Download pdf using QR code application

EPSRC
Engineering and Physical Sciences Research Council

Aston University

A behavioural method for assessing the performance of image denoising algorithms

Daniel H. Baker, Rob J. Summers & Tim S. Meese

Microsoft Research

AGM 22nd May 2012

email: d.h.baker1@aston.ac.uk

Background

- Denoising algorithms remove unwanted noise from images
- They are typically assessed numerically, by calculating RMS error between original and denoised images
- But RMS error does not correlate well with human perception (see Wang & Bovik, 2009)
- Here we develop a behavioural method for assessing the perceptual efficacy of denoising algorithms
- The method is simple, and could be applied to other algorithms such as compression, deblurring etc.

How does denoising work?

- Convolve image with filter bank
- Threshold based on high SF activity
- Re-synthesize denoised image

The filtering and thresholding process is similar to that performed by the early visual system, and may be related to contrast constancy effects (McIlhagga, 2004).

Denoising algorithms are not perfect. When noise is substantial, filter 'ringing' artefacts are common (see lower right image). Also, legitimate image content is often lost, especially at higher frequencies.

Log-Gabor filters have been found to be numerically superior to alternatives (Fischer et al., 2007). However, there are two varieties of log-Gabor filters (see below). Furthermore, filter bandwidth is often arbitrarily chosen. Is it the case that some filter banks produce more salient ringing artefacts than others?

We used a standard denoising algorithm (Fischer et al., 2007) to compare a range of filter bandwidths and two filter types.

