1998-10-01

Analysis of the Interpolation Error Between Multiresolution Images

Bryan S. Morse
morse@byu.edu

Analysis of the Interpolation Error Between Multiresolution Images

Bryan S. Morse

Department of Computer Science, Brigham Young University
3361 TMCB, Provo, UT 84602
morse@cs.byu.edu

Abstract

Many rendering or image-analysis systems require calculation of versions of an image at lesser resolutions than the original. Because the filtering required to perform such calculations accurately cannot typically be done in real time, many systems use interpolation between images at precalculated resolutions. This discrete sampling of the scale component of multiresolution image spaces is analogous to spatial sampling in discrete images. This paper quantifies and bounds the error that can be introduced during such interpolation as a function of the scale-space sampling rate used. A method is presented that uses the diffusion equation to relate spatial derivatives to scale derivatives and from there to an error bound.

1. Introduction

Many graphics and image-processing systems require that an initial high-resolution image be calculated (rendered) at some arbitrary lesser resolution. For example, when interactively viewing large, high resolution images one often needs to view the entire image at some reduced resolution while still being able to, when needed, view parts of the image at higher resolution. For example, in an interactive graphical environment where one can effectively move nearer or farther from an image, movement away from the image corresponds to a decrease in resolution (and corresponding increase in field of view) while movement towards the image corresponds to increasing resolution (and correspondingly decreasing the field of view). Alternatively, one may map an image onto the surface of an object that is visible at varying distances from the viewer (e.g., a receding surface where nearer parts of the object are visible at higher resolution and farther parts are visible at lower resolution).

To avoid aliasing artifacts, such multiresolution rendering obviously requires pre-filtering [1]. However, for many applications, on-the-fly filtering for arbitrary resolutions may not be feasible. The most commonly-used approach to this problem is to precompute versions of the image at a large number of pre-defined resolutions and to interpolate between them when asked to render the image (or some portion) at some arbitrary resolution, thus trading off precomputation and storage for increased interactive performance. When used for graphical texture mapping, this technique is known as *MIP mapping* [2] and is almost universally available in current graphics hardware and software systems.

Similar approaches are also used in multiscale analysis of images, in which a hierarchy of reduced-resolution versions of an image is generated [3, 4, 5, 6]. In many of these techniques, however, one may precompute versions of the image at specified resolutions but may find that desired properties exhibit themselves between these sampled resolutions—thus introducing a scale-space sampling question yet unanswered or even agreed upon in the image-processing community.

While interpolation between multiresolution images has advantages of simplicity and speed, it does not always approximate well the actual change in the value of a pixel under continuous change in resolution. Hybrid methods using filtering of precomputed resolutions instead of interpolation have been proposed [7], but although this approach avoids the inaccuracy of interpolation and is much faster than directly filtering the original image, it still requires further filtering of one of the precomputed images at interactive speeds.

While the limitations of interpolating between multiple resolutions (e.g., MIP mapping) are well known [7], little work has been published that quantifies or bounds the error in such methods. An example of such errors is illustrated in Figures 1 and 2. Clearly, this error can also be made less by more finely sampling the scales used to precompute multiresolution versions of the image. But this leads to an important, fundamental, and yet unanswered question: *what constitutes sufficient sampling of multiple resolutions when interpolating between multiresolution images?*

This paper presents a method for bounding the error in such multiresolution interpolation, thus allowing us to either estimate the resulting error or to find desired sampling rates that limit the error to a desired bound.

Original Image

Three-Quarter Resolution Image

Half-Resolution Image

Interpolated Image

Absolute Difference Image (range 0 to 81)

Figure 1: Example of errors that occur when interpolating between multiresolution images. LEFT COLUMN: original image (top) and half-resolution version (bottom). RIGHT COLUMN: three-quarter resolution (top), interpolated approximation of three-quarter resolution (middle), and absolute difference (bottom). Notice the artificial contrast enhancement and sharpening introduced in the interpolated image and reflected in the difference image. The difference image is normalized for display and has a maximum value of 81 (nearly one-third of the range of the original image).

a)

Gaussian p.s.f. with $\sigma = 1$ and $\sigma = 2$

b)

Gaussian p.s.f. with $\sigma = 1.5$ vs. interpolation between the two curves in (a).

c)

