *confidence image* that gives the viewer valuable information on which features in the image can be trusted.

We present an algorithm called Twinkle that gives the viewer visual insight into uncertainties in the computed image. The algorithm is useful whether the confidence information is in the form of a confidence interval or a distribution of possible values. We demonstrate how to display confidence information in a variety of applications: weather forecasts, intensity of a star, and rating a potential tumor in a diagnostic image. Even so, these examples are merely indicative of the types of situations in which the Twinkle technique can be used.

2 Displaying Confidence Images Based on Pixel Estimates

Suppose that for each pixel in an image, we compute an estimate x_i of its true value as well as a confidence interval $[\ell_i, u_i]$ which is guaranteed with $\alpha\%$ certainty. In other words, if we repeated the data gathering 100 times, we would expect that $\alpha\%$ of the images would have each pixel value within its confidence interval. The collection of confidence intervals forms the *confidence image*. We are left with the task of displaying a confidence image in some useful way.

To do this, we developed an algorithm called Twinkle_image. We generate a sequence of images, each with pixel values contained in the ranges defined by the confidence image. Thus, in each of these images, pixel value i is taken to be a random value chosen from the interval $[\ell_i, u_i]$. We display this sequence of images as a movie, running the frames at a rate so that the change in frame is barely perceptible to the viewer. By comparing the movie with the image x, the viewer can conclude with $\alpha\%$ confidence that features that persist in the frames of the movie are real. Those that appear to flicker (or "twinkle") could be either real or artifact.

We show two examples for illustration, although unfortunately in a static manuscript such as this, the examples require some effort by the reader to imagine what the movie version looks like! A companion electronic version of this article[1] allows readers to view the movies.

Jorge E. Pinzón (private communication) tells us that he is using a similar technique to display uncertainties in the characteristics of ground cover, estimated by satellite images at NASA Goddard.

Example 1: Suspected tumors. Figure 1 is an example computed using Matlab. Suppose that medical imaging has produced this image. A physician might reasonably conclude from this data that there are three masses, and a treatment option appropriate to this conclusion would be chosen.

If the confidence intervals are such that the images in the left side of Figure 2 are typical of those within the 95% confidence level, then this conclusion would be justified. The movie of such images would have barely perceptible flicker, and the three masses would be persistent features. But if there is more noise in the measurements, or more uncertainty in the image reconstruction, then perhaps the images on the right side of Figure 2 are typical. In this case, the existence of three masses is far from certain. As the physician saw the movie of confidence images, at least two

[1] `http://www.mathcs.emory.edu/~nagy/Twinkle.html`

suspected masses would flicker, sometimes distinctly visible, but other times disappearing, and this would demonstrate the uncertainty in their existence.

In both scenarios, the clinician can make better decisions if presented with confidence images rather than just the reconstructed image of Figure 1.

