ELEN 4810
Digital Signal Processing

Introduction

John Wright
Electrical Engineering
Columbia University

Welcome to 4810, Digital Signal Processing

Signal: loosely, any information bearing function:

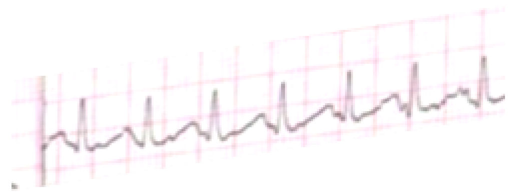

Processing:

Sampling
Denoising / superresolving
Analyzing / interpreting / classifying

Digital Signal Processing:

Acquiring, enhancing, analyzing, and synthesizing signals using a computer, or any other platform that works with

Finitely many samples
Finite precision arithmetic

EKG signal from [Serra et. al., '07]

Why work in the digital domain?

Classical arguments

Flexibility

Programming is usually easier than implementing new hardware

Advances in computing power

Reproducibility:

Same data, same result

Power / precision

Attractive tradeoffs vis. analog implementations

Plus a few recent developments:

New computing platforms

Your phone is a camera, microphone and powerful digital computer!

Data as a resource

The most effective algorithms for signal classification, denoising, superresolution, ect., use massive amounts of training data

What is this class really about?

A few **core mathematical tools** for thinking about discrete time signals and systems.

Introduction: Applications, review of complex math, basic sequences
Time Domain: LTI systems, convolution, impulse response
Frequency Domain: Discrete Time Fourier Transform (DTFT), DFT, FFT, STFT
Sampling: Shannon-Nyquist, Analog-Digital conversion

Midterm exam

Z-transform: Definition, poles, zeros, application in system analysis
Filters: Magnitude, phase response, linear/generalized linear phase
Filter design: FIR design tools, IIR design tools

Special topics: Adaptive filtering
Optimization & learning based signal processing
Applications

Final exam

Motivation (1): Camera Shake Removal

The picture I want:

The picture actually get:

Example from [Wipf + Zhang, *Revising Bayesian Blind Deconvolution*, 2013]

Motivation (1): Camera Shake Removal

The picture I want:

$f(t)$

The picture actually get:

$$= \sum_{\text{translations } \tau} \kappa(\tau) f(t-\tau)$$

Example from [Wipf + Zhang, *Revising Bayesian Blind Deconvolution*, 2013]

Motivation (1): Camera Shake Removal

The picture I want:

The picture actually get:

Example from [Wipf + Zhang, *Revising Bayesian Blind Deconvolution*, 2013]

Motivation (1.5): Scanning Tunneling Spectroscopy

The electron microscope gives:

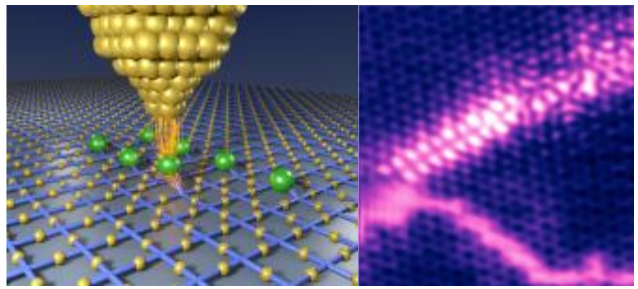

Scanning Tunneling Microscope

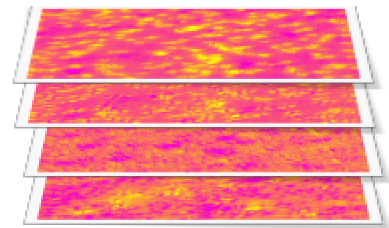

STS Data Cube

Space → Energy

Defects:

