PCMI Lecture # 1

Image Analysis : Introduction and motivation

Image analysis / processing is concerned, to a large degree with the extraction of information from images.

Given: noisy, incomplete measurements of images of objects [Data]

Goal: Deduce something (1) automatically, or (2) that is "hidden" or (3) both from the data about the objects the data came from.

Here are some typical tasks.

Denoising of images: <u>given</u> $d \in L^\infty(S)$ $S \equiv [0,1] \times [0,1]$ where

$d = u + \eta$ $u \in L^\infty(S)$ = true image

η = noise component

<u>Find</u> u.

Example (a) η might be generated by a Gaussian process, u could be a piecewise constant cartoon.

(b) η might again be Gaussian, u might be a smoothly varying ~~[illegible]~~ function.

Image reconstruction I: <u>given</u> $\{d_i\}_{i=1}^N$, $d_i \in \mathbb{R}^m$ $d_i = P_i u$ $u \in L^\infty(S)$

where each P_i is a linear measurement operator, u might be an element of $\mathbb{R}^n$

~~[illegible]~~ (Or u might be $n > m$)

Find u.

stated again: <u>Given</u> $d = Pu$

$d \in \mathbb{R}^{mN}$, $u \in \mathbb{R}^n$, P linear, <u>Find</u> u

Of course the real art is understanding what constraints or prior allow you to effectively find $u \in W$ $W \subset \mathbb{R}^n$ $\dim(W) \ll n$, $\dim(W) \sim mN$.

<u>example:</u> Tomography

Image reconstruction II:

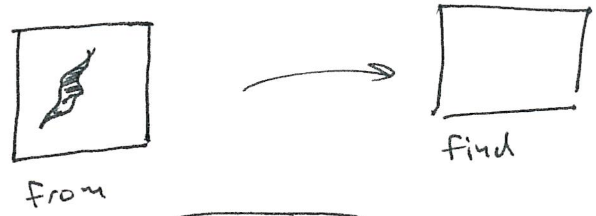

$\boxed{\text{inpainting}}$

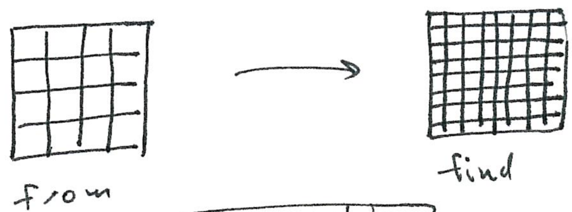

$\boxed{\text{super-resolution}}$

Decomposition I:

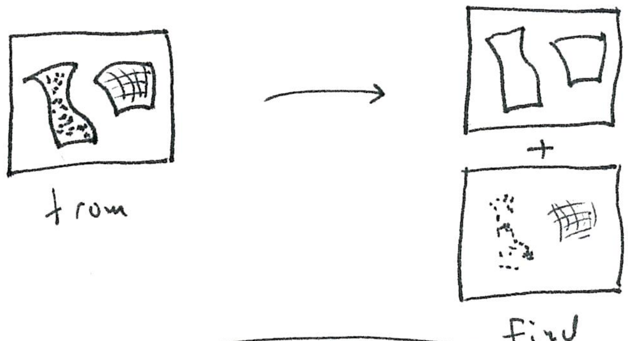

$\boxed{\text{cartoon + texture decomposition}}$

Decomposition II:

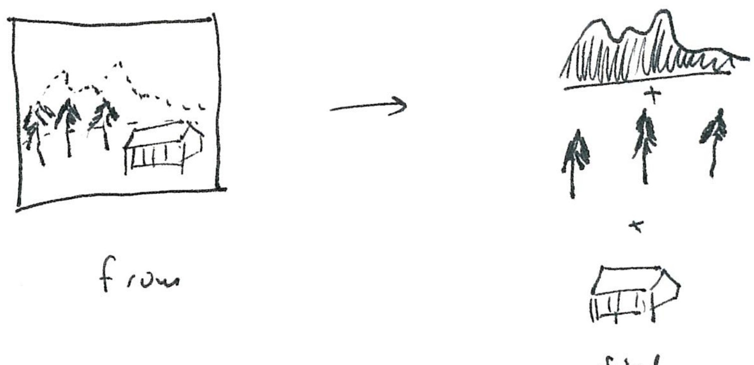

$\boxed{\text{segmentation}}$

Detection: Find faces, cars, trees, crystals, tortoises, etc in the image.

