

# CUR Matrix Decompositions for Improved Data Analysis

---

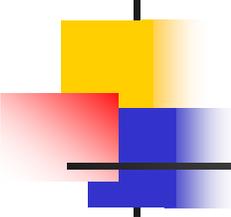
**Michael W. Mahoney**

Yahoo Research

*<http://www.cs.yale.edu/homes/mmahoney>*

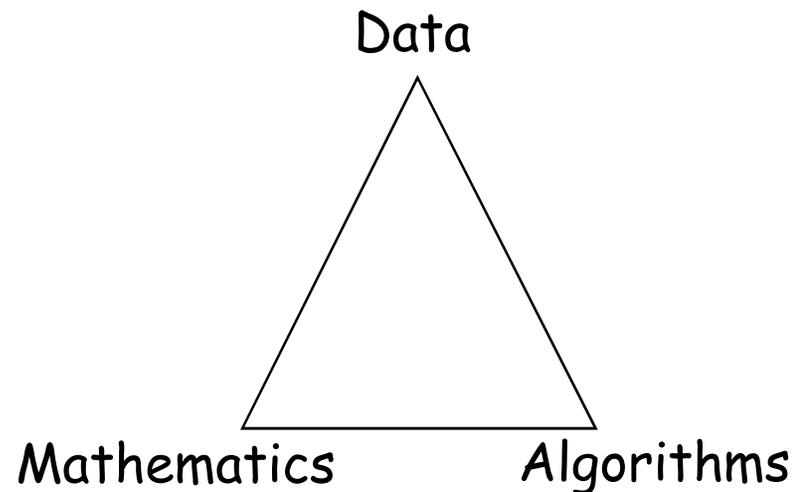
(Joint work with P. Drineas, R. Kannan, S. Muthukrishnan,  
and P. Paschou, K. Kidd, M. Maggioni)

Workshop on Algorithms for Modern Massive Datasets, June 2006



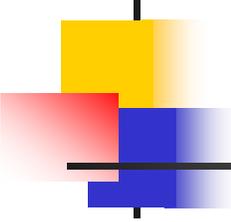
# Modeling data as matrices

---



Matrices often arise with data:

- $n$  objects ("documents," genomes, images, web pages),
- each with  $m$  features,
- may be represented by an  $m \times n$  matrix  $A$ .



# SVD and low-rank approximations

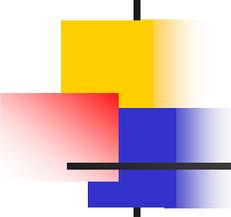
---

**Basic SVD Theorem:** Let  $A$  be an  $m \times n$  matrix with rank  $\rho$ .

- Can express *any* matrix  $A$  as  $A = U \Sigma V^T$ .
- Truncate SVD of  $A$ :  $A_k = U_k \Sigma_k V_k^T$ , get "best" rank- $k$  approximation.

## Properties of truncated SVD:

- Used in data analysis via Principal Components Analysis (PCA) .
- Gives a **very particular structure** (think: rotate-rescale-rotate).
- **Problematic** w.r.t. sparsity, nonnegativity, interpretability, etc.



# Problems with SVD/Eigen-Analysis

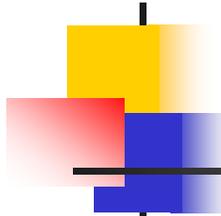
---

**Problems arise:** since **structure in the data** is *not* respected by mathematical **operations on the data**:

- **Sparsity** - is destroyed by orthogonalization.
- **Non-negativity** - is a convex and not linear algebraic notion.
- **Interpretability** - what does a linear combination of 6000 genes "mean."
- **Reification** - maximum variance directions are just that.

**Question:** Do there exist "**better**" **low-rank matrix approximations**.

- "better" **structural properties** for certain applications.
- "better" at **respecting relevant structure**.
- "better" for **interpretability** and **informing intuition**.



# CX and CUR matrix decompositions

---

**Recall:** Matrices are about their **rows** and **columns**.

**Recall:** **Low-rank** matrices have **redundancy** in their rows and columns.

**Def:** A **CX matrix decomposition** is a low-rank approximation *explicitly expressed* in terms of a small number of **columns** of the original matrix  $A$  (e.g.,  $P_C A = C C^+ A$ ).

