ECEU692 Intro to SSI Course Notes

Some selected topics in signal and image processing for SSI(3)

Profs. DiMarzio and Brooks
Northeastern University

Spring 2004

Outline of those Selected Topics

- Impulse responses and convolution (review)
- Fourier Transforms and frequency content (review)
- Convolution / Multiplication properties (review)
- Sampling and quantization and their consequences (review)
- Anti-aliasing filtering (review)
- 2-dimensional signals and all of the above
- The Projection-Slice theorem and tomography: basic idea
- **What does Fourier phase tell us ?**
- **Skip:** Power spectra and autocorrelations
- Resolution and Extent
  - apodization and its dual in time
  - the time-bandwidth product
- Relationships among varieties of Fourier representations
- Hyperspectral Images as signals

Fourier Phase and Information

- As we know, Fourier Transforms are generally complex-valued.
- The magnitude tells us the distribution of energy with frequency:
  - Peaks tells us frequencies where there is a lot of energy.
  - Nulls/valleys tells us frequencies where is no/little energy.
  - The breadth or narrowness of a peak or valley tells us if a component is "broad-band" or "narrow-band".
- For instance, when looking at a "safe" sampling rate, we look at where the magnitude of the Transform is zero (or sufficiently small).
- What about the phase? What does that tell us ?

An Experimental Approach

- It turns out that Fourier phase tells us a lot about timing / location — and this turns out to be particularly information-bearing in images.
- Let's illustrate with an example: a matlab script that

1. Reads in an image into matlab.
2. Computes the Fourier Transform of the image.
3. Replaces the
   (a) phase of the transform with random phase values, while keeping the orignal magnitude,
   (b) magnitude of the image with random magnitude values, while keeping the orignal phase,
   (c) magnitude of the image with a constant magnitude, while keeping the orignal phase.

Some Matlab Code to Do This

```
%% magorphase.m
%%
%% Script to create some figures comparing the importance of magnitude and
%% phase in the Cal band image so familiar to the ECE 1467 W02 sufferers.
%%
%% Requires the file m10067_1.tif to be in the current directory or on
%% Matlab's path.
%%
%% DH Brooks, Feb 2002

%% load image --> variable data and compute its 2d Fourier Transform (FT)

[data,map] = imread('m10067_1');
fdata = fft2(data);

%% create a new FT fmag with same magnitude but random phase

fmag = abs(fdata).*exp(j*(rand(size(fdata))*2*pi - pi));

%% create a new FT frandmagphase with same phase but random magnitude

frandmagphase = rand(size(data)).*exp(j*angle(fdata));

%% create a new FT fphase with constant magnitude but same phase

fphase = 1*exp(j*angle(fdata));

%% in 3 figure windows, show images of inverse FT's of each of these images

figure(1)
imagesc(abs(ifft2(fmag)))
title('Keep magnitude, random phase')
axis image
colormap('default')

figure(2)
imagesc(abs(ifft2(frandmagphase)))
title('Keep phase, random magnitude')
axis image
colormap('default')

figure(3)
imagesc(abs(ifft2(fphase)))
title('Keep  phase, constant magnitude')
axis image
colormap('default')
```

The Image

The Cal Band Entertains through the storm.

And The Results . . .