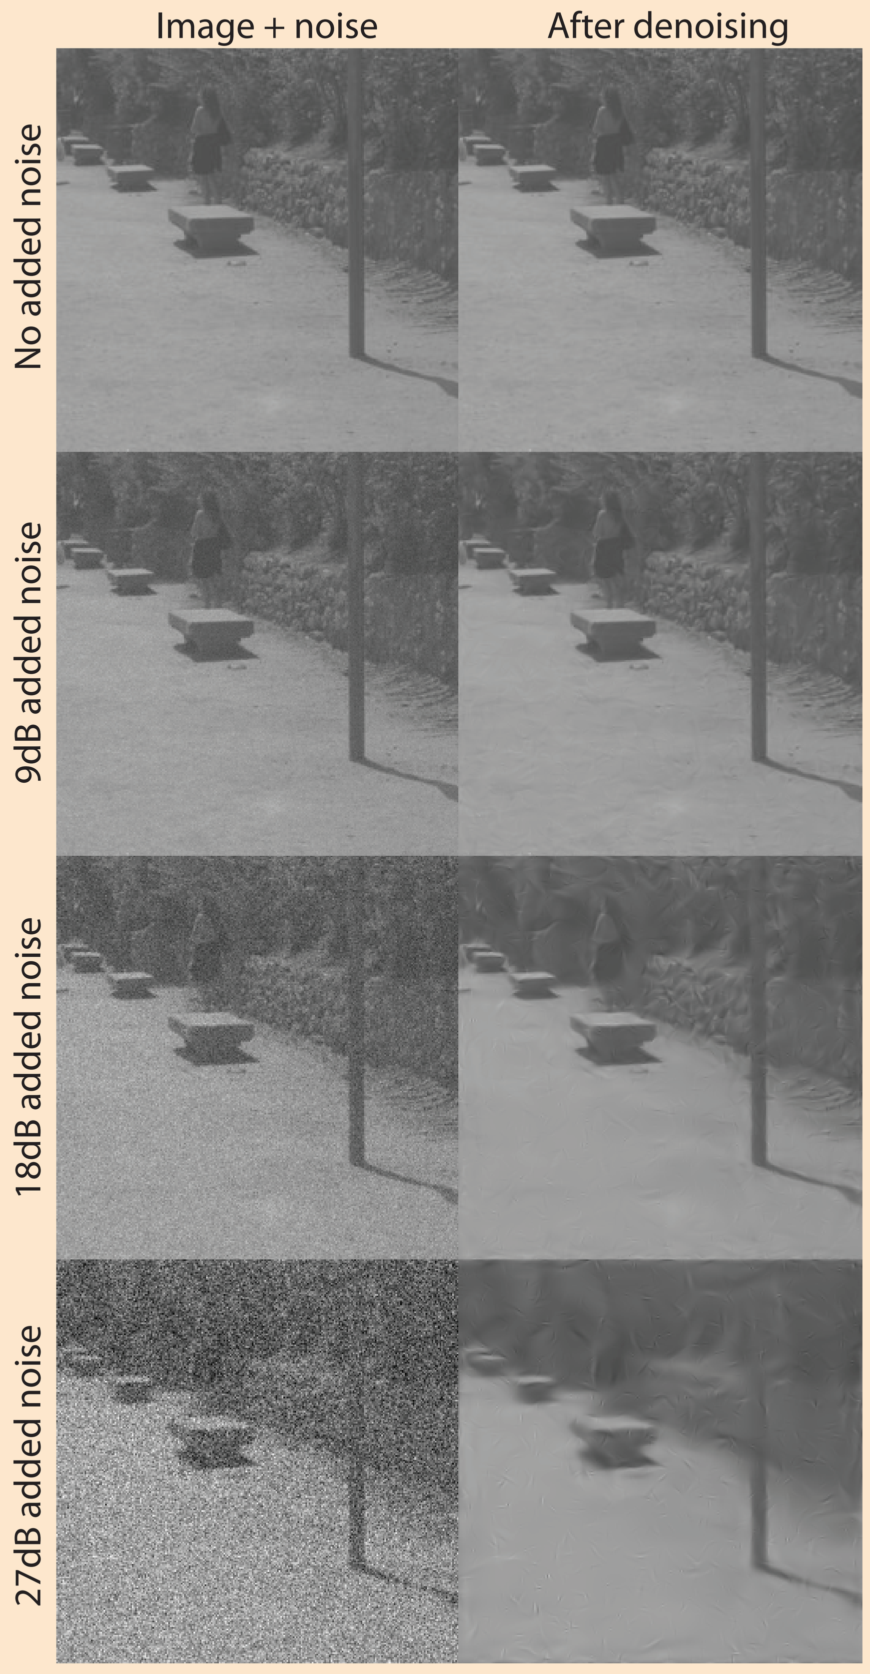

What sort of filters are best?

The spatial frequency profile of Log-Gabor filters is a log-Gaussian in the Fourier domain (Field, 1987), but the orientation profile can be defined in either polar or cartesian coordinates. These two types of log-Gabor (shown here at 3 orientation bandwidths) have different properties:

Polar-separable

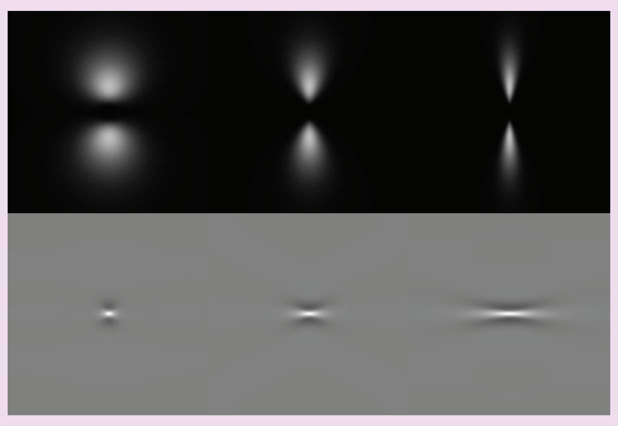

- Defining the orientation component in Polar coordinates produces filters which curve outwards in the spatial domain at narrow bandwidths
- These could produce salient artefacts when used in denoising
- However, polar filters can have very large orientation bandwidths - they can be isotropic

Cartesian-separable

- Cartesian separable filters are more spatially compact, curving inwards. This makes them more useful as a stimulus in experiments (reported elsewhere)
- However, for a given spatial frequency bandwidth, there is a maximum possible orientation bandwidth, where the spatial profile of the filter is one pixel wide. So isotropic Cartesian filters are not possible

Methods

Baseline condition
- Two interval forced choice (2IFC) task
- One interval is an unaltered image
- The other interval is a *different* image with noise added
- Task is to report the noisy image
- Staircase controls contrast of noise

Denoised conditions
- Exactly the same as the baseline, except that both images were denoised
- Observers are detecting the denoising artefacts (see left for examples)
- Range of filter bandwidths and types