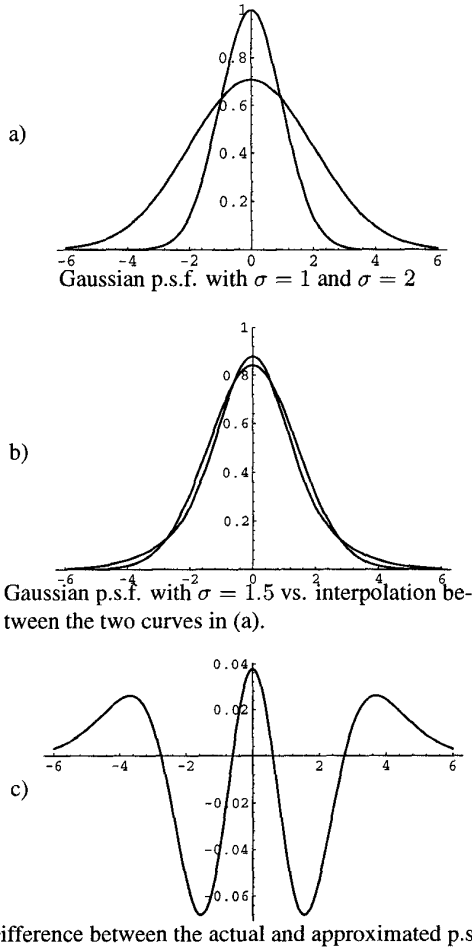Difference between the actual and approximated p.s.f.

Figure 2: Errors that can occur in interpolation of one-dimensional Gaussian point-spread functions. The interpolated function is sharper and higher-contrast than the correct function. Interpolation of two-dimensional point-spread functions behave similarly.

2. Scale Spaces

A useful tool for mathematically representing multiresolution spaces is the concept of a *scale space* (e.g., [4],[5], and many others): the set of all images of the same scene at varying resolutions. If we assume that the multiresolution images are "acquired" (generated) from some base image using a scaled, weighted measurement aperture applied uniformly over the image, such a scale space may be written as the convolution of the basis image with scaled versions of the measurement aperture:

$$L(\overline{x}, \sigma) = L(\overline{x}, 0) * G(\overline{x}, \sigma)$$

where $L(\overline{x}, 0)$ denotes the underlying scene (original image or "zero-scale" basis for the space), $*$ denotes convolution, and $G(\overline{x}, \sigma)$ denotes a measurement aperture with size σ.

It can be shown that in order to avoid artifacts from spurious resolution (temporary increases in sharpness as resolution decreases), the unique selection of aperture weights is the Gaussian [8]:

$$G(\overline{x}, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\overline{x} \cdot \overline{x}/\sigma^2}$$

For this reason, scale spaces are most commonly generated using Gaussian blurring where the blurring parameter σ is the "scale" of the image. Measurement scale (defined in this way) and resolution are inversely related.

An important property of Gaussian-generated scale spaces is that Gaussian blurring with scale σ is equivalent to running the diffusion equation for time t:

$$\frac{\partial}{\partial t} L = \nabla^2 L \qquad (1)$$

where $t = \sigma^2/2$. This property is key as we attempt to determine the error in linear interpolation in the resolution (scale) dimension of these spaces.

3. Error in Interpolation Between Resolutions

The approximation error ε in linear interpolation of a function f between two known values separated by h is

$$\varepsilon \leq \frac{h^2}{8} |f^{(2)}(x')| \qquad (2)$$

where x' is the intermediate point at which the magnitude of the second derivative of f is greatest [9]. Thus, if we can bound the second derivative of the multiresolution image *with respect to scale*, we can bound the error in such interpolation.

The key to bounding these derivatives with respect to scale is the diffusion equation. Using Eq. 1 and substituting $t = \sigma^2/2$ and $dt = \sigma \, d\sigma$,

$$\frac{\partial}{\partial \sigma} L = \frac{1}{\sigma} \nabla^2 L$$

Extending this to second-order derivatives,

$$\frac{\partial^2}{\partial \sigma^2} L = \frac{1}{\sigma^2} \nabla^2 \nabla^2 L$$

The implication of this is that if we can bound the fourth derivatives with respect to our spatial variables, we can likewise bound the second derivative with respect to scale (resolution).