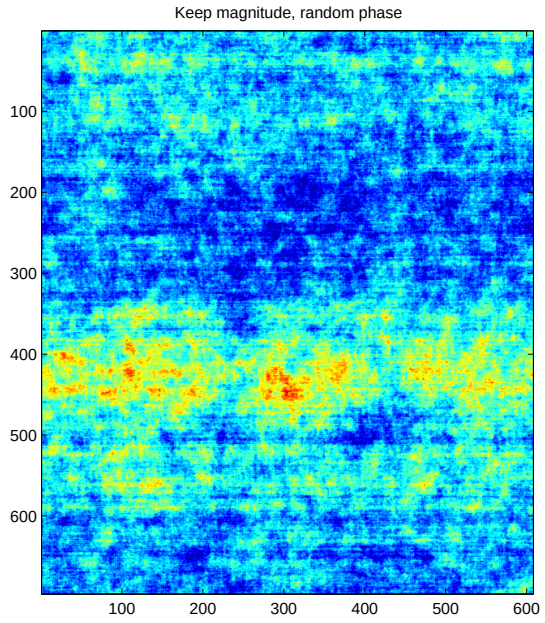

Keep phase, random magnitude

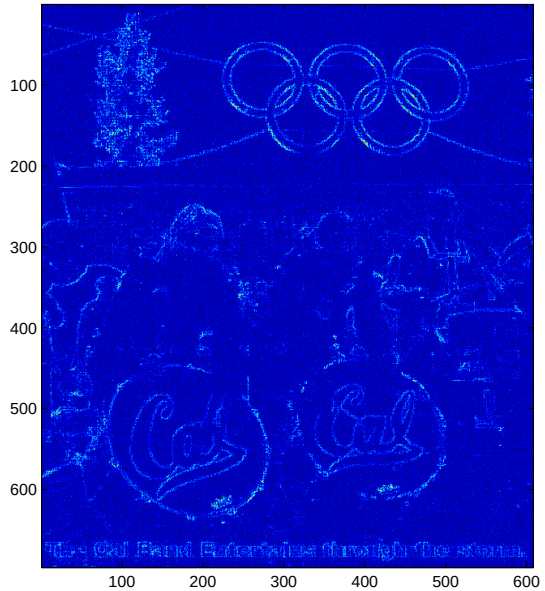

Keep phase, constant magnitude

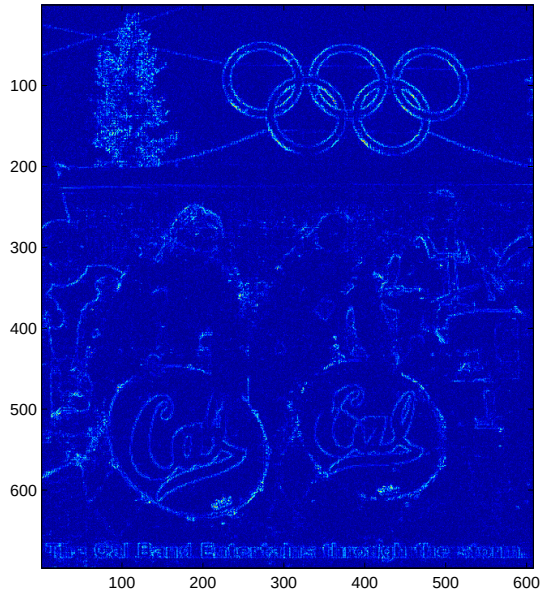

Spectral Resolution

Remember the lecture and homework problem on the diffraction limit ?

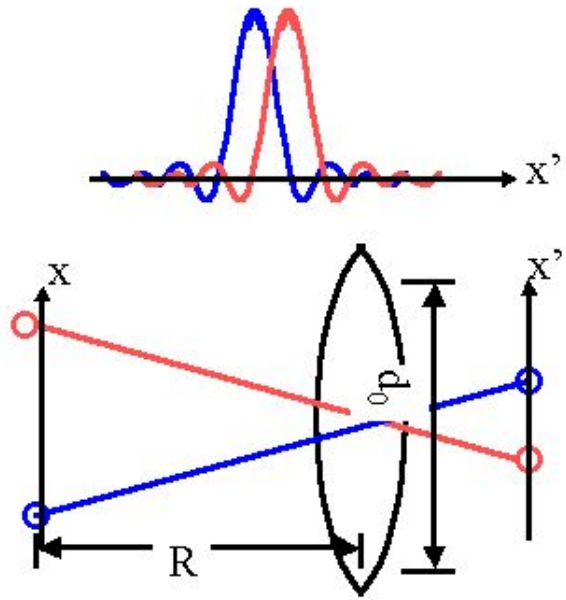

What Are the Implications?

- Two important points to remember here:
  1. What we see in the pupil plane (in other words, at the lens) is like a spatial Fourier Transform of the image that went into the lens,
  2. The physical size of the lens "cuts-off" the higher-frequency components of this Fourier Transform, and
  3. What we see after the lens (in the image plane) is thus a "low-pass filtered" version of the original image.
  4. In particular, the distance between two (spectral) points in the image thats we want to resolve is inversely proportional to how much of the image Fourier Transform "gets through" the lens, that is $\boldsymbol{d_0}$ (actually, in this case it's really inversely proportional to $\boldsymbol{d_0/R}$).
- So this means that the bigger the lens (or the more of the original image that "fits" through the $\boldsymbol{d_0}$ sized aperture), the better our resolution in the space (that is, the farther apart our two peaks are).
- Let's look at a 1D signal analysis version of the same problem (the spectral resolution problem).

