



Opinionated
Lessons
in Statistics

by Bill Press

#42 Wiener (Optimal) Filtering



Norbert Wiener
1894 - 1964

Wiener Filtering (a.k.a. Optimal Filtering)

This general idea can be applied whenever you have a basis in function space that concentrates “mostly signal” in some components relative to “mostly noise” in others.

You *could* just set components with too much noise to zero.

Wiener filtering is better: it gives the optimal way of tapering off the noisy components, so as to give the best (L^2 norm) reconstruction of the original signal.

Can be applied in spatial basis (delta functions, or pixels), Fourier basis (frequency components), wavelet basis, etc.

Different bases are not equivalent, because, in particular problems, signal and noise distribute differently in them. A lot of signal processing is finding the right basis for particular problems – in which signal is most concentrated.

(For simplicity, I’m going to write out the equations as if in a finite-dimensional space, but think infinite dimensional.)

You measure components of signal plus noise

$$C_i = S_i + N_i$$

Let's look for a signal estimator that simply scales the individual components of what is measured

$$\hat{S}_i = C_i \Phi_i$$

find the Φ_i s that minimize $\langle |\hat{\mathbf{S}} - \mathbf{S}|^2 \rangle$

Here is where we use the fact that we are in some orthogonal basis, so the L^2 norm is just the sum of squares of components:

$$\begin{aligned} \langle (\hat{\mathbf{S}} - \mathbf{S}) \cdot (\hat{\mathbf{S}} - \mathbf{S}) \rangle &= \left\langle \sum_i [(S_i + N_i)\Phi_i - S_i]^2 \right\rangle \\ &= \sum_i \{ \langle S_i^2 \rangle (1 - \Phi_i)^2 + \langle N_i^2 \rangle \Phi_i^2 \} - 2 \sum_i \Phi_i (1 - \Phi_i) \langle N_i S_i \rangle \end{aligned}$$

expectations come inside and land on the only two stochastic things, signal and noise

this is zero!

Differentiate w.r.t. Φ and set to zero, giving

$$\Phi_i = \frac{\langle S_i^2 \rangle}{\langle S_i^2 \rangle + \langle N_i^2 \rangle}$$

This is the Wiener filter. It requires estimates of the signal and noise power in each component.

Let's demonstrate in some different bases on this image:

```
IN = fopen('image-face.raw', 'r');  
face = flipud(reshape(fread(IN), 256, 256)');  
fclose(IN);  
bwcolormap = [0: 1/256: 1; 0: 1/256: 1; 0: 1/256: 1]';  
image(face)  
colormap(bwcolormap);  
axis('equal')
```

the file is unformatted bytes, so you read it with `fopen` and `fread`

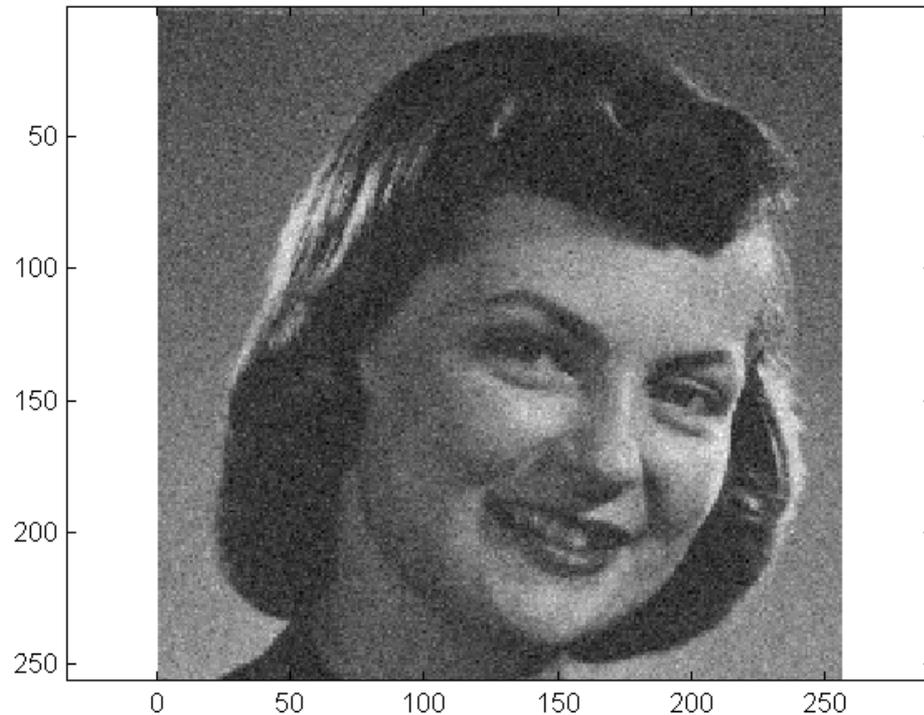


(Favorite demo image in NR. Very retro, it's an IRE test photo from the 1950s, shows film grain and other defects.)