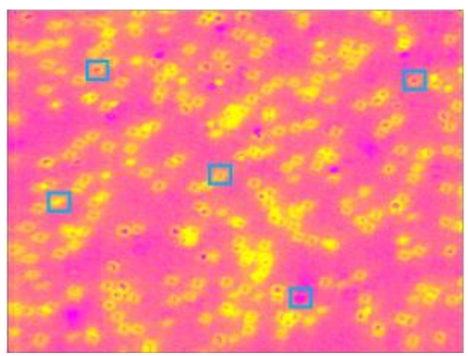

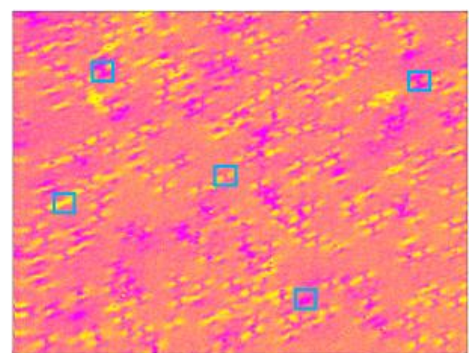

Want to determine where the material defects are located, and what the basic defect "signature" is…

Joint work with Y. Lau, Y. Zhang, S. Cheung, A. Pasuapathy

Motivation (1.5): Scanning Tunneling Spectroscopy

The electron microscope gives:

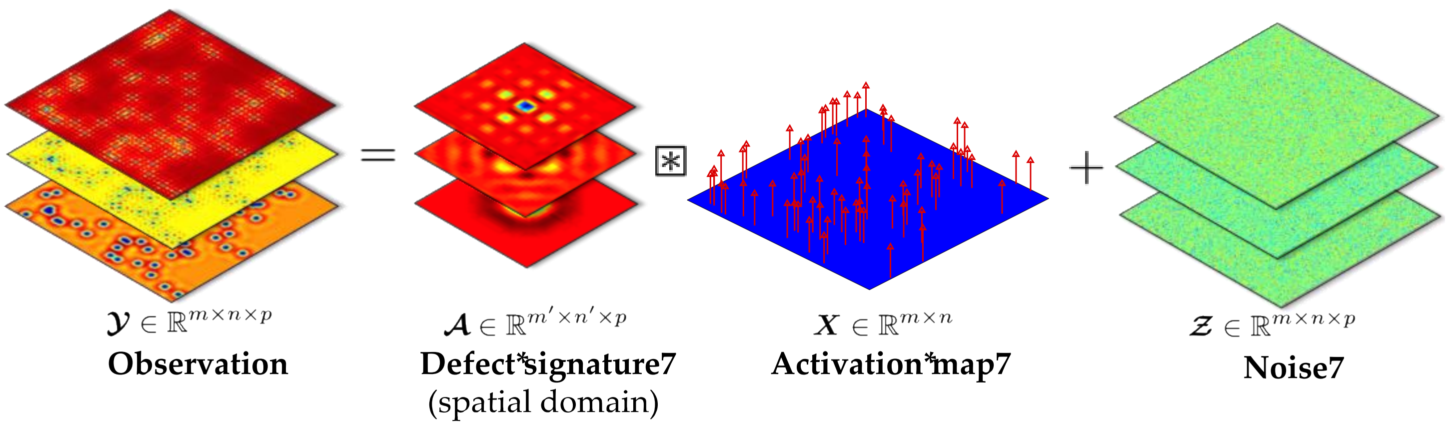

Want to determine where the material defects are located, and what the basic defect "signature" is…

Joint work with Y. Lau, S. Cheung, A. Pasuapathy

Motivation (1.5): Scanning Tunneling Spectroscopy

The electron microscope gives:

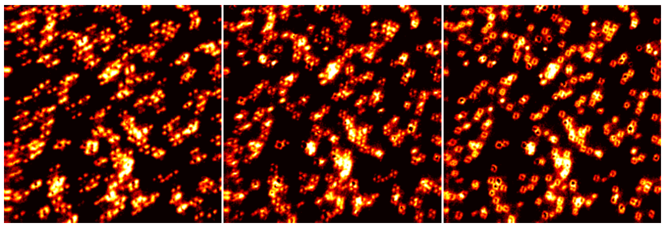

Data $\mathcal{Y}$ (shown: 3 slices of 22 total)

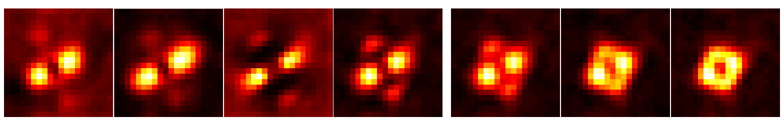

Estimated Defect $\widehat{\mathcal{A}}$
(shown: 7 slices of 22 total)