A Signal Model for A Finite-Length Signal

- Suppose I have a signal which may have two closely-spaced frequency components in it: we'll look at the relationship between the length of time over which we measure a signal and the resolution of its Fourier Transform.
- Let's start with the following model:

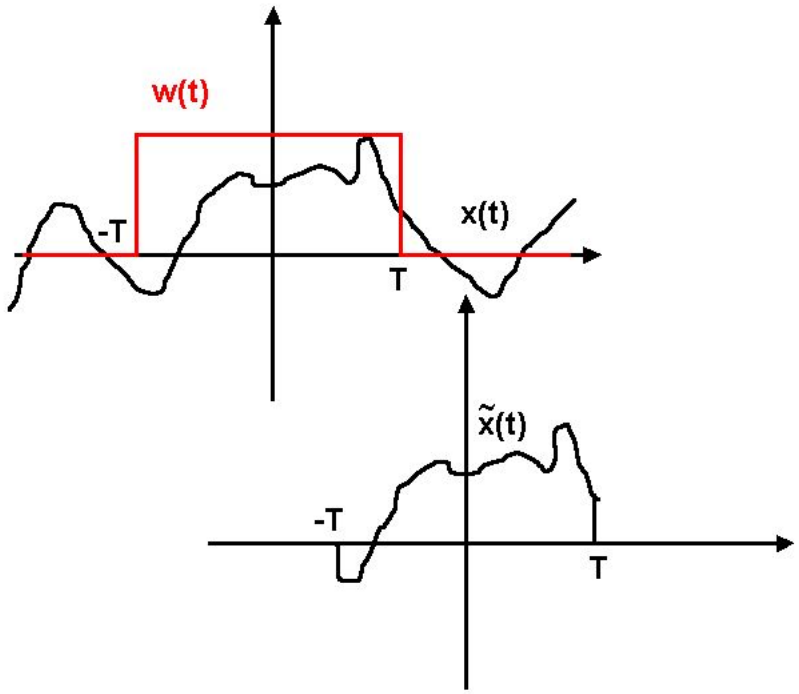

Completing the Time-Domain Model

- We have a signal $x(t)$ which we observe for all time (from $-\infty$ to ∞).
- We multiply $x(t)$ by a "window" signal $w(t)$ which is equal to 1 over a finite time (say, from $t = -T$ to $t = T$), and is 0 elsewhere to get a new signal

$$\tilde{x}(t) = w(t) \cdot x(t).$$

- If a signal's frequency content varies in time (say, speech or music), we can use a succession of shifted windows.
- Note that we can turn the situation on its head:
  1. Start with a finite-length signal $\tilde{x}(t)$.
  2. It could be modeled as **any** signal $x(t)$ which is the same as $\tilde{x}(t)$ on the interval $[-T, T)$.
- This ambiguity turns out to be related to limited resolution in the frequency domain.

What Happens in the Fourier Domain?

- Using the "multiplication / convolution" property, we have

$$\tilde{X}(F) = W(F) * X(F),$$

  where $W(F)$ is the Fourier Transform of the window function $w(t)$.

- Thus $W(F)$ determines how the "true" Fourier Transform $X(F)$ and the transform of the truncated signal, $\tilde{X}(F)$, are lreatd

- So what does $W(F)$, for a "rectangular time window" $w(t)$, look like?

$$W(F) = \frac{1}{\pi F}\sin(2\pi TF) = 2T\mathrm{sinc}(2\pi TF).$$

- Note that
  - Zeros when $2\pi TF = k\pi$ for integer $k \neq 0$.
  - Zeros get closer as observation grows, farther apart as it shrinks.
  - Peak gets higher as observation interval grows, smaller as it shrinks.
  - as $T \to \infty$, $W(F)$ goes to a delta function.

- Be careful about the Matlab `sinc` function !

