


2. Image Formation



3. Image Processing




4. Features



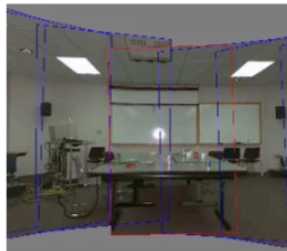
5. Segmentation




6-7. Structure from Motion




8. Motion



9. Stitching




10. Computational Photography



11. Stereo



12. 3D Shape



13. Image-based Rendering




14. Recognition

3.1	Point operators	101
3.1.1	Pixel transforms	103
3.1.2	Color transforms	104
3.1.3	Compositing and matting	105
3.1.4	Histogram equalization	107
3.1.5	<i>Application: Tonal adjustment</i>	111
3.2	Linear filtering	111
3.2.1	Separable filtering	115
3.2.2	Examples of linear filtering	117
3.2.3	Band-pass and steerable filters	118
3.3	More neighborhood operators	122
3.3.1	Non-linear filtering	122
3.3.2	Morphology	127
3.3.3	Distance transforms	129
3.3.4	Connected components	131
3.4	Fourier transforms	132
3.4.1	Fourier transform pairs	136
3.4.2	Two-dimensional Fourier transforms	140
3.4.3	Wiener filtering	140
3.4.4	<i>Application: Sharpening, blur, and noise removal</i>	144
3.5	Pyramids and wavelets	144
3.5.1	Interpolation	145
3.5.2	Decimation	148
3.5.3	Multi-resolution representations	150
3.5.4	Wavelets	154
3.5.5	<i>Application: Image blending</i>	160

Median filters

- A **Median Filter** operates over a window by selecting the median intensity in the window.
- What advantage does a median filter have over a mean filter?
- Is a median filter a kind of convolution?

Comparison: salt and pepper noise



Bilateral filtering




Figure 3.20 Bilateral filtering (Durand and Dorsey 2002) © 2002 ACM: (a) noisy step edge input; (b) domain filter (Gaussian); (c) range filter (similarity to center pixel value); (d) bilateral filter; (e) filtered step edge output; (f) 3D distance between pixels.

Morphological Operators





Figure 3.21 Binary image morphology: (a) original image; (b) dilation; (c) erosion; (d) majority; (e) opening; (f) closing. The structuring element for all examples is a 5×5 square. The effects of majority are a subtle rounding of sharp corners. Opening fails to eliminate the dot, since it is not wide enough.


3.1	Point operators	101
3.1.1	Pixel transforms	103
3.1.2	Color transforms	104
3.1.3	Compositing and matting	105
3.1.4	Histogram equalization	107
3.1.5	<i>Application: Tonal adjustment</i>	111
3.2	Linear filtering	111
3.2.1	Separable filtering	115
3.2.2	Examples of linear filtering	117
3.2.3	Band-pass and steerable filters	118
3.3	More neighborhood operators	122
3.3.1	Non-linear filtering	122
3.3.2	Morphology	127
3.3.3	Distance transforms	129
3.3.4	Connected components	131
3.4	Fourier transforms	132
3.4.1	Fourier transform pairs	136
3.4.2	Two-dimensional Fourier transforms	140
3.4.3	Wiener filtering	140
3.4.4	<i>Application: Sharpening, blur, and noise removal</i>	144
3.5	Pyramids and wavelets	144
3.5.1	Interpolation	145
3.5.2	Decimation	148
3.5.3	Multi-resolution representations	150
3.5.4	Wavelets	154
3.5.5	<i>Application: Image blending</i>	160

Why does the Gaussian give a nice smooth image, but the square filter give edgy artifacts?


Gaussian



Box filter



Why does a lower resolution image still make sense to us? What do we lose?



Thinking in Frequency

Fourier, Joseph (1768-1830)



French mathematician who discovered that any periodic motion can be written as a superposition of sinusoidal and cosinusoidal vibrations. He developed a mathematical theory of [heat](#) 🌡️ in *Théorie Analytique de la Chaleur (Analytic Theory of Heat)*, (1822), discussing it in terms of differential equations.

Fourier was a friend and advisor of Napoleon. [Fourier believed that his health would be improved by wrapping himself up in blankets, and in this state he tripped down the stairs in his house and killed himself.](#) The paper of [Galois](#) which he had taken home to read shortly before his death was never recovered.