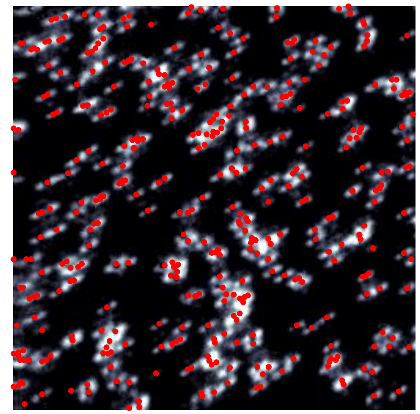

Estimated Defect Locations from $\widehat{\boldsymbol{X}}$
(overlaid on top of one slice of the data $\mathcal{Y}$)

Want to determine where the material defects are located, and what the basic defect "signature" is…

Joint work with Y. Lau, S. Cheung, A. Pasuapathy

Motivation (2): Spectrum Sensing

New communications technology is spectrum-hungry…

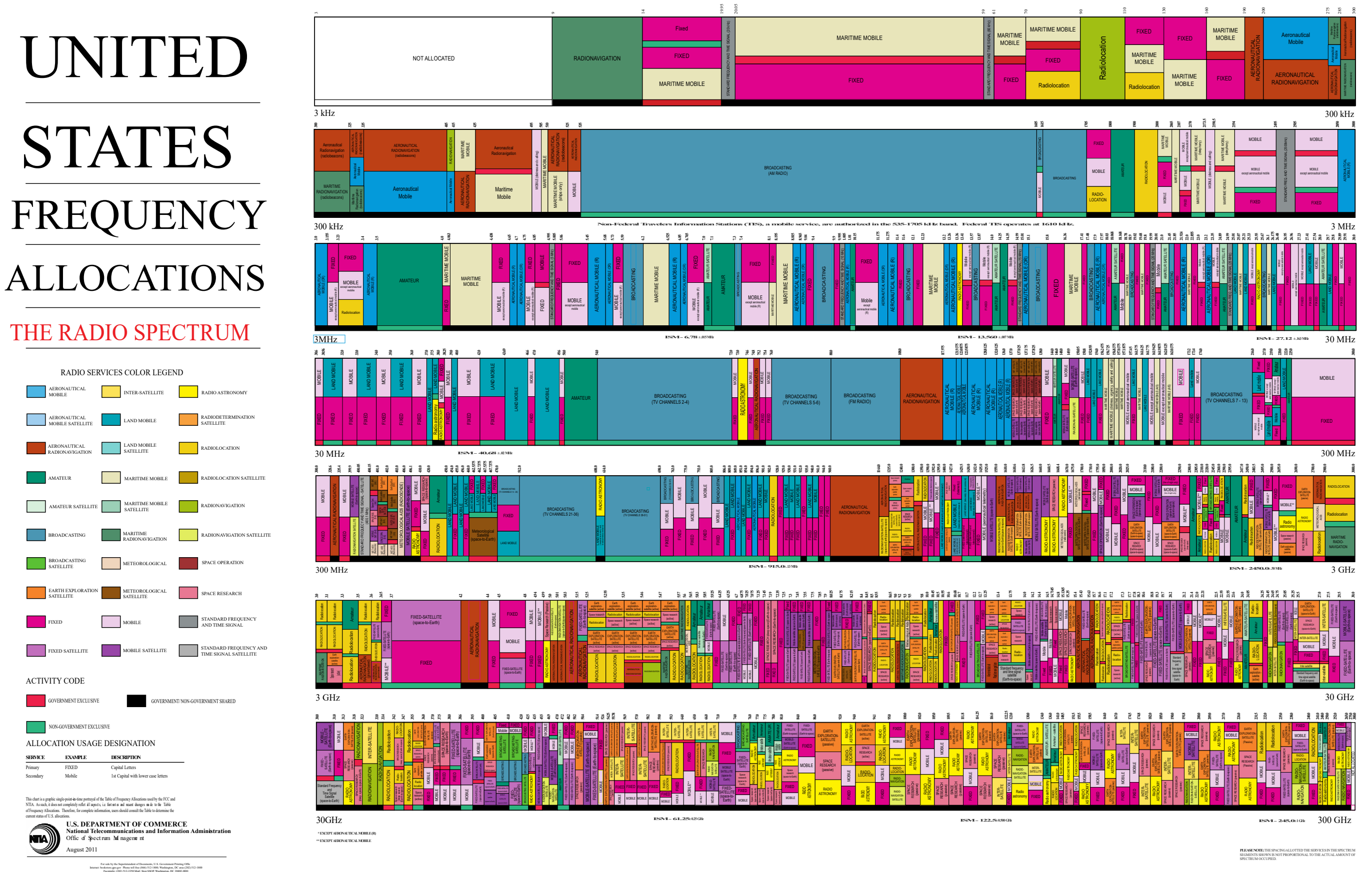

… and available RF spectrum is increasingly scarce.