Add noise (here, Gaussian white noise):

```
noi syface = face + 20*randn(256, 256);  
noi syface = 255 * (noi syface - mi n(noi syface(:))) / (max(noi syface(:)) - mi n(noi syface(:)));  
i mage(noi syface)  
col ormap(bwcol ormap);  
axi s(' equal ')
```

Have to rescale, because noise takes it out of 0-255.
Note reduced contrast resulting.



First example: **Fourier basis**

This will be a “low pass filter” using the fact that the signal is concentrated at low spatial frequencies, while the noise is white (flat).

Actually, Fourier is not a great basis for de-noising most images, since low-pass will reduce the resolution of the picture (blur it) along with de-noising.

```
fftface = fft2(noisyface);
```

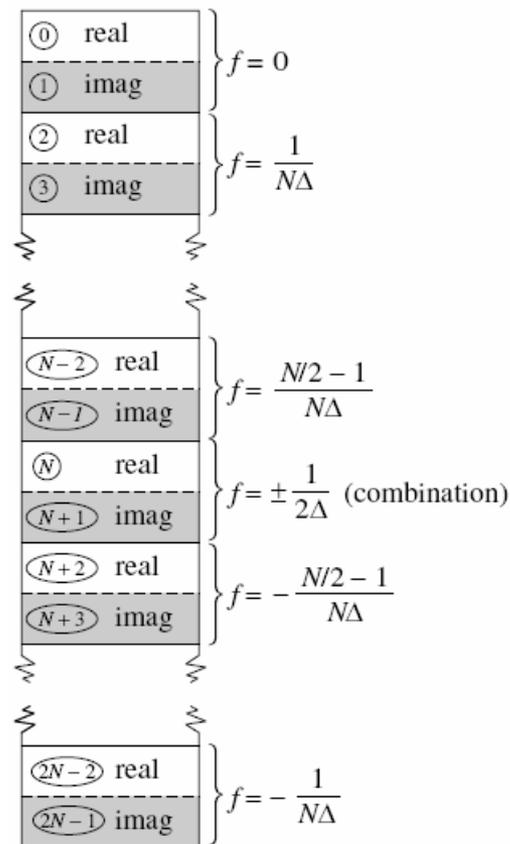
```
[row col] = ndgrid(1:256);
```

```
ftable = [0:127, 128:-1:1];
```

```
freqs = 0.5 * sqrt((ftable(row).^2+ftable(col).^2)/(2*128^2));
```

Nyquist frequency

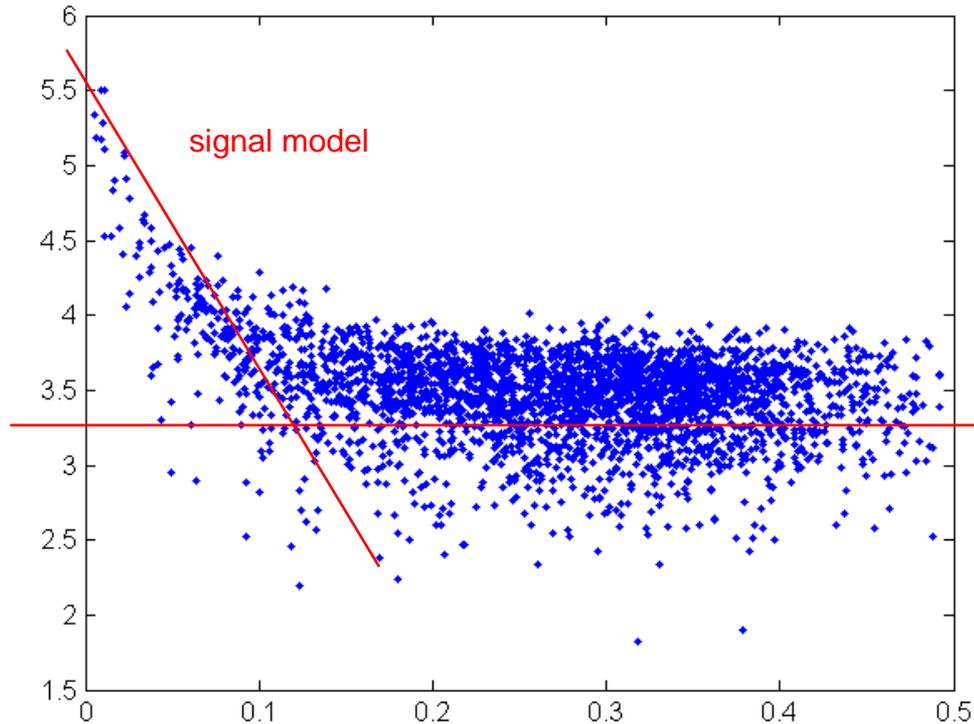
1 1 1... 1 2 3...
2 2 2... 1 2 3...
3 3 3... 1 2 3...