SEE ALSO: [Galois](#)

Additional biographies: [MacTutor \(St. Andrews\)](#), [Bonn](#)


Jean Baptiste Joseph Fourier (1768-1830)

had crazy idea (1807):

Any univariate function can be rewritten as a weighted sum of sines and cosines of different frequencies.



...the manner in which the author arrives at these equations is not exempt of difficulties and...his analysis to integrate them still leaves something to be desired on the score of generality and even rigour.

- Don't believe it?
 - Neither did Lagrange, Laplace, Poisson and other big wigs
 - Not translated into English until 1878!
- But it's (mostly) true!
 - called Fourier Series
 - there are some subtle restrictions




Frequency Spectra





- example : $g(t) = \sin(2\pi f t) + (1/3)\sin(2\pi(3f) t)$



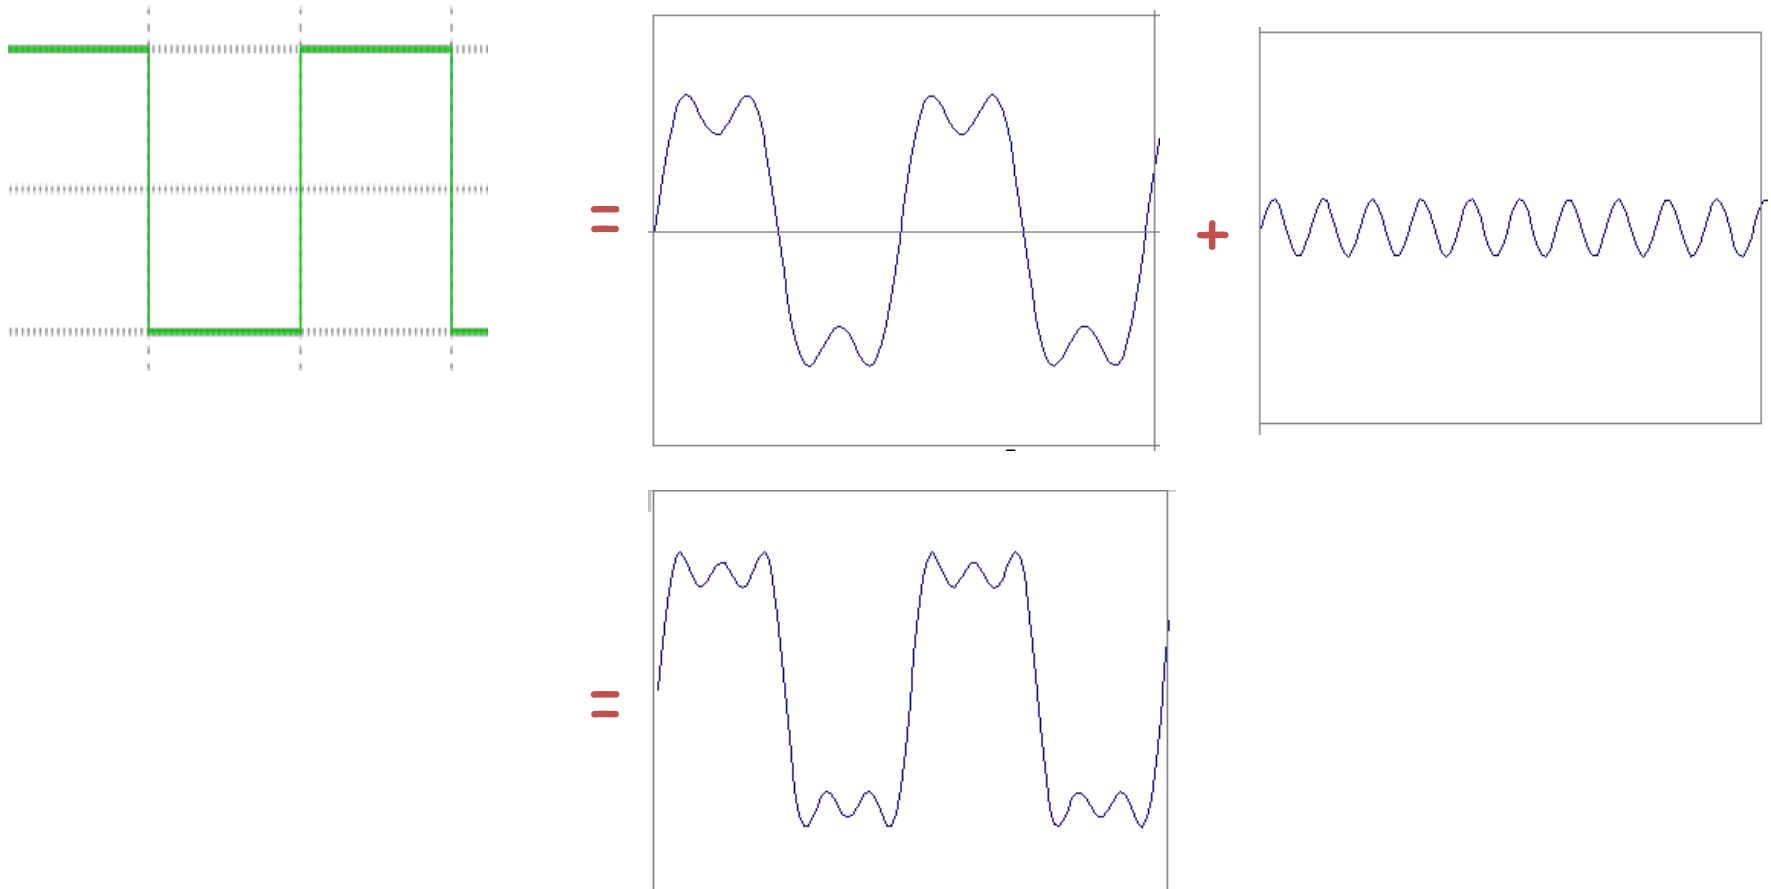
Frequency Spectra



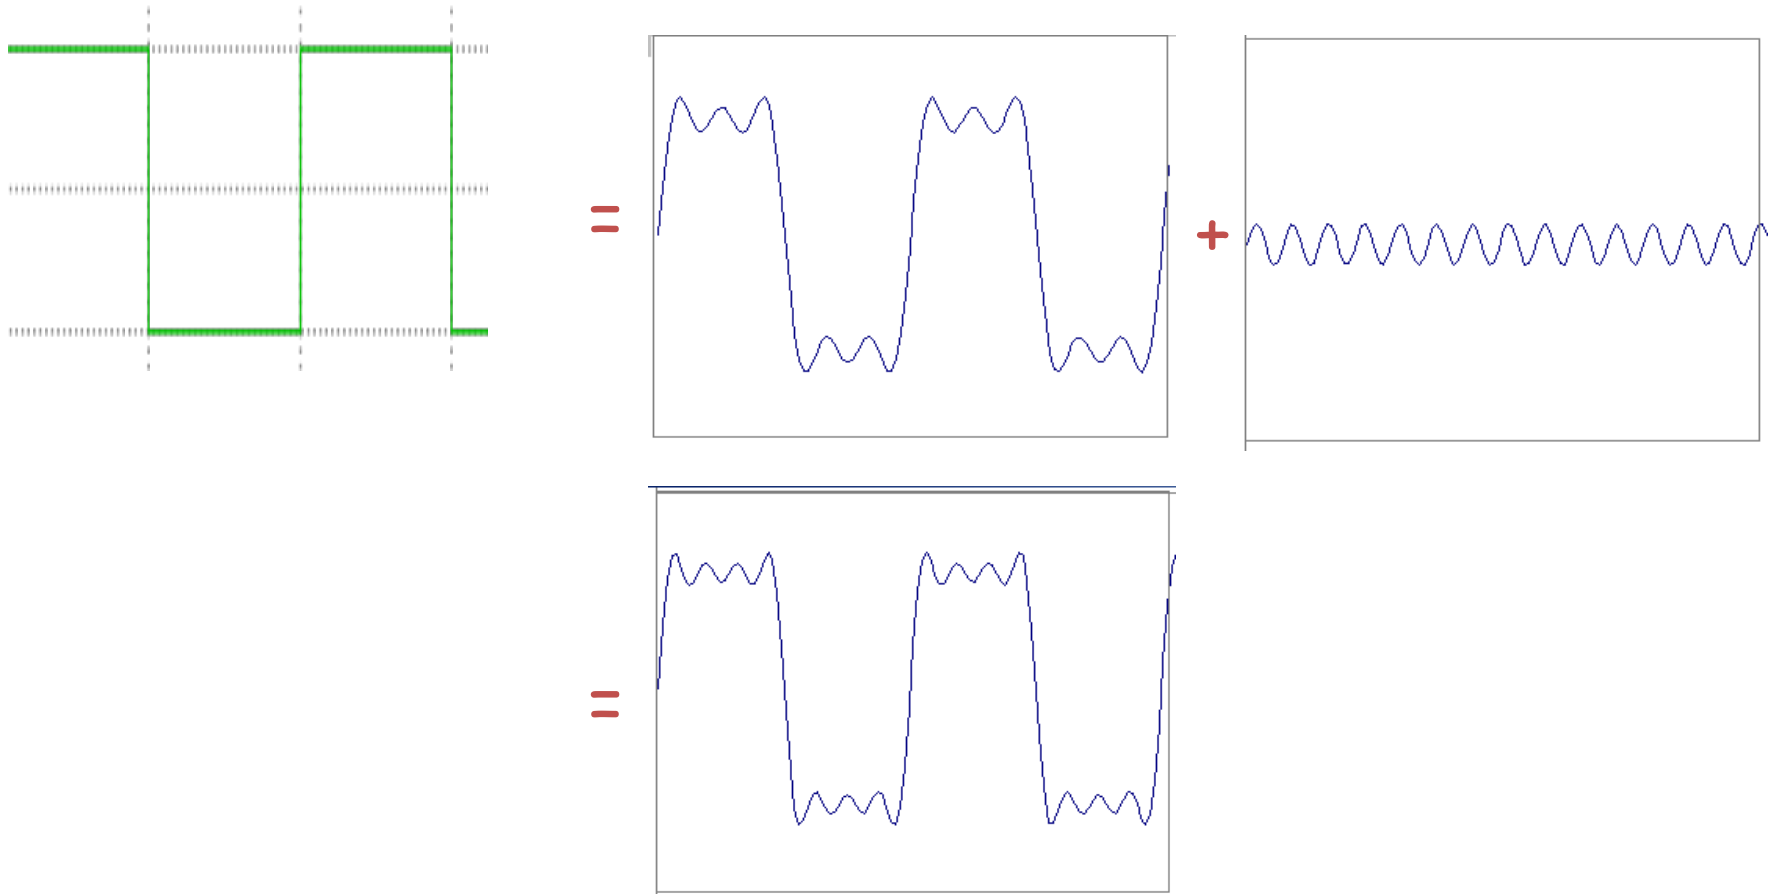
Frequency Spectra




Frequency Spectra




Frequency Spectra




Frequency Spectra




Frequency Spectra



Frequency Spectra





$$= A \sum_{k=1}^{\infty} \frac{1}{k} \sin(2\pi kt)$$






Example: Music

- We think of music in terms of frequencies at different magnitudes






Fourier analysis in images

Intensity Image



Fourier Image



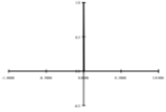
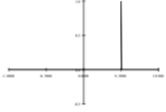



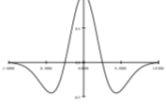
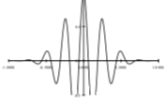
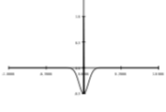
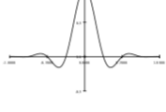
Fourier Transform

- Fourier transform stores the **magnitude** and **phase** at each frequency
 - **Magnitude** encodes how much signal there is at a particular frequency
 - **Phase** encodes spatial information (indirectly)
 - For mathematical convenience, this is often notated in terms of real and complex numbers