Motivation (2): Spectrum Sensing

Sampling theorems inspire new ways of sensing the spectrum:

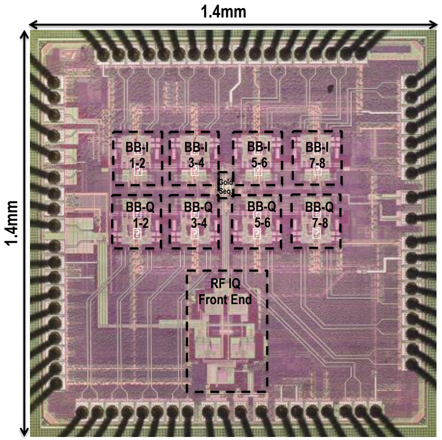

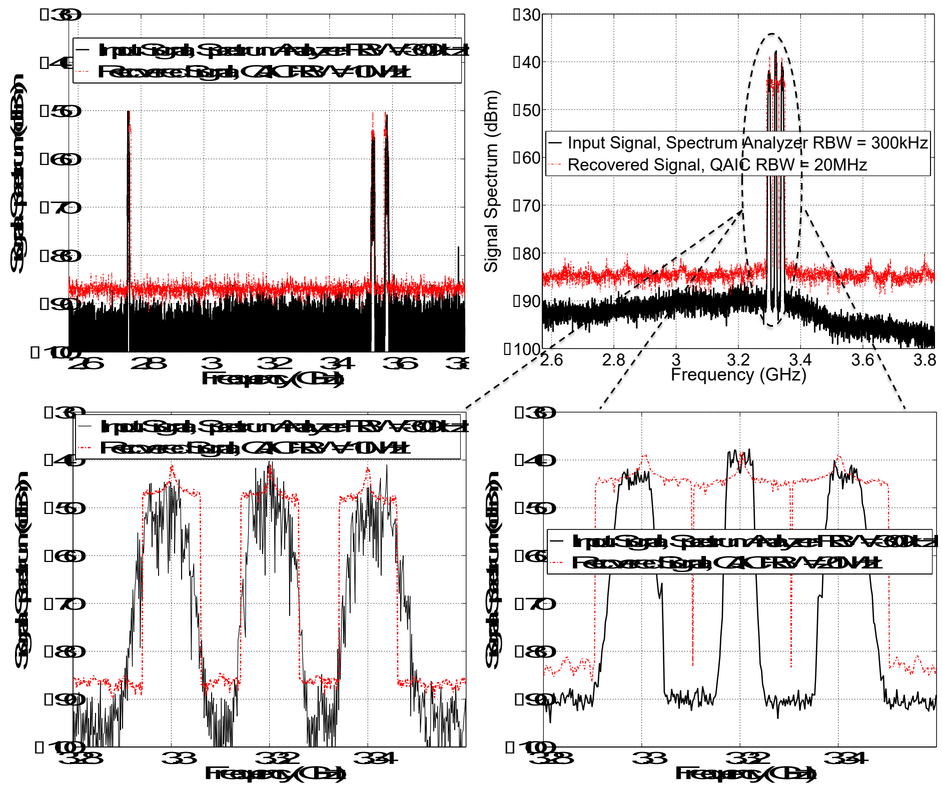

Joint work led by T. Yazicigil, T. Haque, P. Kinget

Motivation (3): Dynamic MRI

… and for other completely different signal modalities, e.g., magnetic resonance imaging:

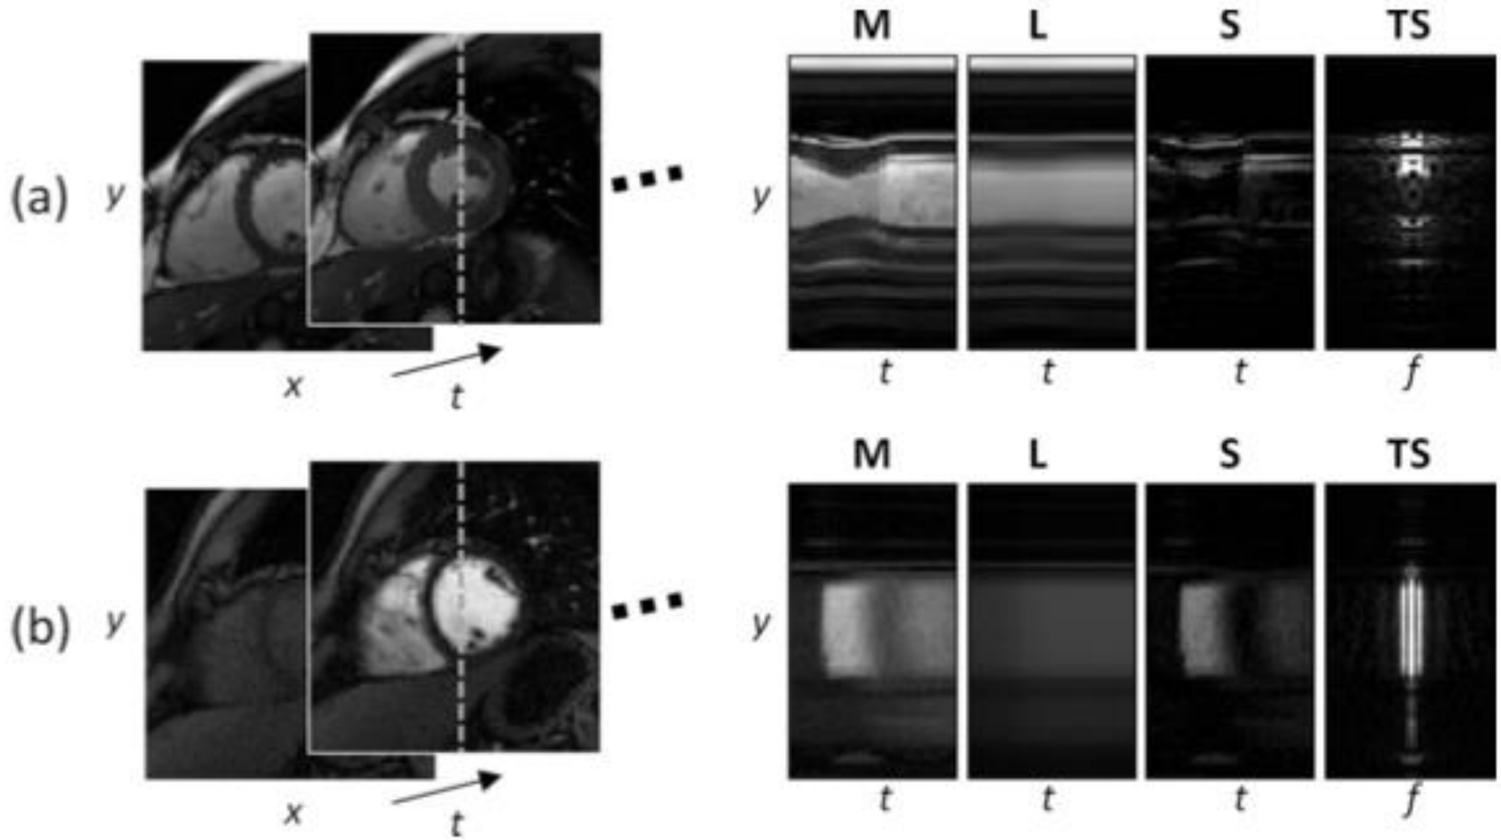

Example from [Otazo+Candes+Sodickson, *Low-Rank and Sparse Matrix Decomposition for Accelerated MRI*, 2014]

Syllabus

Introduction:	Applications, review of complex math, basic sequences
Time Domain:	LTI systems, convolution, impulse response
Frequency Domain:	Discrete Time Fourier Transform (DTFT), DFT, FFT, STFT
Sampling:	Shannon-Nyquist, Analog-Digital conversion