Three Different Examples of $W(F)$

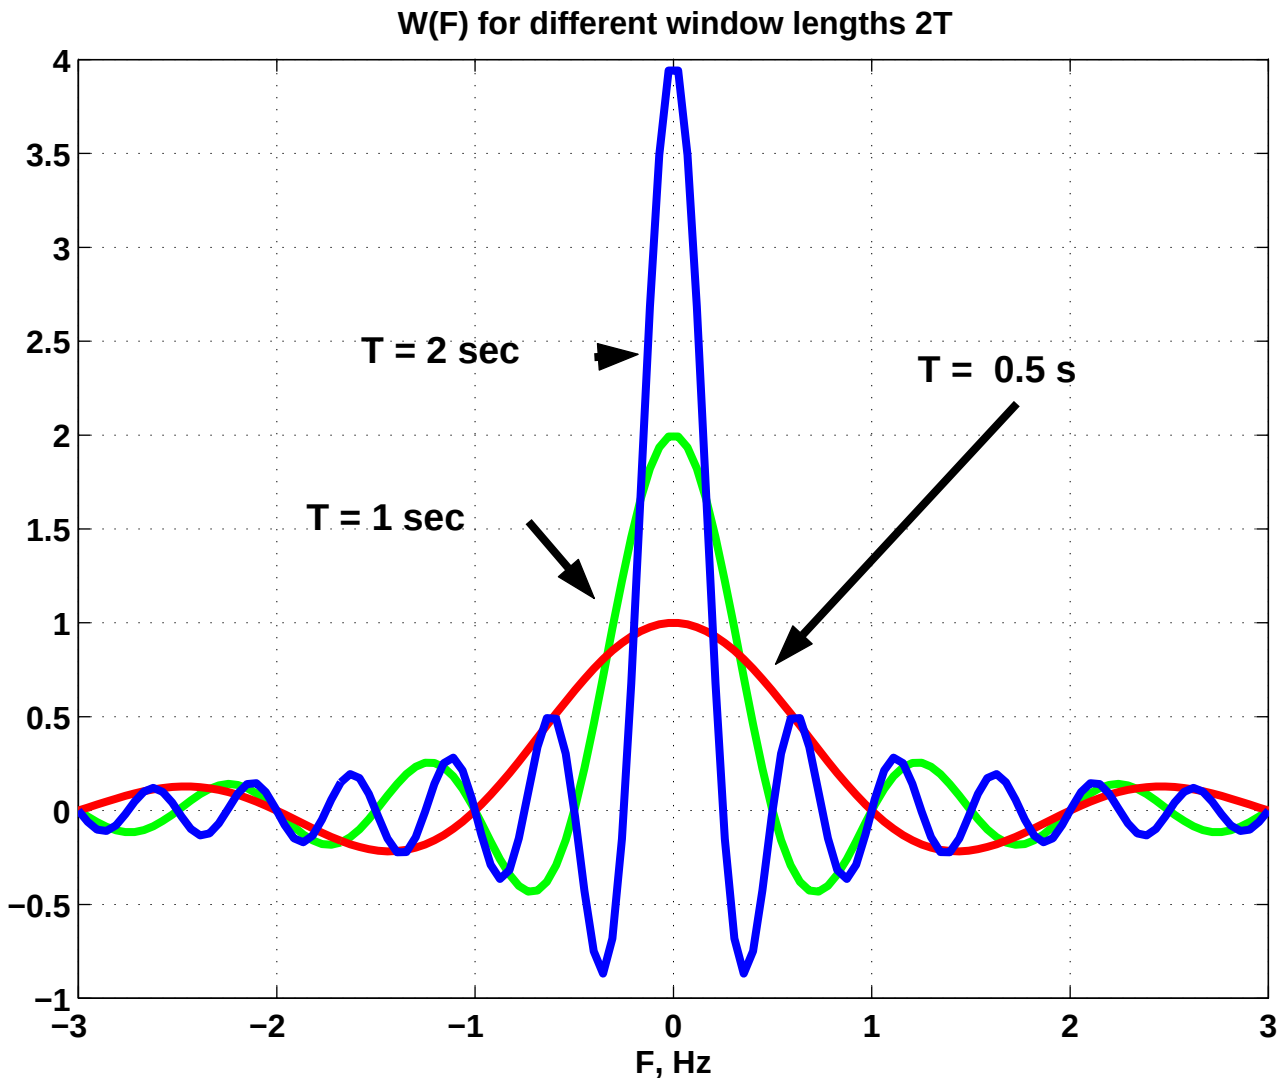

The Effect of $W(F)$ on Resolution

Let's look at a simple case: $T = 2$s, and $x(t)$ is a sum of two cosines at $.85$ and 1.15 Hz.