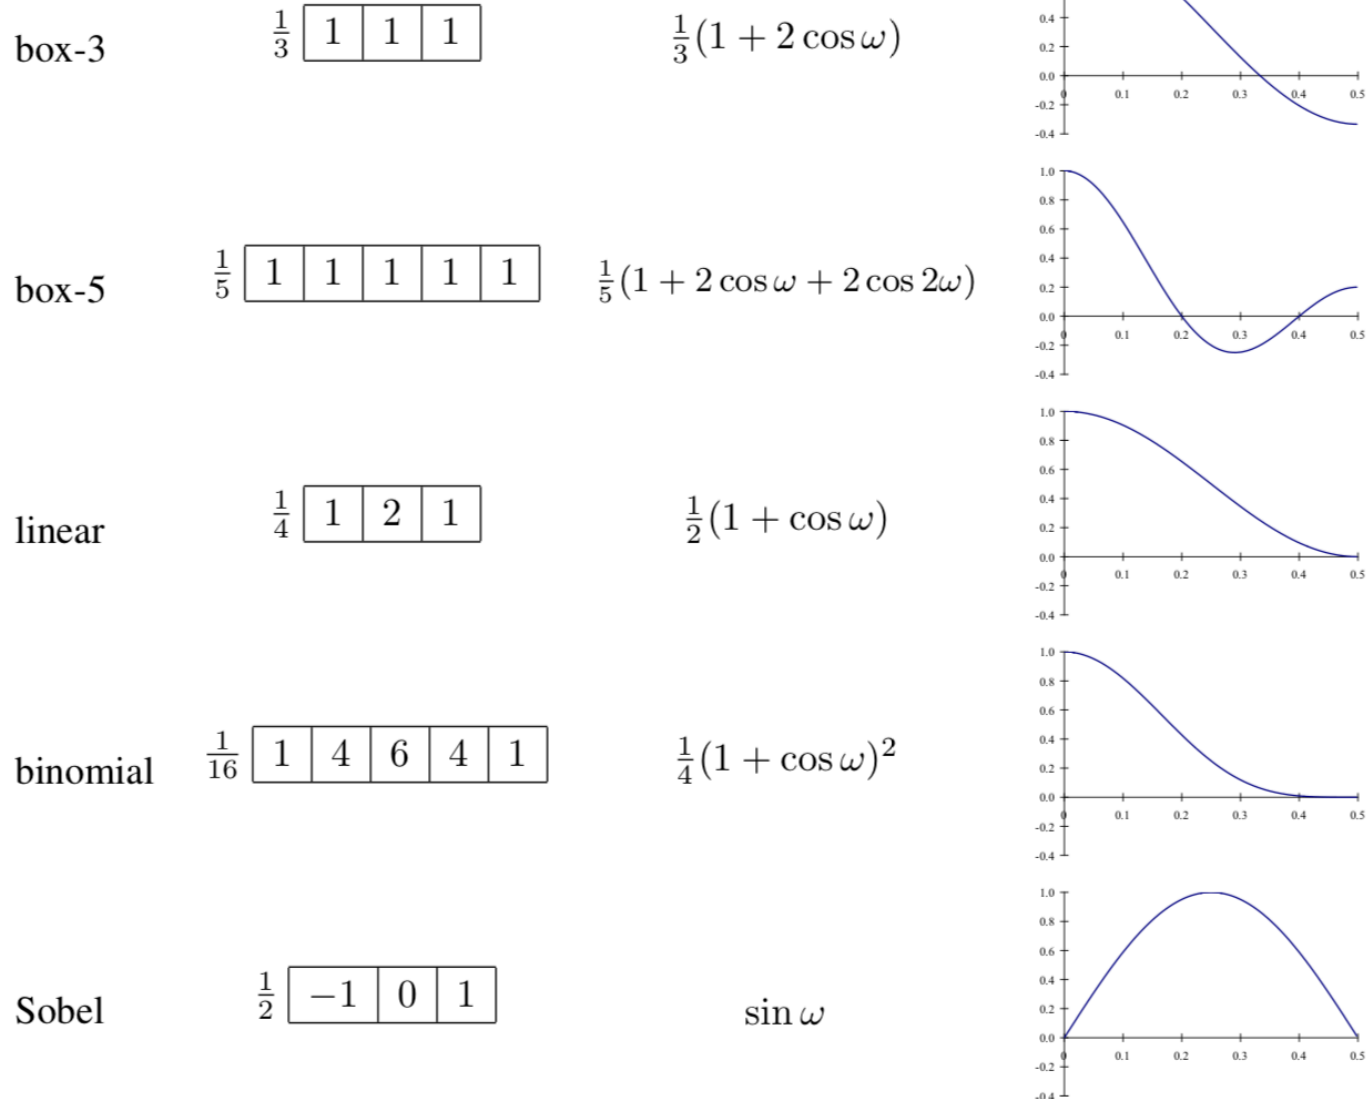
Amplitude: $A = \pm \sqrt{R(\omega)^2 + I(\omega)^2}$