Details
- Noise was 2D white pixel noise
- Images were windowed with a raised cosine envelope, 10 degrees wide
- 100ms presentation duration
- Feedback indicates correct/incorrect
- 3 observers, 4 repetitions each

Equipment
- Nokia MultiGraph 445x monitor
- Gamma corrected, mean lum 60cd/m^2
- Cambridge Research Systems ViSaGe

Analysis
- Fit psychometric functions using Probit analysis to estimate thresholds
- Compare thresholds across conditions

Image set

350 cropped images from the Barcelona Calibrated Image Database (256x256 pixels), converted to greyscale. Images were of both man-made and natural environments. We excluded images with an extremely steep or shallow spectral slope, leaving a mean amplitude spectral slope of -1.45 for 308 images. Images were DC balanced and contrast was reduced by 50% (mean RMS contrast = 21dB). The stimulus set grey-level histogram was skewed slightly to the right.

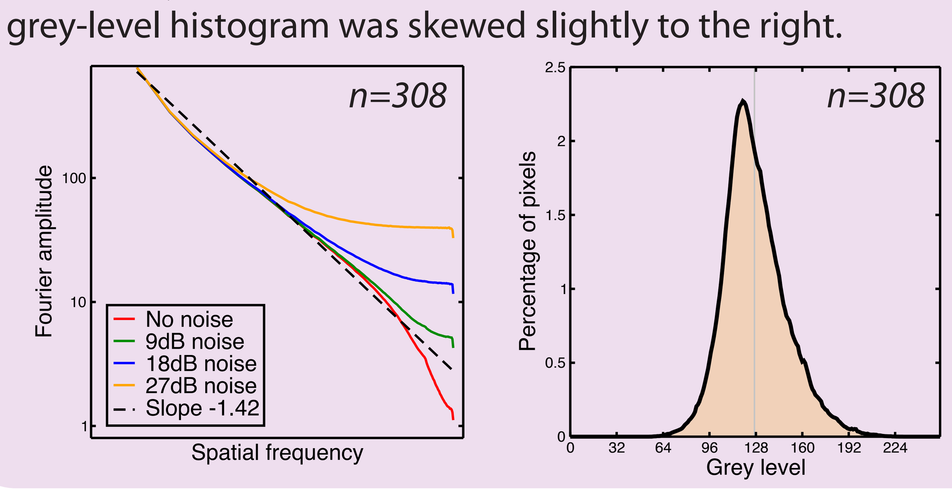

Which image contains noise?