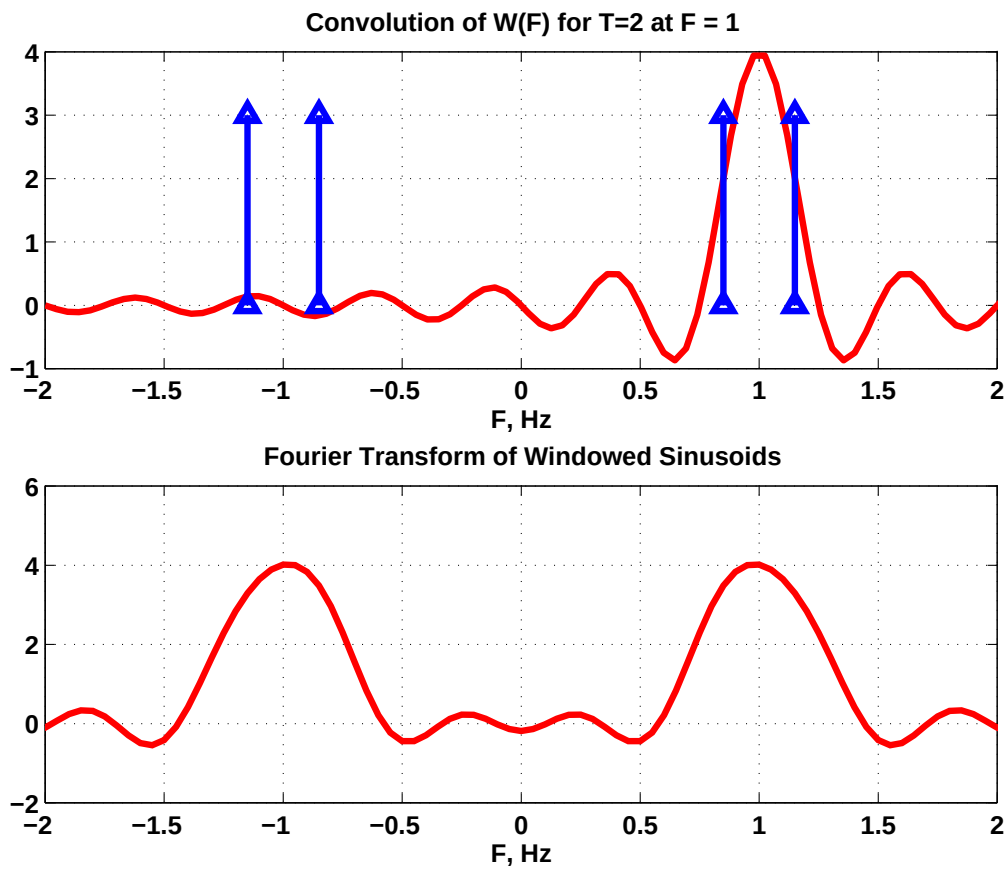

Conclusions

- Spectral resolution is limited by the length of time over which we observe the signal.
- Analogous to the diffraction limit, where the size of the aperture (lens) limits resolution. Here we look at frequency resolution limited by time extent; in diffraction it's space resolution limited by frequency extent.
- Note that this result also implies that a time-limited signal can't be band-limited, since for finite T the convolution of $X(F)$ with $W(F)$ will be non-zero at every F. This has implications for the Shannon sampling theorem.
- In fact, with appropriate definitions we can lower-bound the product of the time-width and the bandwith of a signal. This says that you can't have both arbitrarily small bandwidth (locate the signal in frequency) **and** arbitrarily small time-width (locate it in time too).
- This is one version of what's called the "Uncertainty Principle of the Fourier Transform".

A Word on Signal Types and Fourier Transforms

well, actually, quite a few words . . .

- Let's stick with 1-D signals. Two basic types
  1. continuous in time, $x(t)$, and
  2. discrete in time, $x(nT)$ or simply $x(n)$
- Two flavors of each type
  1. periodic, $x(t) = x(t + T_P)$ for some T_P and all t, equivalent for discrete signals, and
  2. aperiodic
- Two flavors of Fourier representations
  1. Series — sums of discrete (distinct, separated, harmonically related) sinusoids. *Discrete in frequency*
  2. Transforms — *Continuous in frequency*

