

Reminder

- Project 1 due Wednesday
- Intended to be easy intro project
- What if your filtering is slow?
 - Separable filters (Gaussian)
 - Laplacian Pyramid
- What if my results aren't convincing?




Recap: Fourier domain

2d Fourier Transform

$$\hat{F}(k, \ell) = \frac{1}{\sqrt{MN}} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} F(m, n) e^{-j2\pi\left(k\frac{m}{M} + \ell\frac{n}{N}\right)}$$


Fourier Bases

Teases away fast vs. slow changes in the image.




This change of basis is the Fourier Transform

Fourier Bases



in Matlab, check out: `imagesc(log(abs(fftshift(fft2(im)))));`



Man-made Scene



Can change spectrum, then reconstruct




Low and High Pass filtering




Sinc Filter

- What is the spatial representation of the hard cutoff in the frequency domain?



Frequency Domain

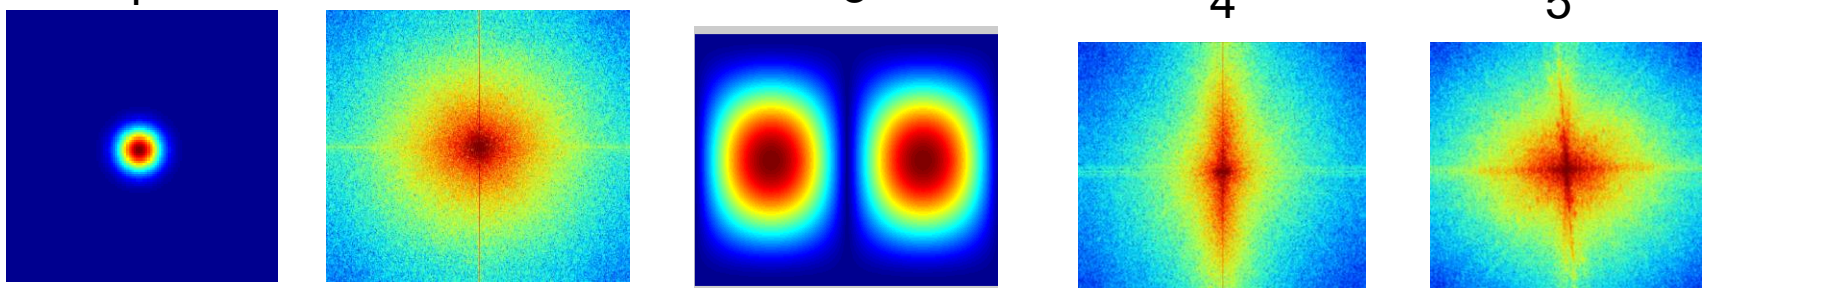


Spatial Domain

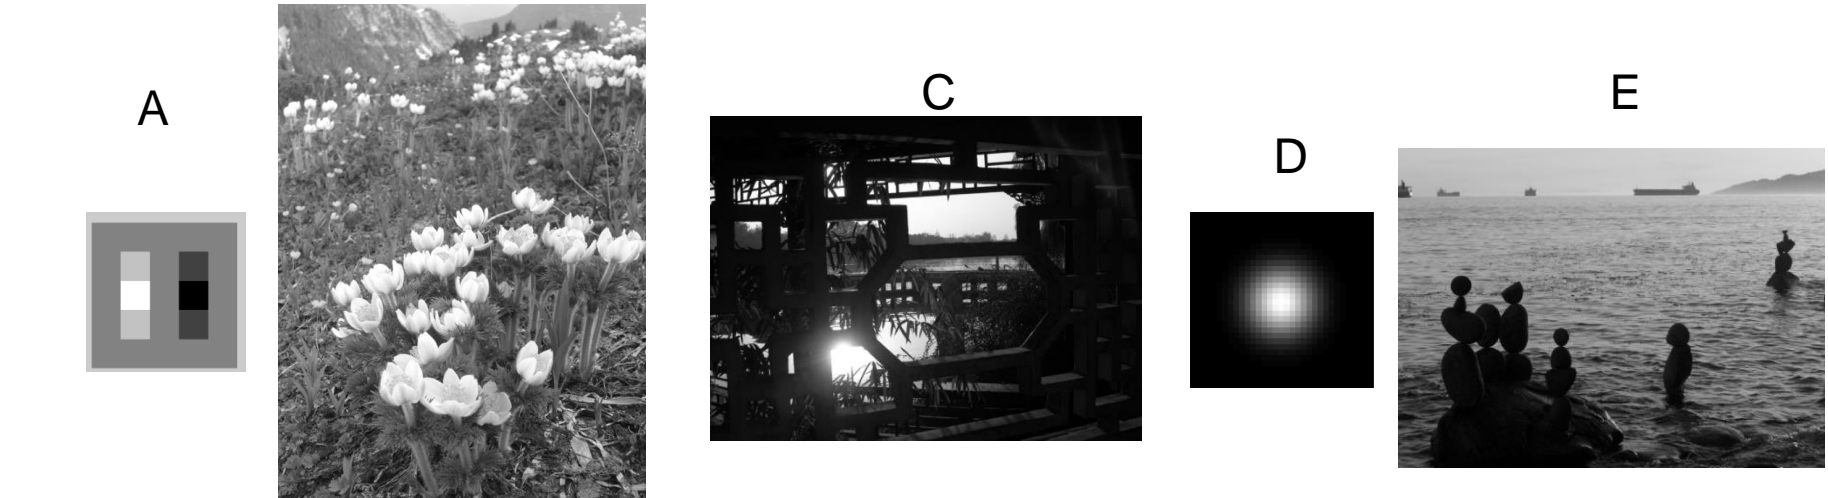
Review

1. Match the spatial domain image to the Fourier magnitude image

1 2 3 4 5



A B C D E



The task is to match the spatial domain images (A-E) to the Fourier magnitude images (1-5). The Fourier magnitude images show the frequency content of the spatial domain images. Image 1 is a single central peak, corresponding to a constant image (A). Image 2 is a central peak with a vertical line, corresponding to a scene with a vertical edge or line (B). Image 3 is two side-by-side peaks, corresponding to a scene with two distinct vertical features (C). Image 4 is a central peak with a vertical line and a horizontal line, corresponding to a scene with both vertical and horizontal features (D). Image 5 is a central peak with both vertical and horizontal lines, corresponding to a scene with both vertical and horizontal features (E).