Numerical assessment

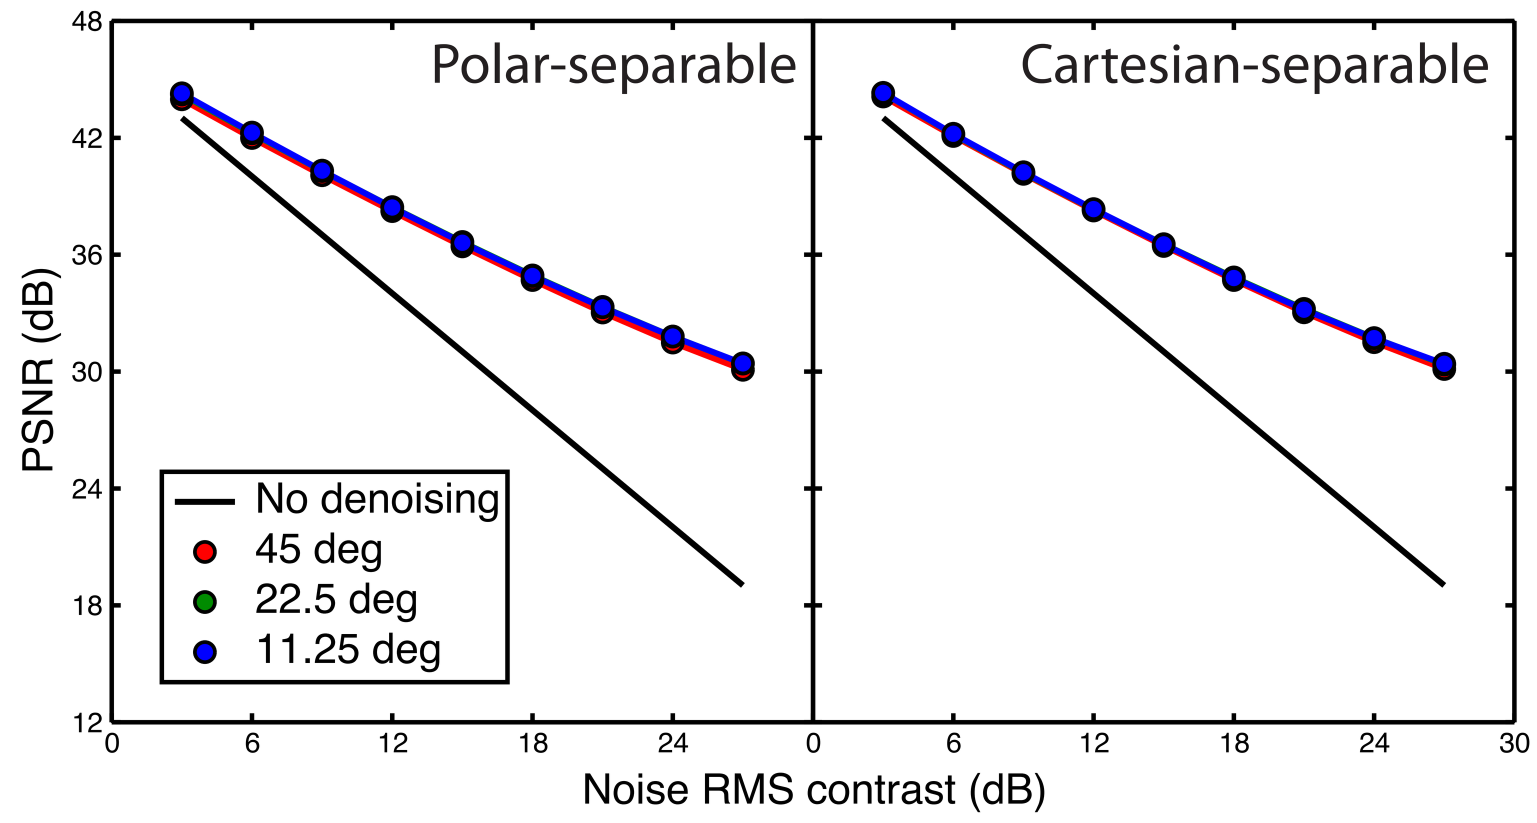

- Calculated peak signal-to-noise ratio (the RMS error between original and noisy image, converted to decibels) for the noisy images (black lines) and the denoised images for various filter regimes (circles)
- For added noise, the PSNR function had unit slope, so doubling noise contrast (an increase of 6dB) halved the PSNR (a decrease of 6dB)
- After denoising, the slope of the PSNR function was around 0.5, showing the effectiveness of the algorithm
- There were no differences between filter types or orientation bandwidths - is this also true perceptually?
- What is the relationship between PSNR and detection threshold?

Results

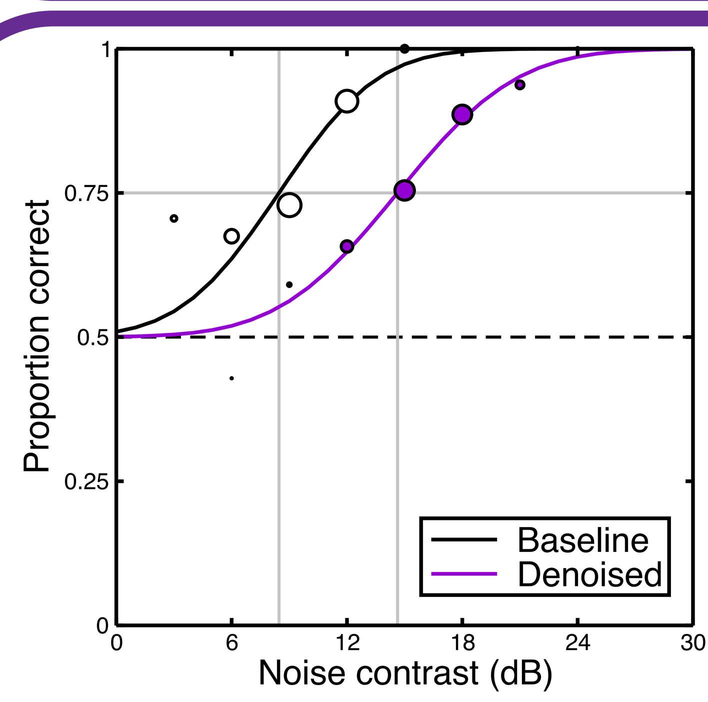

- Example psychometric functions (observer DHB) for the baseline (black) and one denoised condition (purple) - both are monotonic
- The threshold difference of ~6dB indicates that observers require twice as much noise contrast after denoising to reach threshold
- Put another way, the denoising algorithm restores 6dB of visible image quality before any image distortions become perceptually salient