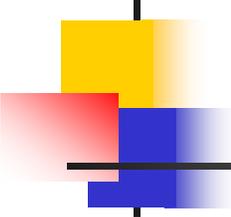
**Def:** A **CUR matrix decomposition** is a low-rank approximation *explicitly expressed* in terms of a small number of **columns and rows** of the original matrix  $A$ .

$$\begin{pmatrix} A \\ m \times n \end{pmatrix} \approx \begin{pmatrix} C \\ O(1) \text{ columns} \end{pmatrix} \cdot \begin{pmatrix} U \\ \text{Carefully chosen } U \end{pmatrix} \cdot \begin{pmatrix} R \\ O(1) \text{ rows} \end{pmatrix}$$



# Previous CUR-type decompositions

Goreinov, Tyrtshnikov, & Zamarashkin (LAA '97, ...)	<p><b>C</b>: columns that span max volume</p> <p><b>U</b>: <math>W^+</math></p> <p><b>R</b>: rows that span max volume</p>	<p>Existential result</p> <p>Error bounds depend on <math>\ W^+\ _2</math></p> <p>Spectral norm bounds!</p>
Berry, Stewart, & Pulatova (Num. Math. '99, TR '04, ...)	<p><b>C</b>: variant of the QR algorithm</p> <p><b>R</b>: variant of the QR algorithm</p> <p><b>U</b>: minimizes <math>\ A-CUR\ _F</math></p>	<p>No a priori bounds</p> <p><math>A</math> must be known to construct <math>U</math>.</p> <p>Solid experimental performance</p>
Williams & Seeger (NIPS '01, ...)	<p><b>C</b>: uniformly at random</p> <p><b>U</b>: <math>W^+</math></p> <p><b>R</b>: uniformly at random</p>	<p>Experimental evaluation</p> <p><math>A</math> is assumed PSD</p> <p>Connections to Nystrom method</p>
Drineas, Kannan & Mahoney (TR '04, SICOMP '06)	<p><b>C</b>: w.r.t. column lengths</p> <p><b>U</b>: in linear/constant time</p> <p><b>R</b>: w.r.t. row lengths</p>	<p>"Sketching" massive matrices</p> <p>Provable, a priori, bounds</p> <p>Explicit dependency on <math>A - A_k</math></p>
Drineas, Mahoney, & Muthukrishnan (TR '06)	<p><b>C</b>: depends on singular vectors of <math>A</math>.</p> <p><b>U</b>: (almost) <math>W^+</math></p> <p><b>R</b>: depends on singular vectors of <math>C</math></p>	<p><math>(1+\epsilon)</math> approximation to <math>A - A_k</math></p> <p>Computable in low polynomial time</p> <p>(Suffices to compute <math>SVD(A_k)</math>)</p>



# Three CUR Data Applications

---

## **Human Genetics:** DNA SNP Data

- Biological Goal: Evaluate intra- and inter-population tag-SNP transferability.

## **Medical Imaging:** Hyperspectral Image Data

- Medical Goal: Compress the data, without sacrificing classification quality.

## **Recommendation Systems:** Customer Preference Data

- Business Goal: Reconstruct the data, to make high-quality recommendations.

# CUR Data Application: Human Genetics

(Joint work with P. Paschou and K. Kidd's lab at Yale University)

Recall, "the" human genome:

- 30,000 - 40,000 genes
- 3 billion base pairs
- The functionality of 97% of the genome is unknown.
- BUT: individual differences (polymorphic variation) at  $\approx 1$  b.p. per thousand.

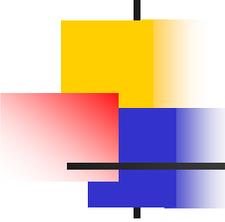
SNPs (Single Nucleotide Polymorphisms):

- The most common type of genetic polymorphic variation.
- They are known locations at the human genome where two (out of A, C, G, T) alternate nucleotide bases (alleles) are observed.

SNPs

individuals

... AG CT GT GG CT CC CC CC CC AG AG AG AG AG AA CT AA GG GG CC GG AG CG AC CC AA CC AA GG TT AG CT CG CG CG AT CT CT AG CT AG GG GT GA AG ...  
... GG TT TT GG TT CC CC CC CC GG AA AG AG AG AA CT AA GG GG CC GG AA GG AA CC AA CC AA GG TT AA TT GG GG GG TT TT CC GG TT GG GG TT GG AA ...  
... GG TT TT GG TT CC CC CC CC GG AA AG AG AA AG CT AA GG GG CC AG AG CG AC CC AA CC AA GG TT AG CT CG CG CG AT CT CT AG CT AG GG GT GA AG ...  
... GG TT TT GG TT CC CC CC CC GG AA AG AG AG AA CC GG AA CC CC AG GG CC AC CC AA CG AA GG TT AG CT CG CG CG AT CT CT AG CT AG GT GT GA AG ...  
... GG TT TT GG TT CC CC CC CC GG AA GG GG GG AA CT AA GG GG CT GG AA CC AC CG AA CC AA GG TT GG CC CG CG CG AT CT CT AG CT AG GG TT GG AA ...  
... GG TT TT GG TT CC CC CG CC AG AG AG AG AG AA CT AA GG GG CT GG AG CC CC CG AA CC AA GT TT AG CT CG CG CG AT CT CT AG CT AG GG TT GG AA ...  
... GG TT TT GG TT CC CC CC CC GG AA AG AG AG AA TT AA GG GG CC AG AG CG AA CC AA CG AA GG TT AA TT GG GG GG TT TT CC GG TT GG GT TT GG AA ...