This is a higher level segmentation

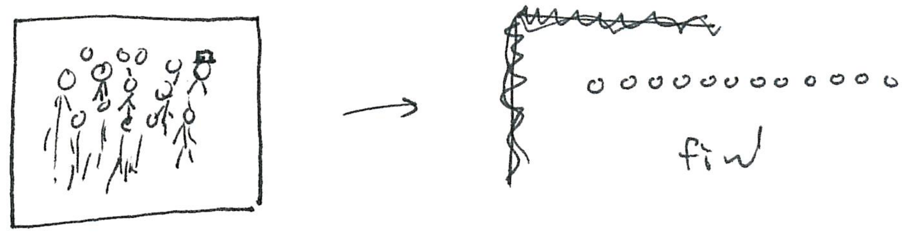

from

from

Recognition: find particular faces, cars, trees, crystals, tortoises,
or identify the faces, cars, trees, crystal, tortoises that you detect.

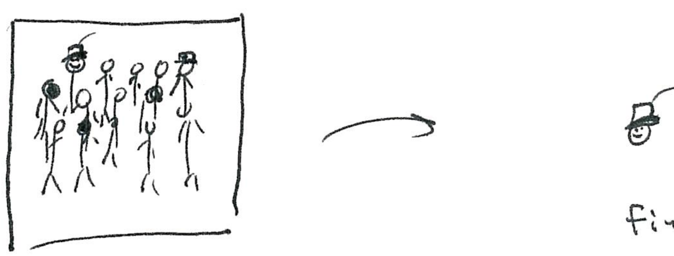

from

?

Jack Jabwnoky

We will not consider topics like

Compression:

Sparse Representation

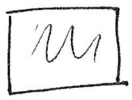 $= f_1 + f_2 + f_3$

$\{f_i\}_{i=1}^{N}$ standard library of basis functions

But of course both are connected to previous problems. For example if $u = f_1 + f_2 + f_3$ and $d = u + \eta$, then we can use compressed sensing methods to extract u from d. We will not study this, but there are others at the summer school who will concentrate on this currently hot topic. E. G. Jared Tanner

Discussion of prior knowledge

Solving any of the problems above requires the use of prior knowledge (one can compress without prior knowledge at least losslessly) ~~compression~~ Lets look at a one dimensional example.

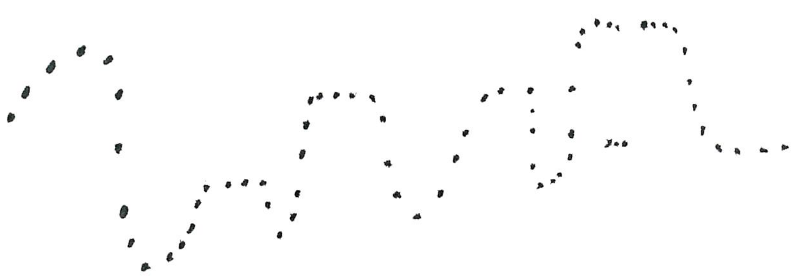

or

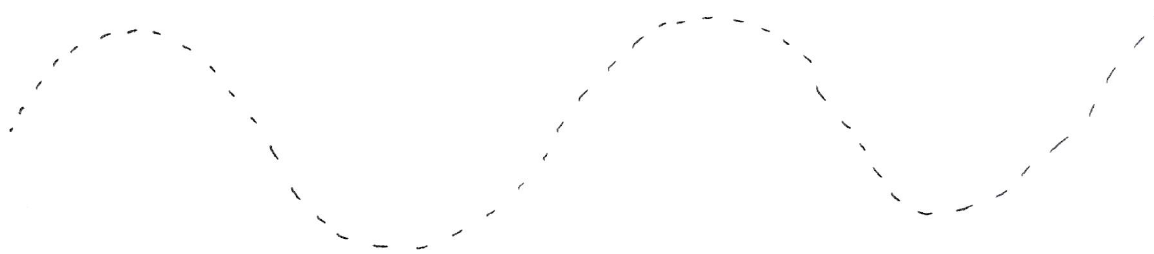

$\{d\} = \{u(i)\}_{i=1}^{n} \in \mathbb{R}^n$

<u>Task</u>: Find signal between samples

$\boxed{\text{Key observation}}$: without more information than the samples all we have is an (ordered) ~~heap~~ list of numbers. We know <u>nothing</u> about the signal between samples