Yes, we can see a separation between signal and noise:

```
samp = randsample(256*256, 3000);  
plot(freql(samp), log10(abs(fftface(samp))), ' . ')
```

abs is here the complex modulus

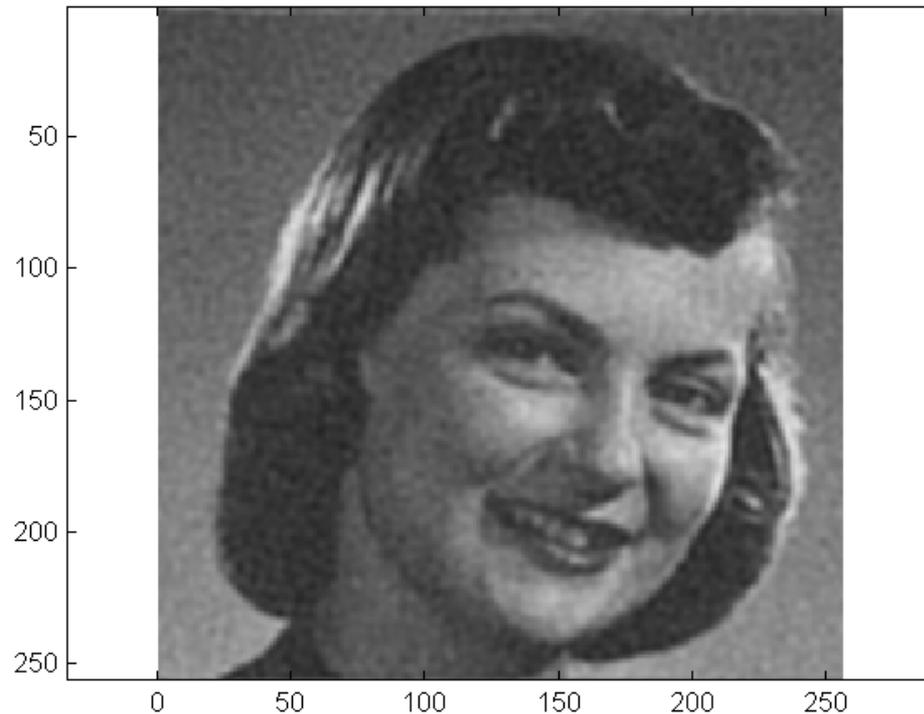


noise model

Why the line a bit low?
I played around with it to
make the picture look better!

So my reconstruction is not
exactly best least squares
reconstruction, but will be
less blurred.

de-noised



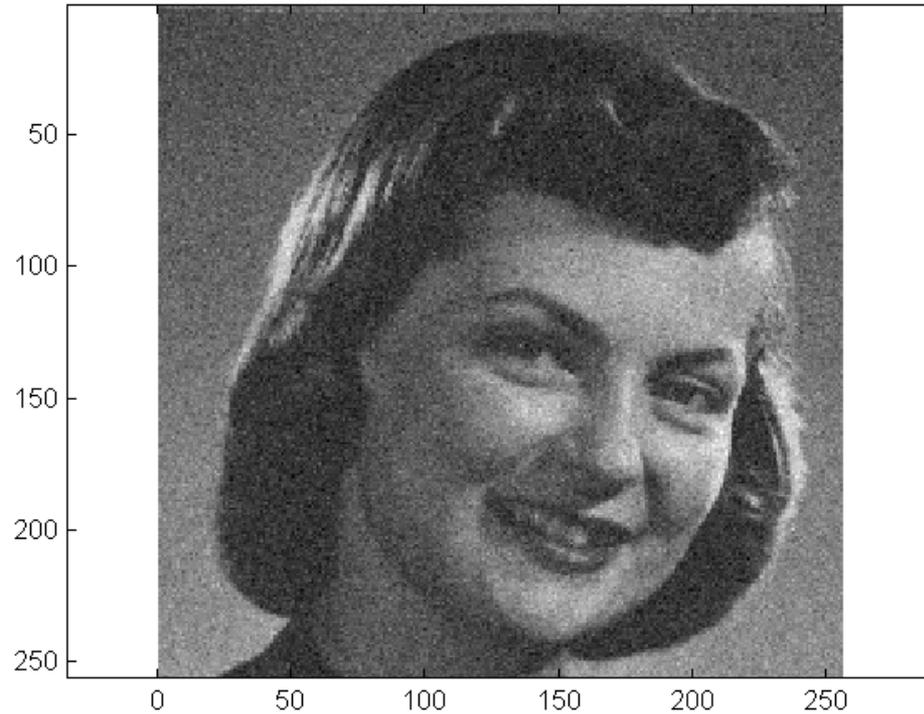
```
sig = 10.^(5.5 - 15 .* freqs);  
noi = 10.^3.2;  
fftfiltface = fftface .* (sig.^2 ./ (sig.^2 + noi.^2));  
reface = ifft2(fftfiltface, 'symmetric');  
image(reface)  
colormap(bwcolormap);  
axis('equal')
```

