Title:
Does Median Filtering Truly Preserve Edges Better Than Linear Filtering

Author(s):
Ery Arias-Castro and David L. Donoho

Technical Report number (Dept. of Statistics, Stanford Univ.):
2007-6

Date:
December 2007

Abstract Image processing researchers commonly assert that "median filtering is better than linear filtering for removing noise in the presence of edges". Using a straightforward large-n decision-theory framework, this folk-theorem is seen to be false in general. We show that median filtering and linear filtering have similar asymptotic worst-case mean-squared error (MSE) when the signal-to-noise ratio (SNR) is of order 1, which corresponds to the case of constant per-pixel noise level in a digital signal. To see dramatic benefits of median smoothing in an asymptotic setting, the per-pixel noise level should tend to zero (i.e. SNR should grow very large).

We show that a two-stage median filtering using two very different window widths can dramatically outperform traditional linear and median filtering in settings where the underlying object has edges. In this two-stage procedure, the first pass, at a fine scale, aims at increasing the SNR. The second pass, at a coarser scale, correctly exploits the nonlinearity of the median.

Image processing methods based on nonlinear partial differential equations (PDEs) are often said to improve on linear filtering in the presence of edges. Such methods seem difficult to analyze rigorously in a decision-theoretic framework. A popular example is mean curvature motion (MCM), which is formally a kind of iterated median filtering. Our results on iterated median filtering suggest that some PDE-based methods are candidates to rigorously outperform linear filtering in an asymptotic framework.