Substituting this into Eq. 2,

$$\varepsilon \leq \frac{h^2}{8} \frac{1}{\sigma^2} |\nabla^2 \nabla^2 L(x')| \qquad (3)$$

where x' is now the intermediate point at which the magnitude of the second derivative *with respect to scale* (fourth derivative with respect to space) is greatest.

If we sample scale exponentially, as is usually the case in scale-space implementations [5] and multiresolution displays [2], the scale σ_i at step i of the resolution is $\sigma_i = \sigma_0 b^i$ for some exponential base b (the multiplying factor from scale to scale). The difference between one sampled scale σ and the next is thus $h = \sigma(b - 1)$. Substituting this into Eq. 3 gives

$$
\begin{aligned}
\varepsilon &\leq \frac{[\sigma(b-1)]^2}{8} \frac{1}{\sigma^2} |\nabla^2 \nabla^2 L(x')| \\
&= \frac{(b-1)^2}{8} |\nabla^2 \nabla^2 L(x')|
\end{aligned}
$$

If we bound $|\nabla^2 \nabla^2 L|$ by some value B, this becomes

$$
\boxed{\varepsilon \leq \frac{(b-1)^2}{8} B} \tag{4}
$$

where $b - 1$ is the percentage increase from one sampled scale to another and B is the estimated bound on the fourth-order spatial derivatives.

There are thus two values that control the error in interpolation between multiple resolutions:

1. The percentage increase between scales $(b - 1)$, and

2. The estimated bound on our fourth spatial derivatives $B \geq |\nabla^2 \nabla^2 L|$.

Example 1:

For a discretely-sampled initial image (i.e., ignoring any derivatives higher than those captured by the image discretization), the derivative bound B is four times the number of image grey-levels N. If we sample resolution through successive doubling ($b = 2$) as is often used in multiresolution methods [2, 3], the error bound per pixel is thus

$$
\varepsilon \leq \frac{N}{2}
$$

or as much as *one-half* of the image range.

Example 2:

Suppose that instead of simply determining the error, we wish to determine a sampling factor that ensures a desired bound on the error. With $N = 256$ intensity levels, again using $B = 4N$, and desiring a maximum error of a single intensity level ($\varepsilon < 1$), solving for sampling factor b gives

$$
b \leq 1.088
$$

This implies that to ensure an error bound of a single intensity level, one can reduce the height and width of the image by no more than 8.8% at a time, far less than successive halving of each dimension and much closer to the 1.1 or $\sqrt{1.1}$ scale multipliers reportedly used in recent multiscale research [6].

4. Conclusion

Using the diffusion equation as a way to tie second-order spatial derivatives to first-order scale derivatives in scale spaces, we have turned bounds on fourth-order spatial derivatives into a bound on the error in linear interpolation across resolutions. Although the potential for error in the interpolation of resolution has been appreciated for several years [7], the methods presented here provide a basis for quantitative analysis of this error. Similar techniques could also be used for higher-order interpolation functions.

References

[1] E. A. Feibush, M. Levoy, and R. L. Cook. Synthetic texturing using digital filters. In *SIGGRAPH*, 1980.

[2] L. Williams. Pyramidal parametrics. In *SIGGRAPH*, 1983.

[3] Peter J. Burt and Edward H. Adelson. The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31:532–540, 1983.

[4] Andrew P. Witkin. Scale space filtering. In *Proc. International Joint Conference on Artificial Intelligence (Karlsruhe, W. Germany)*, pages 1019–1023, 1983.

[5] Bart M. ter Haar Romeny and Luc Florack. A multiscale geometric model of human vision. In B. Hendee and P. N. T. Wells, editors, *Perception of Visual Information*. Springer-Verlag, Berlin, 1991.

[6] Stephen M. Pizer, Bryan S. Morse, David Eberly, and Daniel S. Fritsch. Zoom-invariant vision of figural shape: The mathematics of cores. *CVIU*, 1998.

[7] P. S. Heckbert. Filtering by repeated integration. In *SIGGRAPH*, 1986.

[8] J. Babaud, A. P. Witkin, M. Baudin, and R. O. Duda. Uniqueness of the Gaussian kernel for scale-space filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):26–33, 1986.

[9] Stephen M. Pizer and Victor L. Wallace. *To Compute Numerically*. Little, Brown, and Company, 1983.