I just read the constants off by eye from the previous chart

note square, to get power

this tells Matlab that you intend the inverse FFT to be real-valued

noisy



Actually, you might prefer the noisy image, because your brain has good algorithms for adaptively smoothing! But it is a less accurate representation of the original photo in L^2 norm!

Second example: **spatial (pixel) basis**

It doesn't make sense to use the pure pixel basis, because there is no particular separation of signal and noise separately in each pixel.

But a closely related method is to decompose the image into a smoothed background image, and then to take deviations from this as estimating signal power + noise power:

$$X_{ij} = \frac{1}{N_{\text{hood}}} \sum_{\text{hood}} x_{ij}$$

“hood” might be a 5x5 neighborhood centered on each point

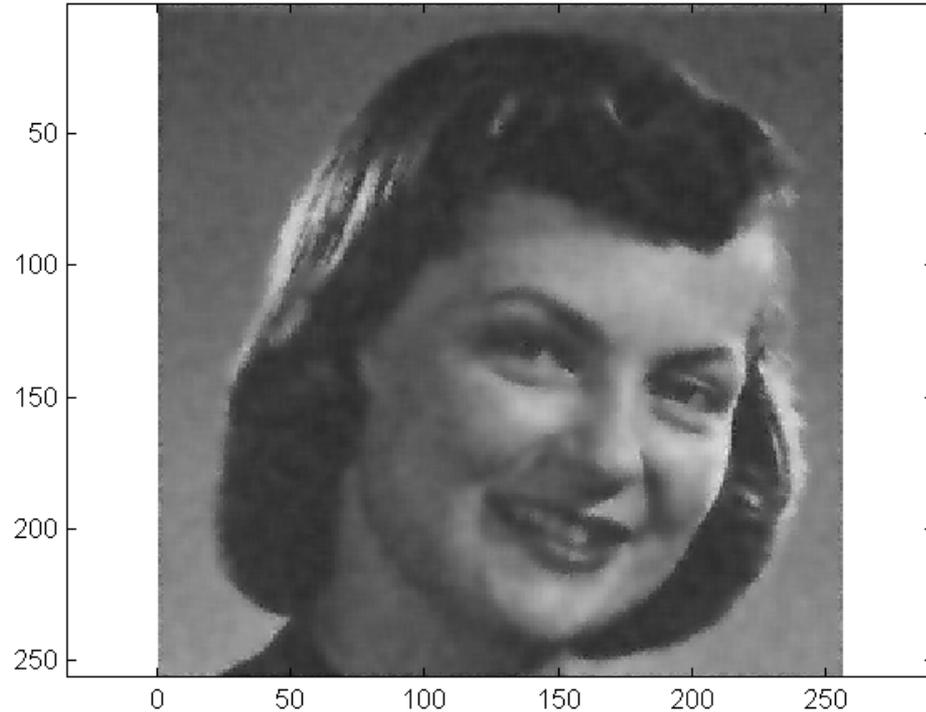
$$\langle S_{ij}^2 + N_{ij}^2 \rangle = \frac{1}{N_{\text{hood}}} \sum_{\text{hood}} (x_{ij} - X_{ij})^2$$

$$\hat{x}_{ij} = X_{ij} + \frac{\langle S_{ij}^2 + N_{ij}^2 \rangle - \langle N_{ij}^2 \rangle}{\langle S_{ij}^2 + N_{ij}^2 \rangle} (x_{ij} - X_{ij})$$

this is the Wiener part

and let the user adjust $\langle N_{ij}^2 \rangle$ as a parameter.