Phase: $\phi = \tan^{-1} \frac{I(\omega)}{R(\omega)}$


Fourier Transform Pairs

Name	Signal	Transform
impulse	 $\delta(x)$	1
shifted impulse	 $\delta(x - u)$	$e^{-j\omega u}$
box filter	 $\text{box}(x/a)$	$a\text{sinc}(a\omega)$
tent	 $\text{tent}(x/a)$	$a\text{sinc}^2(a\omega)$
Gaussian	 $G(x; \sigma)$	$\frac{\sqrt{2\pi}}{\sigma} G(\omega; \sigma^{-1})$
Laplacian of Gaussian	 $(\frac{x^2}{\sigma^4} - \frac{1}{\sigma^2})G(x; \sigma)$	$-\frac{\sqrt{2\pi}}{\sigma} \omega^2 G(\omega; \sigma^{-1})$
Gabor	 $\cos(\omega_0 x)G(x; \sigma)$	$\frac{\sqrt{2\pi}}{\sigma} G(\omega \pm \omega_0; \sigma^{-1})$
unsharp mask	 $(1 + \gamma)\delta(x) - \gamma G(x; \sigma)$	$(1 + \gamma) - \frac{\sqrt{2\pi}\gamma}{\sigma} G(\omega; \sigma^{-1})$
windowed sinc	 $\text{rcos}(x/(aW)) \text{sinc}(x/a)$	(see Figure 3.29)


Fourier Transforms of Filters





Man-made Scene



Can change spectrum, then reconstruct



Low and High Pass filtering



The Convolution Theorem

- The Fourier transform of the convolution of two functions is the product of their Fourier transforms

$$F[g * h] = F[g]F[h]$$


- **Convolution** in spatial domain is equivalent to **multiplication** in frequency domain!

$$g * h = F^{-1}[F[g]F[h]]$$


Filtering in spatial domain

1	0	-1
2	0	-2
1	0	-1


intensity image




Filtering in frequency domain




FFT





\times



$=$




Inverse FFT




Filtering

Why does the Gaussian give a nice smooth image, but the square filter give edgy artifacts?