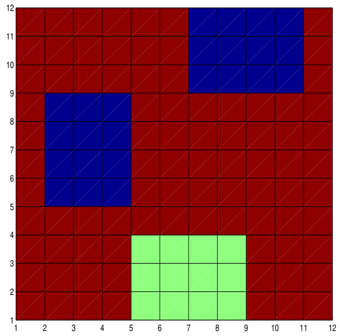

Figure 1: A model medical image

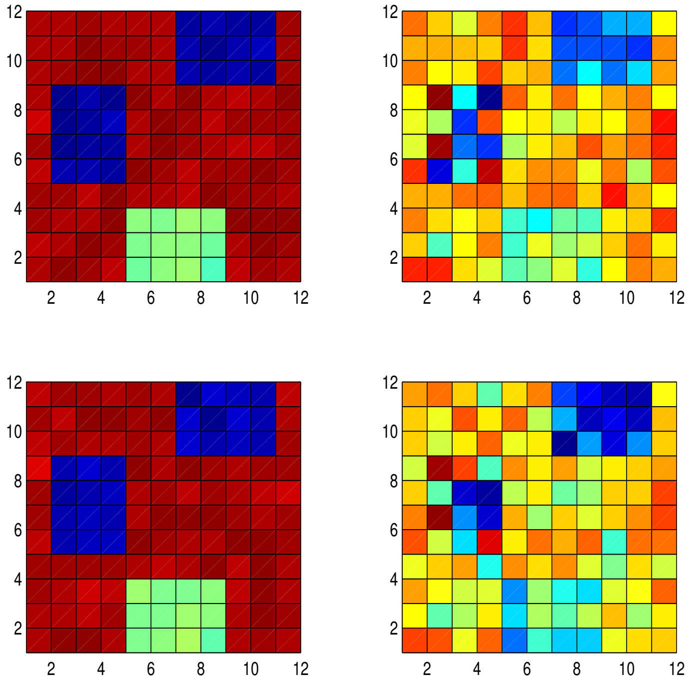

Figure 2: Two low-noise reconstructions of the medical image (left) and two high-noise reconstructions (right).

Example 2: Star intensity. Errors in the design of the original optics for the Hubble Space

Telescope Wide-Field Planetary Camera mean that rather complicated deconvolution techniques need to be applied in order to reconstruct images taken before the repairs. We obtained data from the Space Telescope Science Institute FTP server[2], intended to simulate a star cluster as it would appear in the camera image. The left image of Figure 3 is the true image, while the right is the camera image. We focus on an isolated star in the upper right corner. Sophisticated techniques are necessary to deblur the image: we need to take into account the spacial variation of the blurring function [1] and, in order to avoid "ringing" artifacts in the reconstruction, we need to add side constraints that the pixel values should be nonnegative. The result of such a reconstruction is shown in Figure 4. Bright stars can oversaturate an image, washing out dimmer stars and other objects, so we display images of the actual pixel values along with images of the log of the pixel values. The top two images show the true subimage, the second two are the reconstructed one, while the remaining are 99.99% confidence images [2], where the noise has been assumed to be normally distributed with mean zero and standard deviation one. Although the images containing the log pixel values show some change in the confidence images, the confidence images of the actual pixel values are virtually indistinguishable from the reconstructed image, giving high confidence in the reconstruction of the star and enabling estimates of derived quantities such as star intensity, as well as error estimates for these quantities.

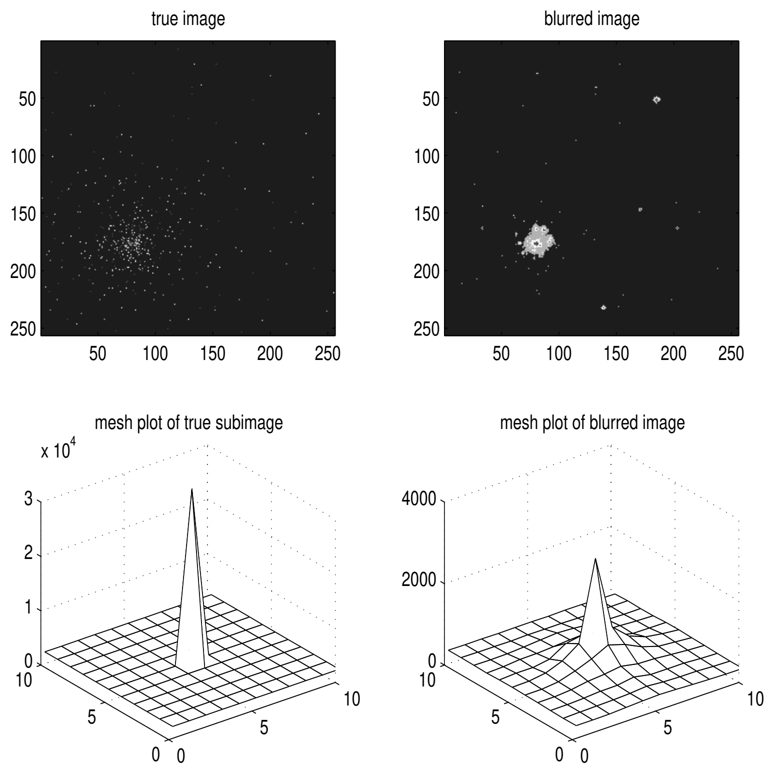

Figure 3: True star image and blurred image, with subimage extracted from the upper right corner.