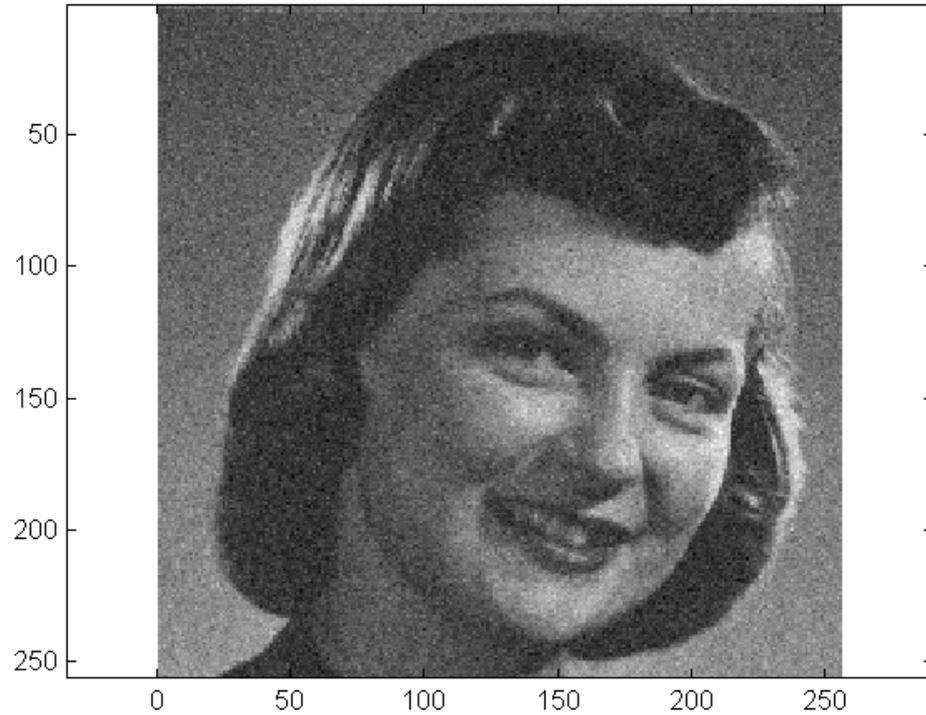
Matlab has a function for this called wiener2



```
wi ener = wi ener2(noi syface, [5, 5]);  
i mage(wi ener)  
col ormap(bwcol ormap);  
axi s(' equal ')
```

you can put your noise estimate as another argument, or you can let Matlab estimate it as some kind of heuristic minimum of values seen for S^2+N^2 over the image

noisy



Most fun of all is the **wavelet basis**.

You don't even have to know what it is, except that it is an (orthogonal) rotation in function space, as is the Fourier transform.

(Its basis is localized both in space and in scale.)

Matlab has a Wavelet Toolbox which I find completely incomprehensible! (I'm sure it's only me with this problem.) So, I'll do a mexfunction wrapper of the NR3 wavelet transform.

```
#include "nr3_matlab.h"
#include "wavelet.h"

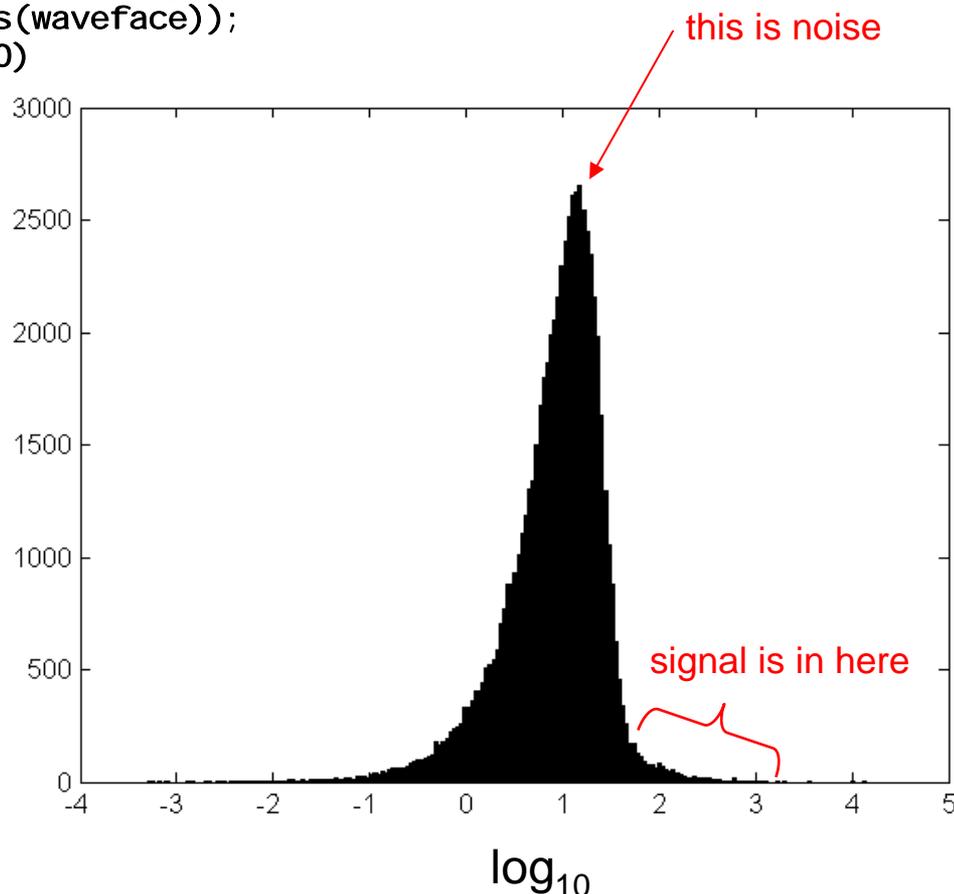
/* Matlab usage:
    outmatrix = wavelet2(inmatrix, isign)
*/

void mexFunction(int nlhs, mxArray *plhs[], int nrhs, const mxArray *prhs[]) {
    MatDoub ain(prhs[0]);
    VecInt dims(2);
    Int mm=(dims[0]=ain.nrows()),nn=(dims[1]=ain.ncols());
    Int isign = Int(mxScalar<Doub>(prhs[1]));
    Int i,mn = mm*nn;
    Daub4 daub4;
    if (nrhs != 2 || nlhs != 1) throw("wavelet2.cpp: bad number of args");
    if ((nn & (nn-1)) != 0 || (mm & (mm-1)) != 0)
        throw("wavelet2.cpp: matrix sizes must be power of 2");
    VecDoub a(mn);
    for (i=0;i<mn;i++) a[i] = (&ain[0][0])[i];
    wtn(a,dims,isign,daub4);
    MatDoub aout(mm,nn,plhs[0]);
    for (i=0;i<mn;i++) (&aout[0][0])[i] = a[i];
    return;
}
```