<u>Given</u>: prior assumptions, we can say something.

<u>Spectral assumption</u>: highest frequency in signal is f and sampling frequency is at least $2f$. <u>Then</u> d contains u! In fact we can get u from a simple convolution operation.

<u>Smoothness</u>: polygonal or spline or some ~~[illegible]~~ kernel estimate method gives good approximations. This is precisely what we naturally do when we look at the above data.

<u>Continuity</u>: obviously not enough. There is a sense in which the data only restricts you to a codimension 10 subspace of $C([0,L])$.

$C_\alpha([0,L]) \equiv$ continuous functions on $[0,L]$ which are zero at the sampling points or grid α.

$P^\alpha([0,L]) \equiv$ piecewise linear functions on $[0,L]$ which are permitted corners at the sampling points α. $P^\alpha([0,L])$ "=" $\mathbb{R}^{|\alpha|}$

$$C([0,L]) = C_\alpha([0,L]) \oplus P^\alpha([0,L])$$

But this is a longer discussion we will continue in lecture 4.

The main point here is that naked data is just naked data without additional knowledge — prior knowledge — we can't do anything (except lossless compression).

In lectures 2 and 3, we will talk about metrics, ubiquitous and critical for any applied task. For example, even a move from lossless compression to lossy compression requires us to decide what is a small error in approximation.

—

The rest of lecture will look at examples of image and shape based problems.

—

Shock tube Data

In this experiment, a shock is propagated down a rectangular (cross section) tube

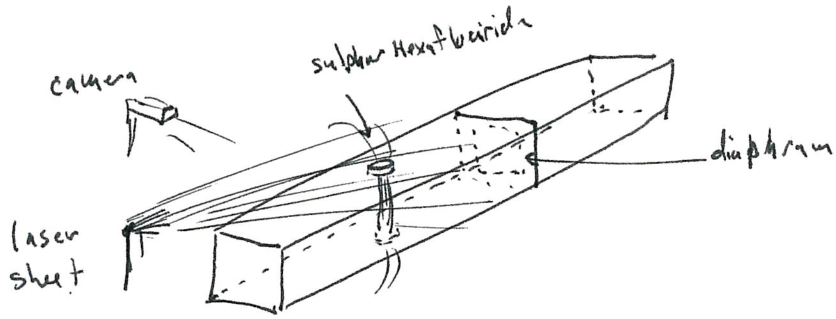

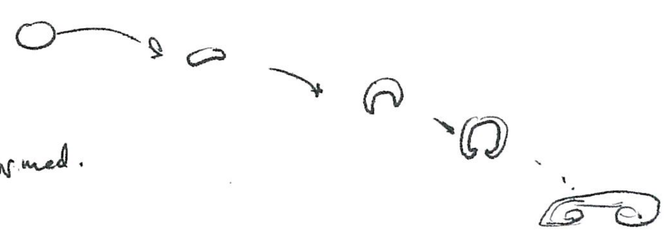

simulations are also performed.

Comparisons between the experiment and simulations are made in order to ~~measure~~ generate a measure of quality of the sim. to validate the simulation. The first question that must be determined is where in the simulated time stream the real data fit show pictures of real and simulated cross sections of the SF_6 column.

X-ray laser ablation experiment.

In This (complicated) experiment lasers are used to generate intense x-rays which drive shockwaves into materials (aluminum, polystyrene, gold) which then jet into each other. The experiments are observed using x-rays and are also simulated. The results are compared, again for validation of complex "hydro-codes".

show pictures of experimental setup

show pictures of experimental / simulated jets.