[2]`ftp://ftp.stsci.edu/software/stsdas/testdata/restore/sims/star_cluster`

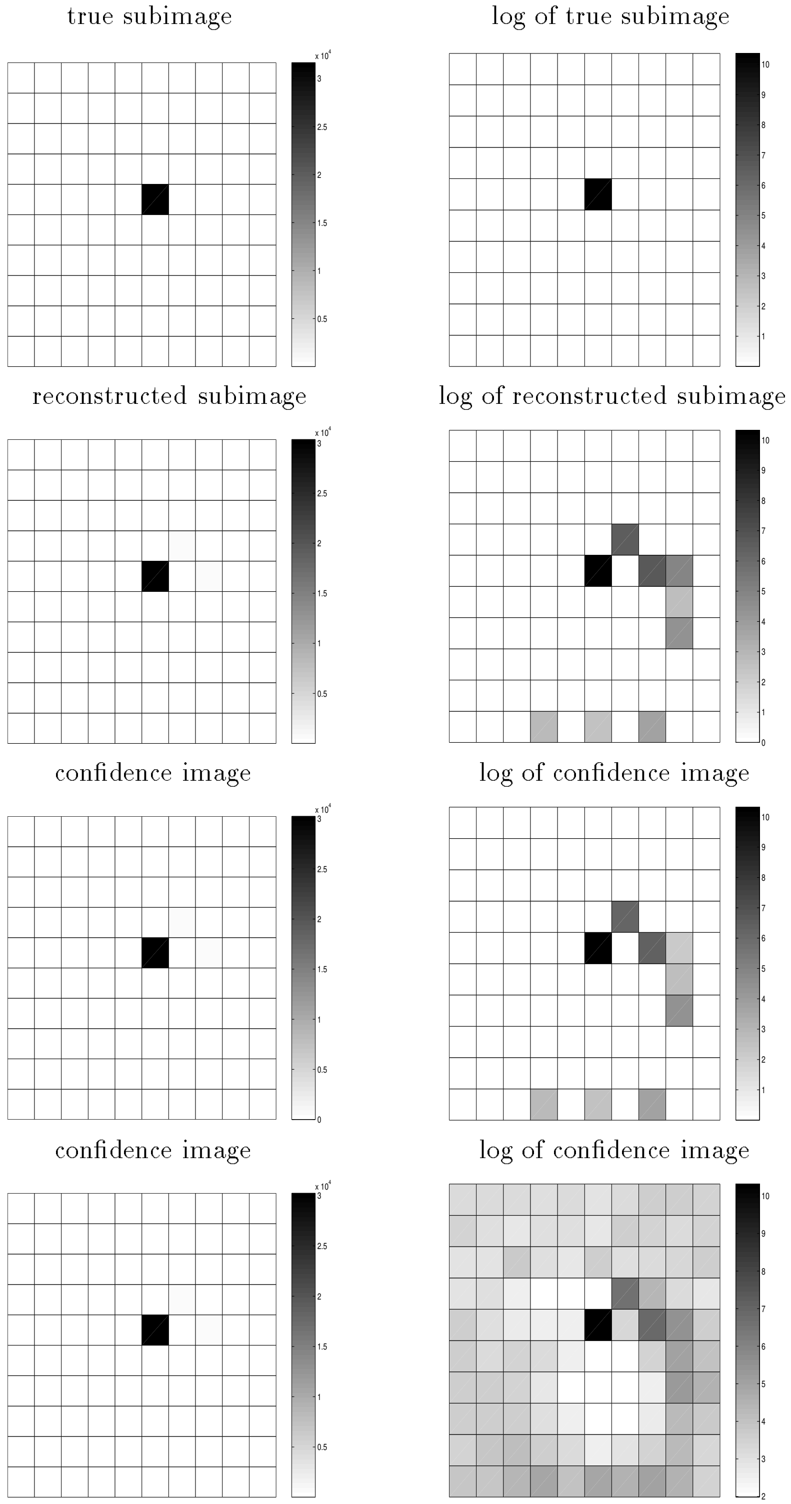

Figure 4: The true and reconstructed star subimages, with two confidence images. Pictures on the left show the actual pixel values, and images on the right show the log of the pixel values.

3 Displaying Confidence Images Based on Feature Estimates

The previous version of Twinkle is useful when we have confidence intervals on the pixel values in the image. In some cases, though, it is the *features* in the image that have uncertainty. For example, a fluid dynamics computation might produce flow lines that have uncertainty. Or a topographical map or estimate of a function might have uncertainties in the locations of contour lines.