← this is the whole point,
the NR3 wavelet transform

Take the wavelet transform and look at the magnitude of the components on a log scale:

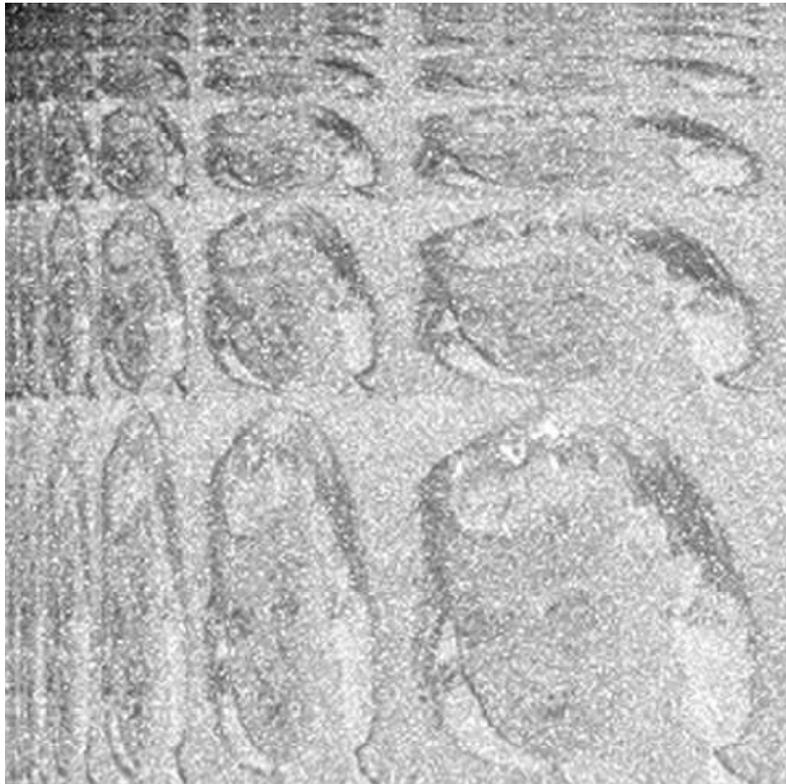
```
waveface = wavelet2(noisyface, 1);  
dist = log10(abs(waveface));  
hist(dist(:), 200)
```



Notice the difference in philosophy from Fourier: There we used frequency (“which component”) to estimate S and N. Here we use the magnitude of the component directly, without regard to which component it is.