Midterm exam

Z-transform:	Definition, poles, zeros, application in system analysis
Filters:	Magnitude, phase response, linear/generalized linear phase
Filter design:	FIR design tools, IIR design tools

Special topics:	Adaptive filtering
	Optimization & learning based signal processing
	Applications

Final exam

Course Organization

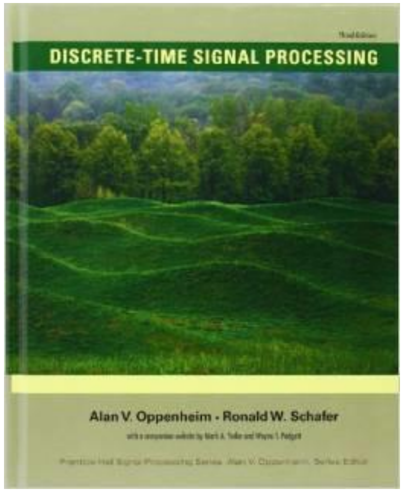

Discrete-Time Signal Processing
Oppenheim + Schafer, 3rd Ed.
[Yes, digital/pdf version is fine.]

Instructor: John Wright
johnwright@ee.columbia.edu
Office hours: 11-1 PM Thursday (408 Mudd)

TA: TBD

Website: Course materials (notes, videos, assignments) on CourseWorks

Grades

Homework	15%
Midterm	30%
Final	30%
Course Project	25%

Exams

Closed-book. *No phones, calculators, or other computing devices.*

Questions: mostly analytical and conceptual. We strive to test understanding; heroic feats of mathematical proof or hand calculation will probably not be necessary.

Students often find the exams challenging. Please prepare appropriately! We will have review sessions before both exams.

Midterm: Tentative date **October 27**. One sheet of handwritten notes (front+back) allowed.

Final: Tentative date **December 15, 1:10 PM-4 PM**. Two sheets of handwritten notes (front+back) allowed. The final is comprehensive in scope, but contains more material from the second half of the semester.

Homework

Analytical problems (Oppenheim-Schafer) + Matlab exercises.

Homework will start heavy and then get much lighter at the end of the semester.

You can get **Matlab** for free through SEAS:
https://portal.seas.columbia.edu/matlab/

If you don't know Matlab, a good resource for getting started is:
http://www.ee.columbia.edu/ee-matlab

Feel free to discuss with peers, but you must write up your solutions independently. Please don't cheat yourself out of learning the material!

Late homework will not be accepted. We *will* drop your lowest homework score, as a buffer against illness, travel, hungry pets, meteor strike, ect.

Course Project

Process some kind of real signal.

Teams:

3-4 students. You can self-organize around common interests via the course forum on CourseWorks.
Students who cannot find teammates will be randomly grouped.

1-2 student teams are also acceptable, **with instructor approval.**

Programming language:

Any programming language/platform is ok. We recommend Matlab or Python.

Deliverables:

Topic + team in writing, by **Oct 15**. Details in the syllabus and on Courseworks.

Code and **Project report** (~4 pages) by **December 23**, 11:59 PM.

Course Project

Process some kind of real signal.

Project ideas (more online):

Audio: speech classification, pitch detection, guitar effects
Images: superresolution ("magic" zoom), categorization, camera shake removal
Videos: transform coding / compression
Depth data: denoising / interpolation
Other types of signals are great:
Bioinformatic data, financial time series, seismic surveys

Good projects are often **fun and relevant** to your everyday life.

Advanced tools are also great …

Machine learning techniques, deep learning
Nonlinear features,
Compressed sensing, sparse approximation,
Source separation techniques

… as long as you're applying them to some type of real data.

Syllabus

Introduction:	Applications, review of complex math, basic sequences
Time Domain:	LTI systems, convolution, impulse response
Frequency Domain:	Discrete Time Fourier Transform (DTFT), DFT, FFT, STFT
Sampling:	Shannon-Nyquist, Analog-Digital conversion

Midterm exam

Z-transform:	Definition, poles, zeros, application in system analysis
Filters:	Magnitude, phase response, linear/generalized linear phase
Filter design:	FIR design tools, IIR design tools

Special topics: Adaptive filtering
Optimization & learning based signal processing
Applications

Final exam