Signals and Transform Pairs

- A Fourier representation for each flavor of each type of signal
  1. continuous in time, periodic: (Continuous) Fourier Series: discrete in frequency and itself aperiodic
  2. continuous in time, aperiodic: Continuous Time Fourier Transform: continuous in frequency and itself aperiodic
  3. discrete in time, periodic: Discrete Fourier Series: discrete in frequency and itself periodic
  4. discrete in time, aperiodic: Discrete-Time Fourier Transform: continous in frequency and itself periodic
- From a practical viewpoint, we usually calculate the last one, and only at a finite, discrete set of frequencies. Usual trick
  - sample it at regularly spaced (harmonically related) discrete frequencies.
  - result is called Discrete Fourier Transform (DFT): closely related to Discrete Fourier Series: discrete in frequency and itself periodic
  - Usually calculated via a fast algorithm: Fast Fourier Transform (FFT)

Some General Observations

- Frequency for discrete signals always repeats with period 2π: $e^{j(\omega+2\pi k)m} = e^{j\omega m}$
- Discretization and periodicity:
  1. Discrete (sampled) in one domain means periodic in the other
  2. Continuous in one domain means aperiodic in the other
  3. Any time we have sampling in one domain, ie periodicity in the other, we need to worry about aliasing in the "other" (periodic) domain.
  4. Note: if the variable is discrete, the formula has a sum. If it is continuous, the formula has an integral.
  5. Note: if variable is periodic, sum (integral) only over a period.
- This extends pretty much perfectly to 2 or 3 dimensional signals
- What you saw in Discrete Systems: the Discrete-Time Fourier Transform

$$X(\omega) = \sum_n x(n)e^{-j\omega n}$$

$$x(n) = \left(\frac{1}{2\pi}\right) \int_{-\pi}^{\pi} X(\omega)e^{j\omega n}$$

Another Difference Between Discrete and Continuous

- For discrete systems, we defined an "impulse sequence" $\delta(n)$ $(\delta(m,n))$
- Response to this "simple input" tests system, we apply it to all inputs.
- What signal plays this role in continuous signal/system analysis ?
- The problem is that if a continuous signal is non-zero only at one isolated time instant (say, $t = 0$), then it has zero energy and is not physically realizable.
- So we need a signal which has non-zero energy in the limit as you go to smaller and smaller time intervals: $\Rightarrow$ "height" goes to ∞.
- The result is a signal called the "impulse" or "Dirac delta", denoted $\delta(t)$. Mathematically it's quite complicated: called a "generalized" or "distribution" function and technically defined through its action via integration—think of the signal whose integral is a unit step function: $u(t) = 1$, $t > 0$ and $u(t) = 0$, $t < 0$.
- However the intuition developed about impulse responses for discrete time follow the same way for continous time.

And Finally . . .

Some comments on notation:

1. If no limits given on integral or sum, it is generally assumed to go from $-\infty$ to ∞.
2. Following Proakis and Manolakis book, we will try to use
   (a) upper case $\Omega = 2\pi F$ for frequencies of continuous signals (ie aperiodic in frequency)
   (b) lower case $\omega = 2\pi f$ for frequencies of discrete signals (ie periodic in frequency)

HSI Data Cubes as Signals

Three types of variables:

- Space (over the image plane)

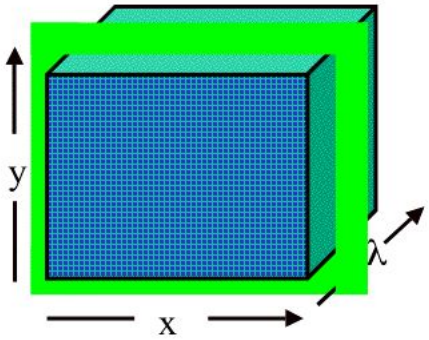

- Wavelength
- Time
  - within one frame
  - at one pixel across data cubes
  - exact details depend on how data is acquired (eg snapshot, pushbroom, . . .)
  - here we'll assume OSL HSI imager adquisition

Some Signal Processing Operations of Interest

- We want a representation we can store on a computer. This means we need to sample in time, space, and wavelength
- We may want to smooth or filter out temporal or spatial noise (low-pass filter)
- We may want to enhance rapid changes in time or space (high-pass filter)
- We may want to take more/less or even average / subsample in wavelength.
- We may want to look for dominant features
  - in space at one wavelength
  - in each pixel across wavelengths
  - in each pixel and wavelength across time
  - various combinations of the above
- and more . . .