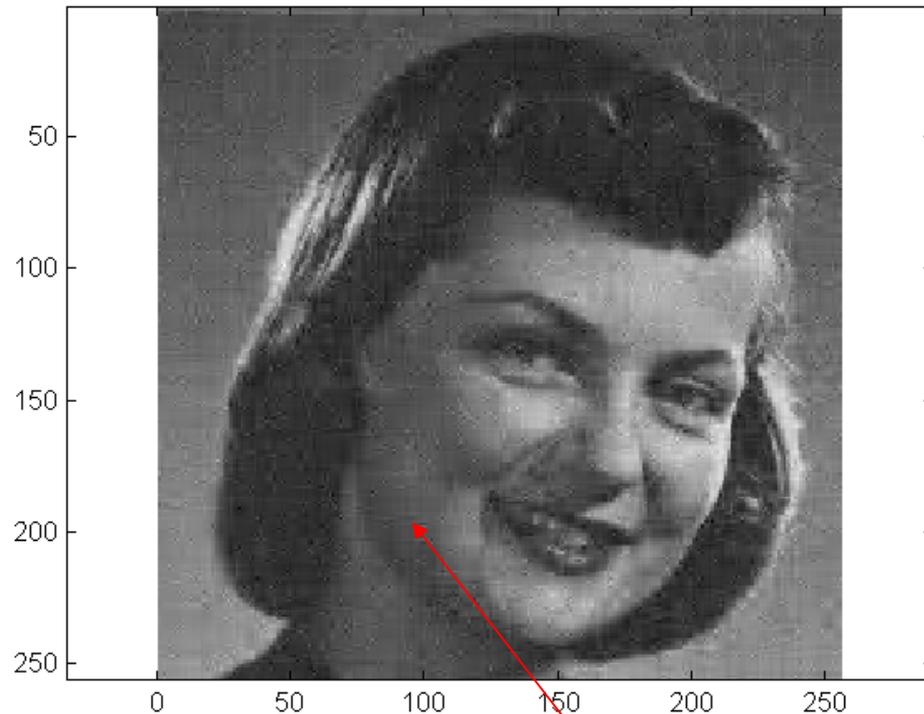
If you fiddle around with mapping the gray scale (zero point, contrast, etc.) of the matrix “waveface” you can see how the wavelet basis works

low resolution
information is in this
corner



high resolution
information is in this
corner

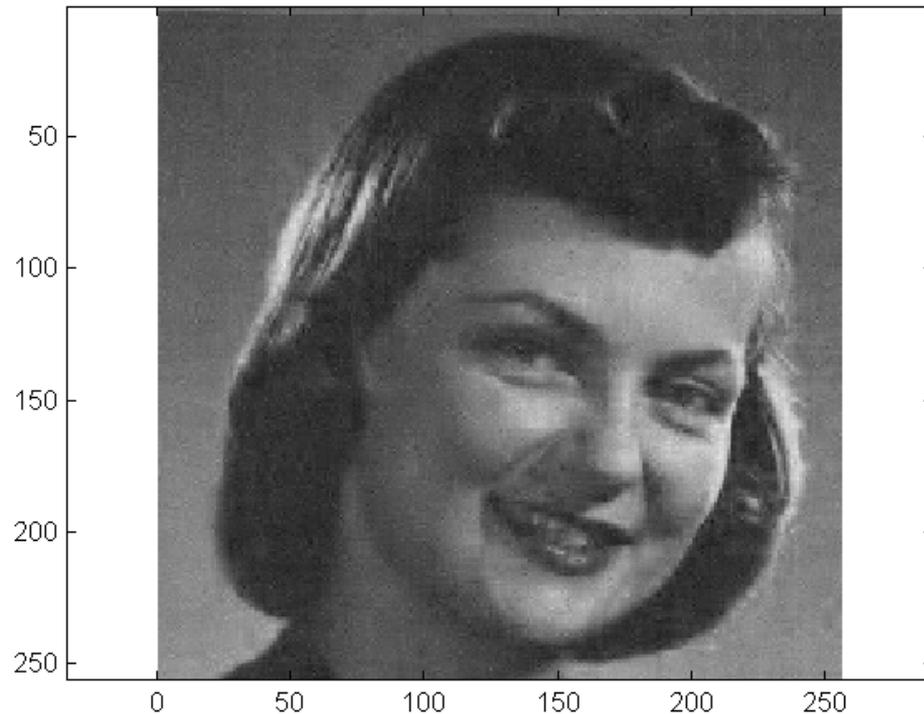
Truncate-to-zero components with magnitude less than 30.
This is not a true Wiener filter, because it doesn't roll off smoothly.



```
fwaveface = waveface;  
fwaveface(abs(waveface)<30) = 0. ;  
werface = wavelet2(fwaveface, -1);  
image(werface)  
colormap(bwcolormap)  
axis('equal')
```

Notice the “wavelet plaid” in the image. You sometimes see this on digital TV, because MPEG4 uses wavelets for still texture coding.