In such cases, we apply a variant of the algorithm, Twinkle_feature. Each feature is represented by its distribution. For instance, we might know that the location of a contour has some expected value, some variance, and that the uncertainty is normally distributed. Twinkle_feature would then generate a sequence of contours chosen from this distribution and display them as a movie.

forecasted isotherms | one of the twinkled isotherms

one of the twinkled isotherms | one of the twinkled isotherms

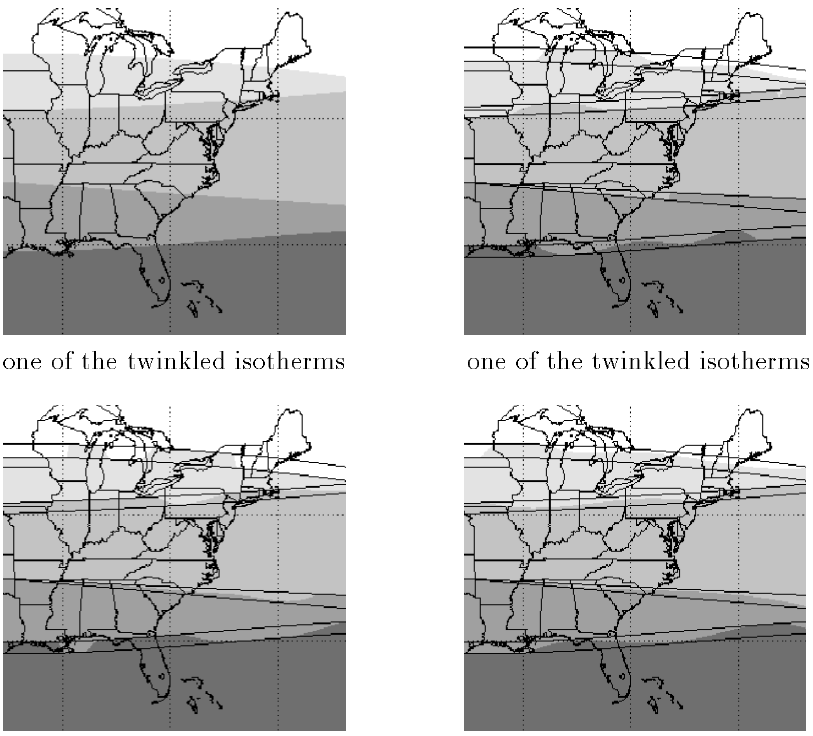

Figure 5: Forecasted temperatures (top left), and three images illustrating forecast uncertainties.

Example 3: Weather Forecasting. Weather forecasts of such things as high/low pressure regions, temperatures, precipitation, etc., are often displayed as isopleths; that is, lines connecting points of equal value of the particular variable in question. These isopleths are usually shown on maps as fixed lines which do not indicate any uncertainty in the forecast. An isotherm example is shown in the top left image of Figure 5. Uncertainty in temperature is not easy to visualize from this single image. Suppose, as an example, that the vertical uncertainty in the isotherms is

- 5 pixels in the top isotherm
- increasing from 0.2 to 5.6 pixels across the second isotherm,

- 0.05 times the y-coordinate of the third isotherm, and
- 6 vertical pixels in the bottom isotherm.

We display the bounds for the isotherms as solid lines. By generating sample locations of these isotherms within these bounds, we obtain pictures like the other three images of Figure 5. Displaying them as a movie gives the viewer a much clearer idea of what will be the expected temperature. Note, for example, that the uncertainty in northern Florida is nicely illustrated by the movie samples.

4 Conclusions

We have demonstrated a useful tool, Twinkle, for visualizing statistical confidence that we can have in pixel values or in the features they form. The technique is useful whether the confidence information is in the form of a confidence interval or a distribution of possible values. Further work remains to be done on efficient algorithms for computing such confidence information, since the information displayed is only as reliable as the confidence calculations.

References

[1] James G. Nagy and Dianne P. O'Leary. Restoring images degraded by spatially-variant blur. *SIAM Journal on Scientific Computing*, 19:1063–1082, 1998.

[2] James G. Nagy and Dianne P. O'Leary. Image restoration through subimages and confidence images. Technical Report CS-TR-4162, Department of Computer Science, University of Maryland, College Park, MD, 2000.