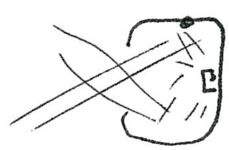

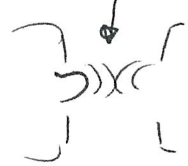

CMODI face recognition

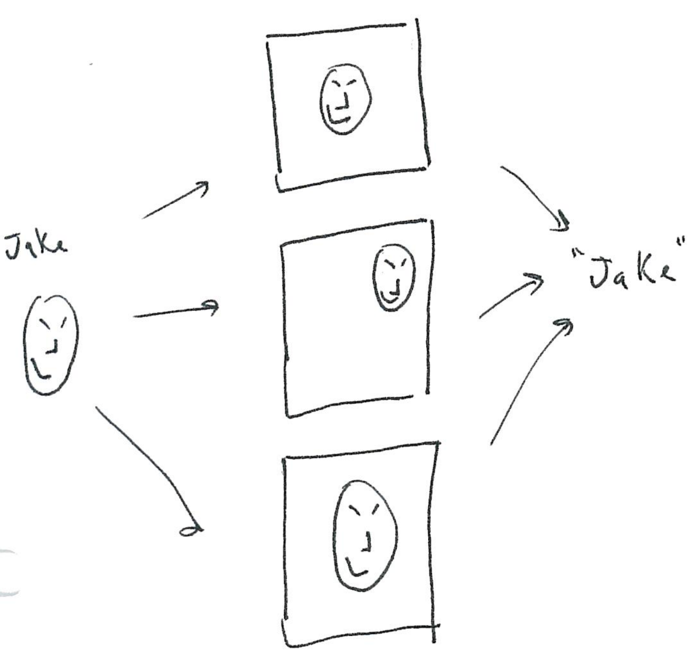

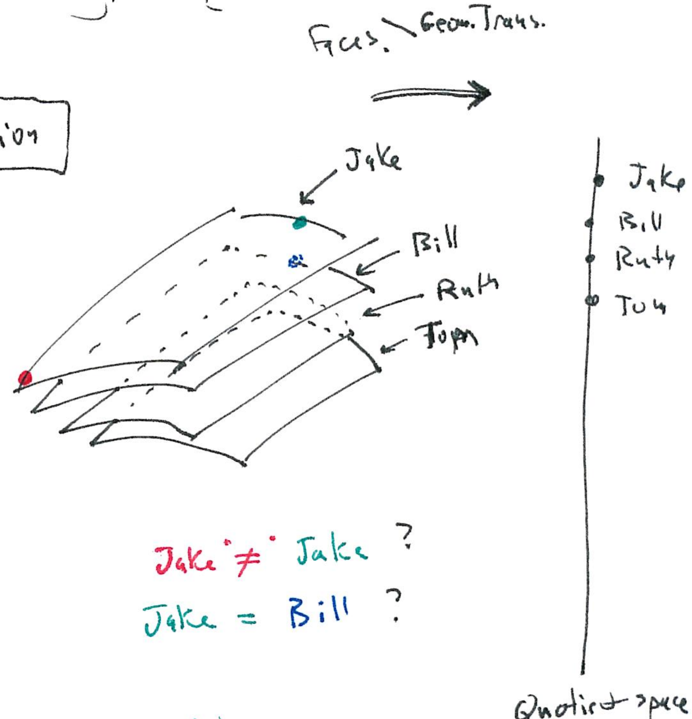

Jake "≠" Jake ?

Jake = Bill ?

- given data
- given truth

show CMODI face regognition picture

Shape recognition task

which are spoons, which are forks.

This is a trivial task for humans and less trivial for computers... though this one, in particular is not too difficult.

We want to separate shape from metric issues here. Don't care about scale, translation, rotation. 3D rotation makes this much more complicated.

Material Science Problems

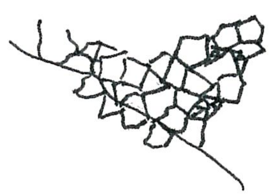

automatically compute statistics on crystalline domains, finding anomolies, cracks, unusual shapes, etc.

...