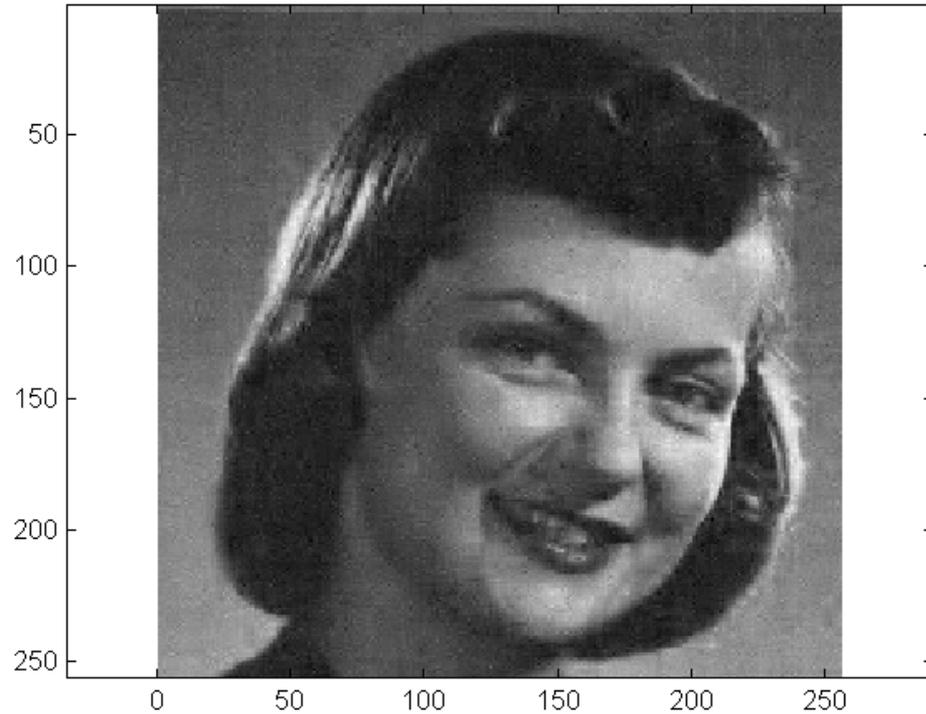
Compare to Wiener filter (smooth roll-off)



```
fwaveface = waveface .* (waveface.^2 ./ (waveface.^2 + 900));  
wrecface = wavel et2(fwaveface, -1);  
i mage(wrecface)  
col ormap(bwcol ormap)  
axi s(' equal ')
```

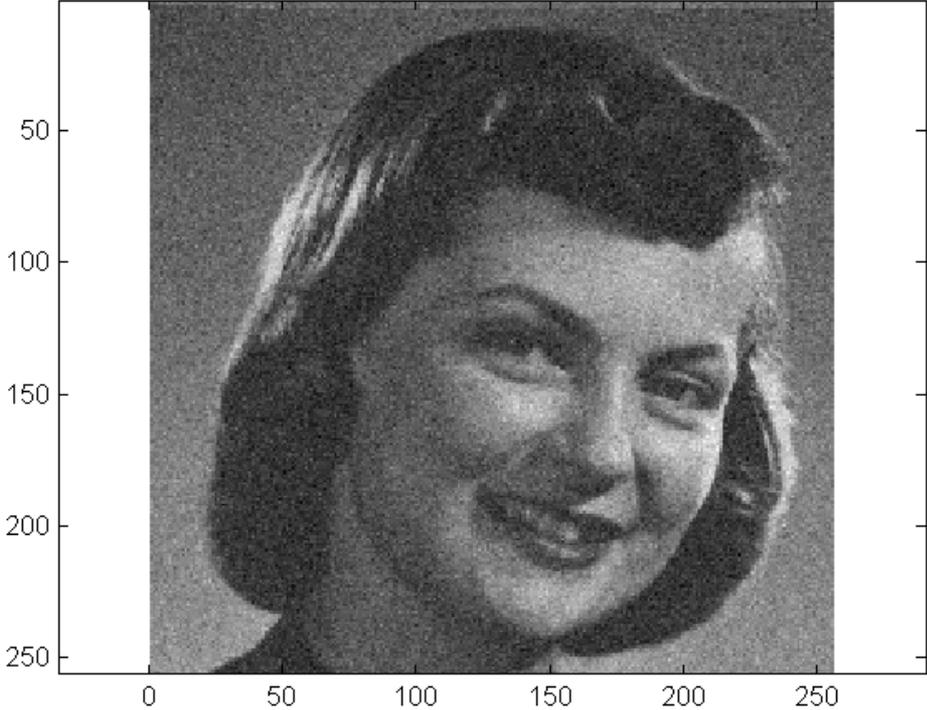
i.e., noise amplitude 30 in the
previous histogram

Even better if we restore the contrast



```
werecface = 255*(werecface - min(werecface(:)))/(max(werecface(:))-min(werecface(:)));  
image(werecface)  
colormap(bwcolormap)  
axis('equal')
```

Compare to what we started with (noisy)



Of course, we can never get the original back:
information is truly lost in the presence of noise



The moral about Wiener filters is that they work in any basis, but are better in some than in others. That is what signal processing is all about!