Empirical Filter Estimation for Subpixel Interpolation and Matching

Bill Triggs, CNRS-INRIA, Grenoble, France — www.inrialpes.fr/movi/people/Triggs

Contributions

1 Empirically designed linear filters for subpixel image interpolation.
- Maximizes real performance on training images, not theoretical design criteria.
- Use for subpixel correlation matching, resampling.
- Similar to super-resolution, but uses only one input image and estimates what the camera would have seen if shifted a fraction of a pixel, not an enhanced image.

2 An experimental study of the stability & performance of the designed interpolators.

Method

1 Direct minimization of filter interpolation error on real training images
— various window sizes and filter types (separable, non-separable, symmetric)
— L1, L2 and robust error norms
— large optimization problem – use Limited Memory Quasi-Newton method.

2 Generate training data by ***subsampling large images at subpixel shifts***
— exploits scale-invariance & edge-dominated high-freq. spectra of natural images
— to mimic real zooms, subsample using a realistic ***Pixel Response Function (PRF)***

Conclusions

1 The optimal subpixel interpolators:
— oscillate like sinc functions but have much more rapidly decaying tails
— depend on the width of the PRF, but only weakly on its shape
— depend very little on the training images and robust cost norm used
— have noticeable derivative discontinuities at zero crossings
— are slightly non-separable (5-10% of amplitude).

2 Interpolation error is caused mainly by ***aliasing*** (jaggies) not noise
— once present, aliasing can not be removed by any linear interpolator
— slight (~ 0.5 pixel) optical blurring *before* sampling can greatly improve interpolation accuracy.

What's wrong with sinc interpolation?

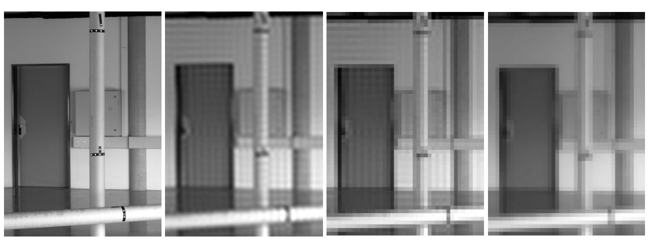

original image — band-limit — band-limit + subsample — blur + subsample

Real images have high-frequency spectra dominated by step edges. They are not exactly band-limited, and abrupt bandwidth truncation causes Nyquist frequency ringing that confuses the signal near edges. Blurring before subsampling reduces this, but even then sinc interpolation is not optimal for realistic images.

Pixel spatial Response Functions (PRF's)

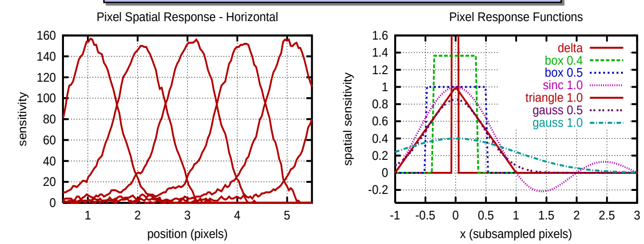

The PRF (discrete Point Spread Function) gives the response of a pixel as the (blurred & sampled) image of an ideal point source moves across it. It is easily measured from an image of a slightly non-aligned step edge. We designed interpolators for several PRF's to study what difference they make.

Optimal Subpixel Interpolators (separable case)

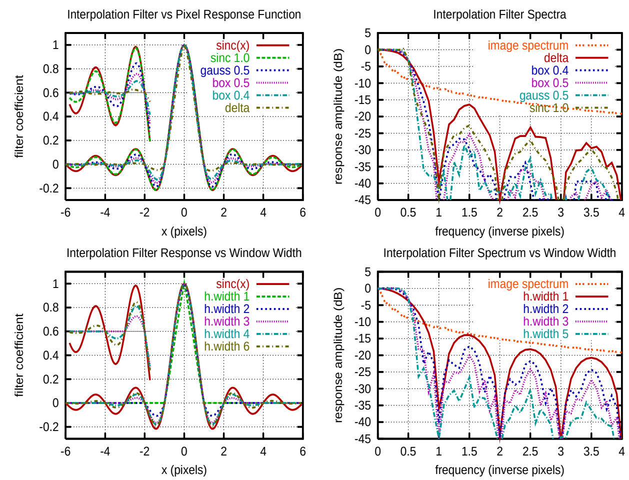

Optimal interpolator impulse responses and spectra for separable L2-norm filters designed for various PRF's (top, half-width $w = 6$) and filter widths (bottom, gauss 0.5 PRF).

Interpolation error is mainly due to aliasing not noise

Decimating image edge (a) produces strong jaggies (b). Even an optimal interpolation filter can not recover from these (c). The prediction error (d) is concentrated along strong edges, in phase with the jaggies. A more smoothly subsampled image (e) has no visible jaggies and can be reconstructed with far less error (f).

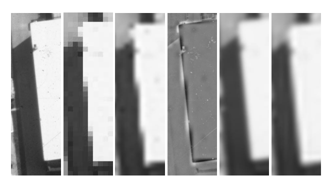

(a) (b) (c) (d) (e) (f)

Interpolation Filter Performance

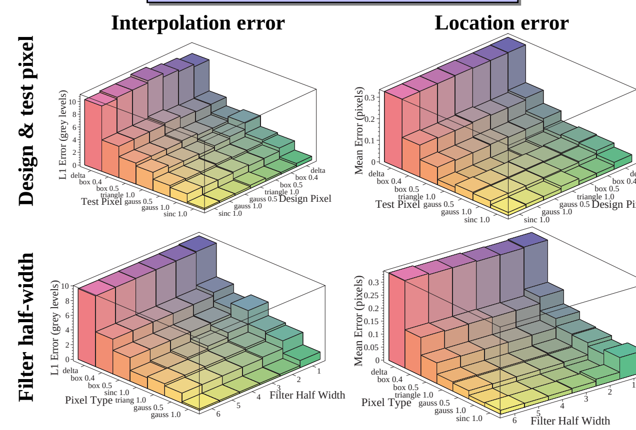

Location Error vs. Post-sampling Smoothing

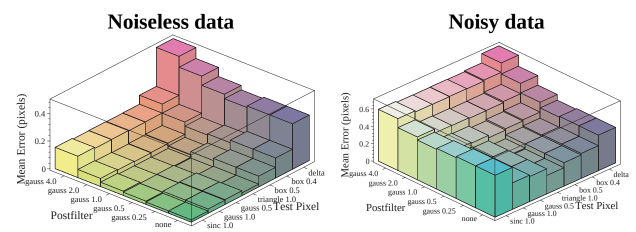

- Smoothing *after* sampling does not improve localization accuracy
- Sinc interpolators have poor noise resistance owing to their large L2 norm.

Gradient discontinuities and implied windowing function

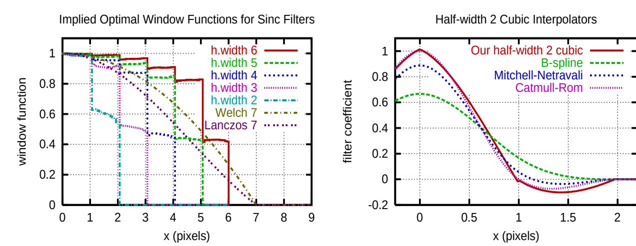

- Optimal interpolators have small gradient discontinuities at zero crossings.
- Optimal windows for sinc interpolators are staircases, last 2 lobes $\sim 80\%, 40\%$ amplitu

Acknowledgements Partly supported by European Union FET-Open project VIBES.