Gaussian



Box filter




Gaussian


intensity image




filter: gaussian




filtered image




log ft magnitude of image



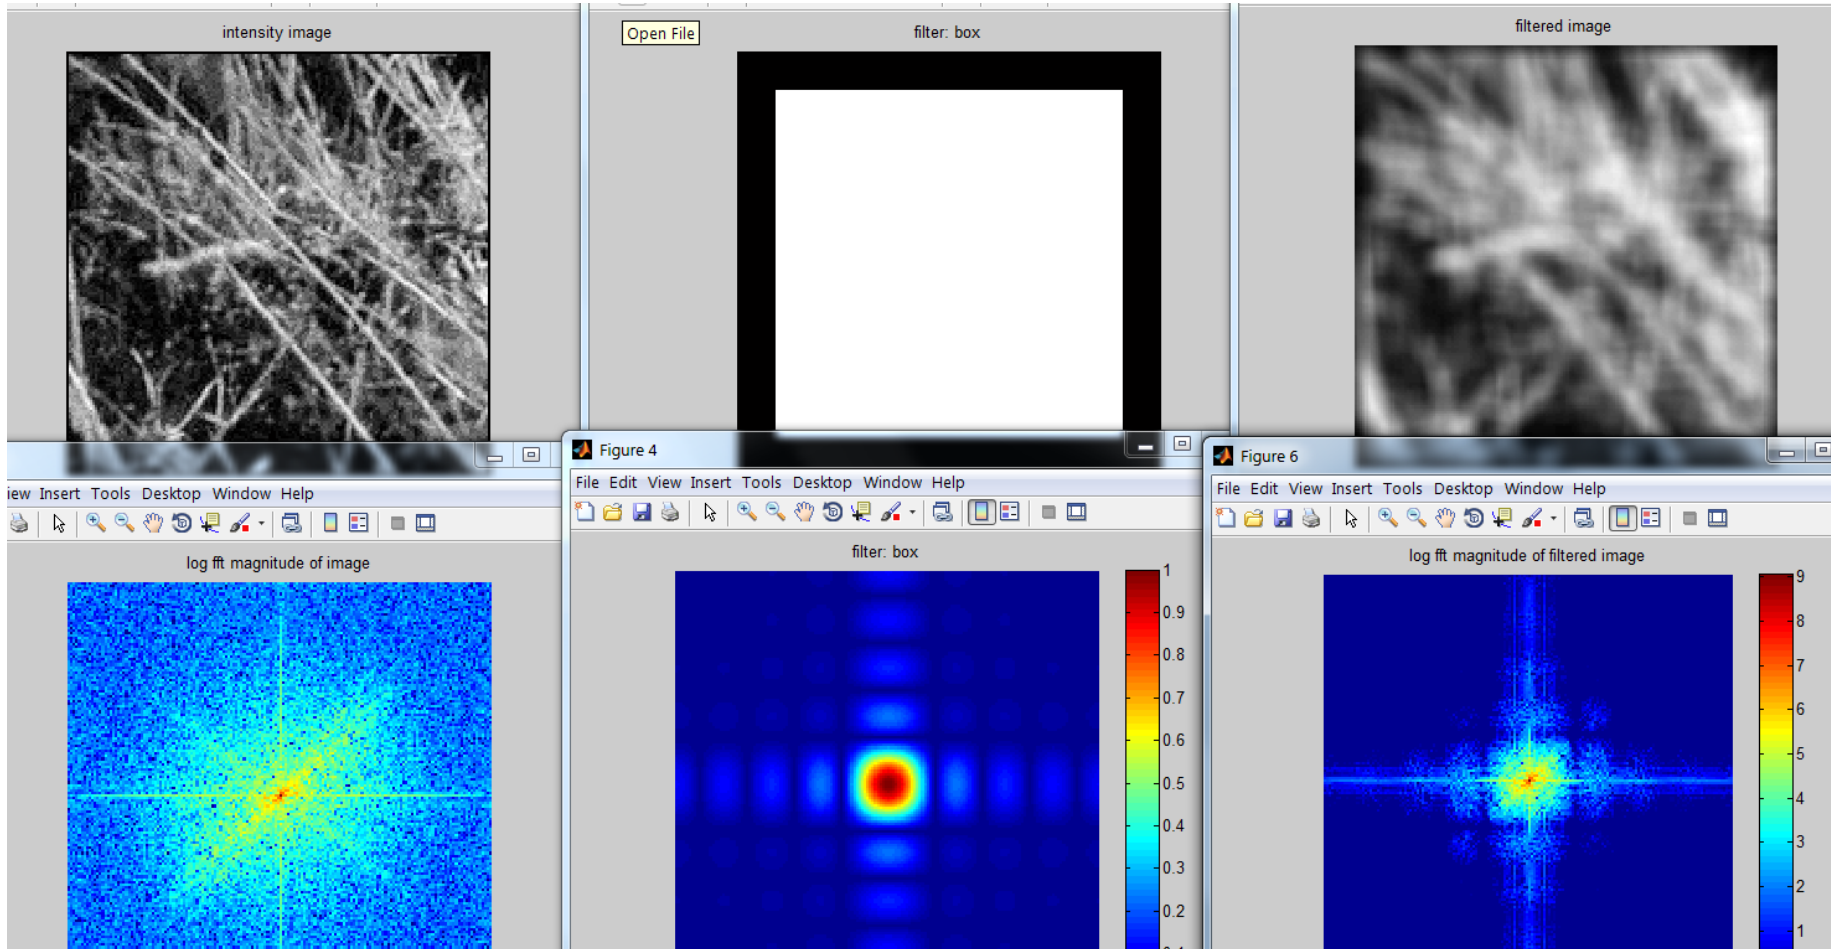
filter: gaussian



log ft magnitude of filtered image



Box Filter



Is convolution invertible?

- If convolution is just multiplication in the Fourier domain, isn't deconvolution just division?
- Sometimes, it clearly is invertible (e.g. a convolution with an identity filter)
- In one case, it clearly isn't invertible (e.g. convolution with an all zero filter)
- What about for common filters like a Gaussian?


But you can't invert multiplication by 0

- But it's not quite zero, is it...


intensity image




filter: gaussian




filtered image




log fft magnitude of image



filter: gaussian




log fft magnitude of filtered image



Let's experiment on Novak



Convolution




=




FFT ↓

FFT ↓


iFFT ↑




*



=




Deconvolution?




iFFT 

FFT 


FFT 




=



· /



But under more realistic conditions




Random noise, .000001 magnitude




iFFT 

FFT 


FFT 




=



· /



But under more realistic conditions




Random noise, .0001 magnitude




iFFT 

FFT 


FFT 



=





· /



But under more realistic conditions


Random noise, .001 magnitude




iFFT ↑

FFT ↓


FFT ↓




=



· /



With a random filter...






Random noise, .001 magnitude



iFFT 

FFT 

FFT 



Deconvolution is hard

- Active research area.
- Even if you know the filter (non-blind deconvolution), it is still very hard and requires strong *regularization*.
- If you don't know the filter (blind deconvolution) it is harder still.