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




Fourier Matlab demo


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 fit magnitude of image



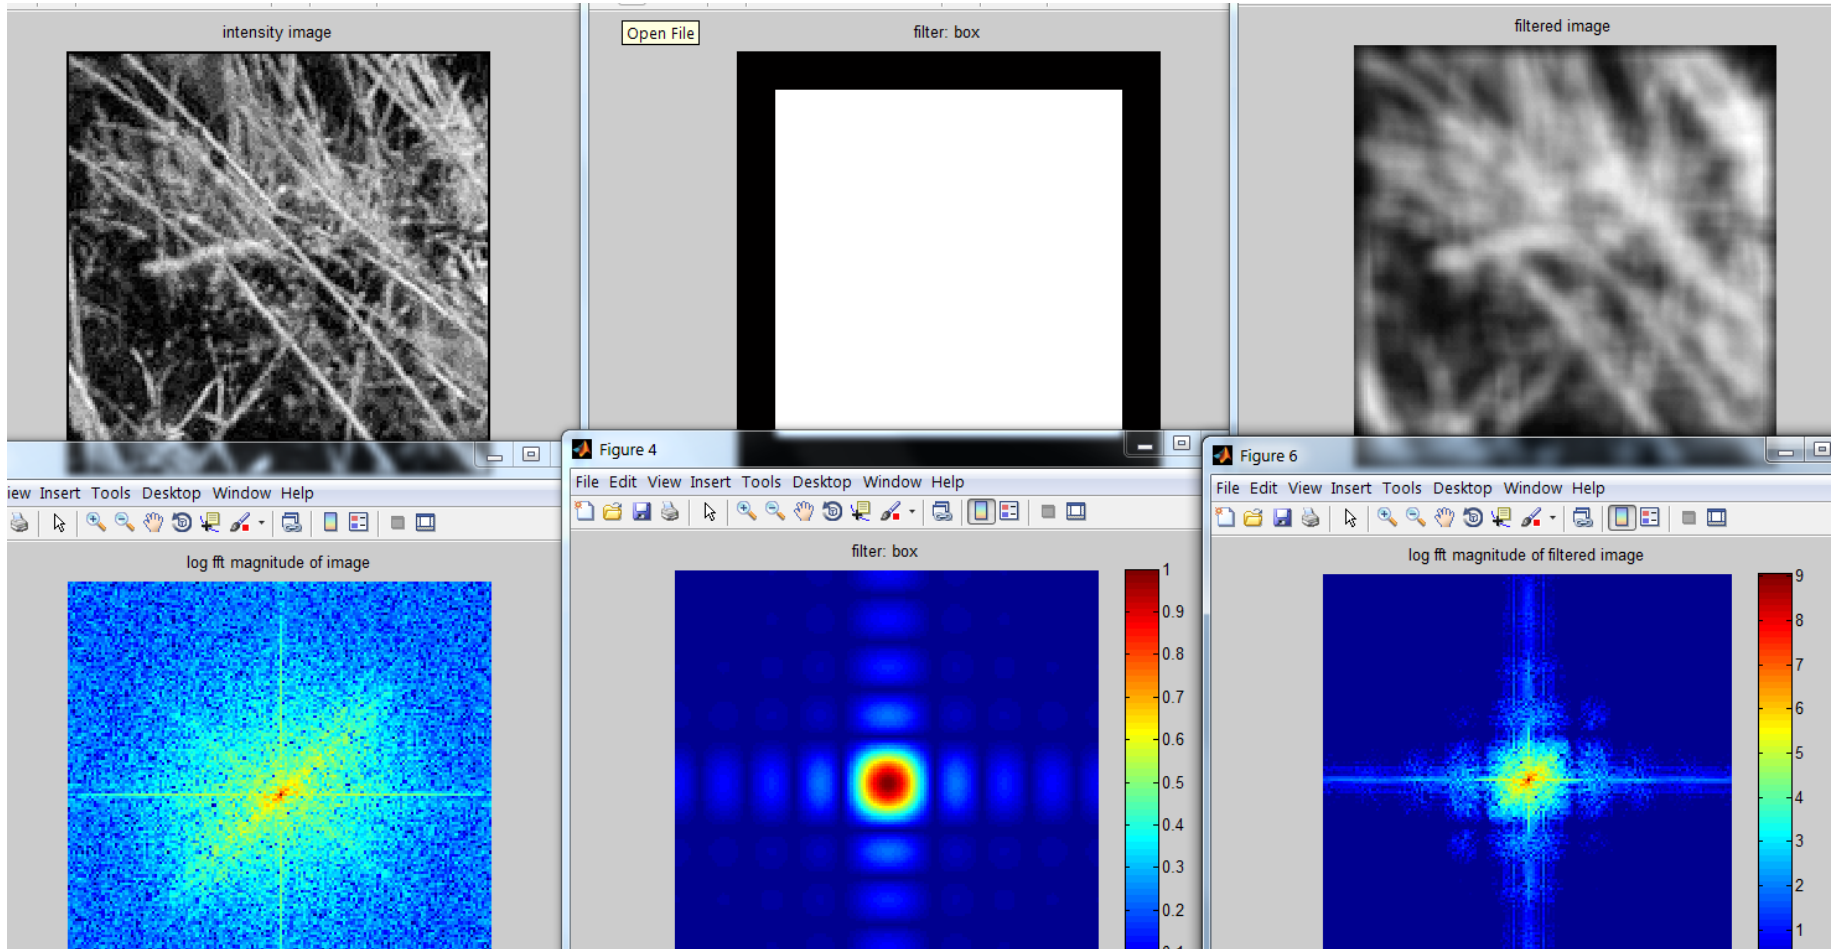
filter: gaussian



log fit 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?

Let's experiment on Novak



Convolution



*


=




FFT ↓

FFT ↓


iFFT ↑




X



=




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.

Sampling

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




Subsampling by a factor of 2



Throw away every other row and column to create a 1/2 size image


Aliasing problem

- 1D example (sinewave):



Aliasing problem

- 1D example (sinewave):




Aliasing problem

- Sub-sampling may be dangerous....
- Characteristic errors may appear:
 - “car wheels rolling the wrong way in movies”
 - “Checkerboards disintegrate in ray tracing”
 - “Striped shirts look funny on color television”

Aliasing in video

Imagine a spoked wheel moving to the right (rotating clockwise).
Mark wheel with dot so we can see what's happening.

If camera shutter is only open for a fraction of a frame time (frame time = $1/30$ sec. for video, $1/24$ sec. for film):



Without dot, wheel appears to be rotating slowly backwards!
(counterclockwise)

Aliasing in graphics



Sampling and aliasing


256x256

128x128

64x64


32x32

16x16



Nyquist-Shannon Sampling Theorem

- When sampling a signal at discrete intervals, the sampling frequency must be $\geq 2 \times f_{\max}$
- f_{\max} = max frequency of the input signal
- This will allow to reconstruct the original perfectly from the sampled version



Anti-aliasing

Solutions:

- Sample more often
- Get rid of all frequencies that are greater than half the new sampling frequency
 - Will lose information
 - But it's better than aliasing
 - Apply a smoothing filter

Algorithm for downsampling by factor of 2

1. Start with image(h, w)

2. Apply low-pass filter

```
im_blur = imfilter(image, fspecial('gaussian', 7, 1))
```

3. Sample every other pixel

```
im_small = im_blur(1:2:end, 1:2:end);
```

Anti-aliasing


256x256

128x128

64x64

32x32

16x16




256x256

128x128

64x64

32x32

16x16



Subsampling without pre-filtering



1/2



1/4 (2x zoom)



1/8 (4x zoom)

Subsampling with Gaussian pre-filtering



Gaussian $1/2$



G $1/4$



G $1/8$