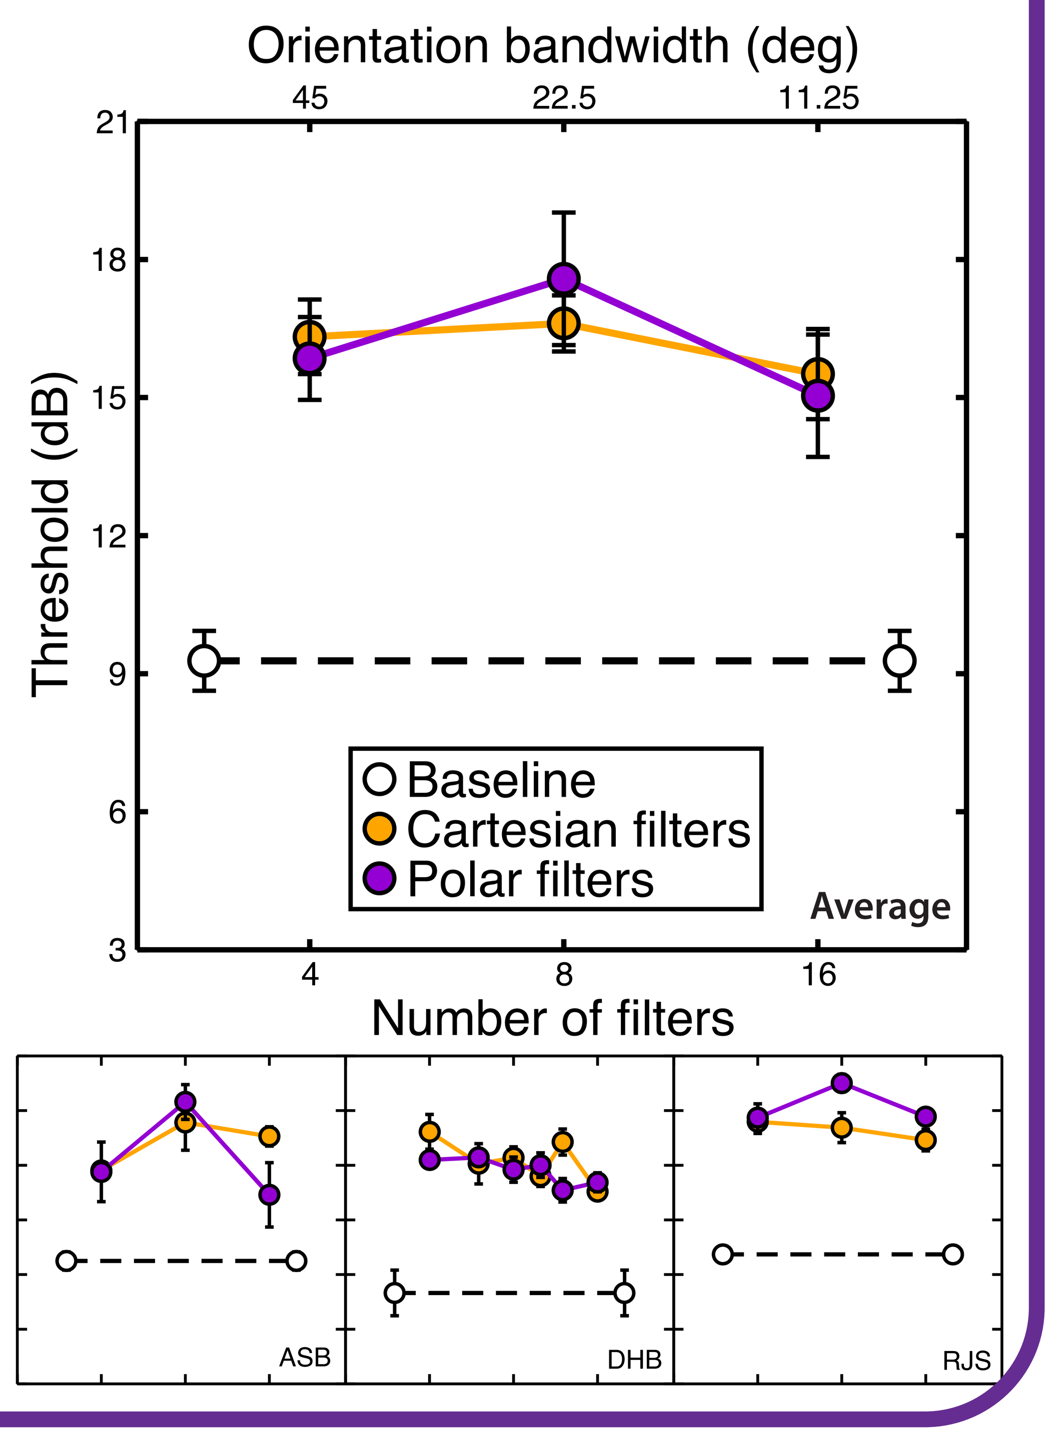

- Average data show that thresholds in the denoised conditions are not strongly dependent on filter type or bandwidth
- This confirms the findings of the numerical analyses (previous panel) using the PSNR (RMS error) metric
- Thresholds (dashed lines below) seem to occur at a fixed PSNR of around 36dB for noisy and denoised images
- This provides a useful heuristic for assessing when denoising artefacts will be perceptually salient

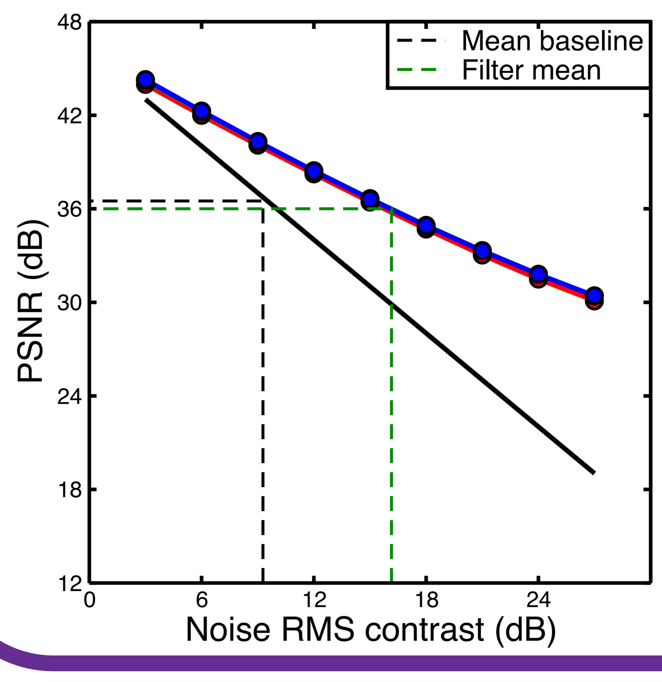

Conclusions

- Performance of image denoising algorithms can be assessed behaviourally using an objective detection task
- Behavioural data correspond well to numerical measures, but provide further data on when noise or artefacts will be perceptually relevant
- There are no clear differences between filter types or bandwidths, so using broader bandwidths (and hence fewer filters) is preferred
- A similar technique should work for assessing compression artefacts (e.g. JPEG artefacts)

References

- Barcelona Calibrated Image Database: http://www.cvc.uab.es/color_calibration/
- Field, D.J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *J Opt Soc Am A*, 4: 2379-2394.
- Fischer, S., Sroubek, F., Perrinet, L., Redondo, R. & Cristobal, G. (2007). Self-invertible 2D log-Gabor wavelets. *Int J Comp Vis*, 75: 231-246.
- McIlhagga, W. (2004). Denoising and contrast constancy. *Vision Res*, 44: 2659-2666.
- Wang, Z. & Bovik, A.C. (2009). Mean squared error: love it or leave it? *IEEE Signal Processing Magazine*, 26: 98-117.