# SNP Biology

---

SNPs carry **redundant information**:

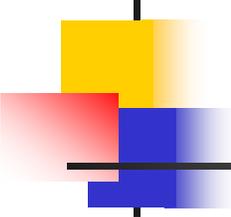
- Human genome is organized into **block-like structure**.
- Strong, but nontrivial, intra-block **correlations**.
- Can focus only on "tagging SNPs," or **tSNPs**.

Different patterns of SNP frequencies/correlations in different populations (e.g., European, Asian, African, etc.):

- Can track **population histories** and **disease genes**.
- Effective markers for **genomic research**.

**International HapMap Project**:

- Create a haplotype map of human genetic variability.
- Map all 10,000,000 SNPs for 270 individuals from 4 different populations.



# SNP Pharmacology

---

## Disease association studies:

- Locate causative genes for **common complex disorders** (e.g., diabetes, heart disease).
- **Identify association** between affection status and known SNPs.
- **Don't need:** knowledge of function of the genes or etiology of the disorder.
- **Investigate candidate genes** in physical proximity with associated SNPs.

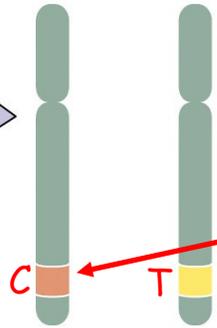
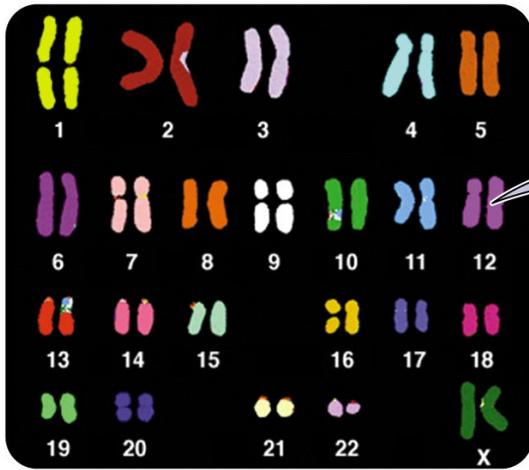
## Develop the "next generation" of drugs:

- "*population-specific,*" eventually "*genome-specific,*" not just "*disease-specific*".

## Funding:

- HapMap project (~\$100,000,000 from NIH, etc.).
- Funding also from pharmaceutical companies, NSF, the DOJ\*, etc.

\*Is it possible to identify the ethnicity of a suspect from his DNA?



Two copies of a chromosome (father, mother)

Focus at a specific locus and assay the observed alleles.

SNP: exactly two alternate alleles appear.

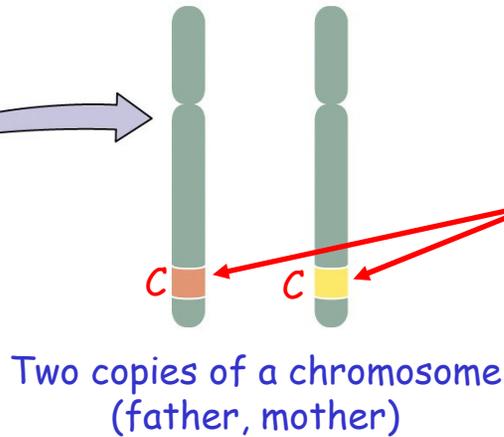
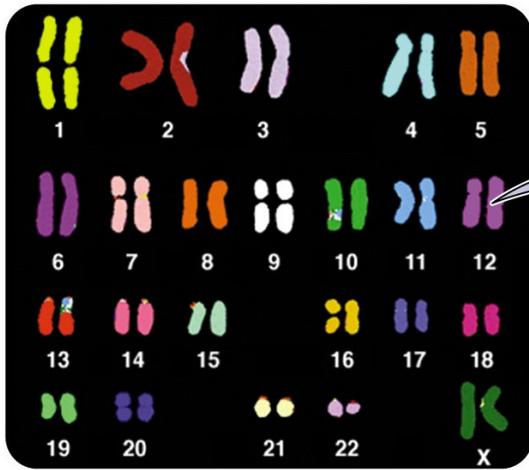
An individual could be:

- Heterozygotic (in our study, CT = TC)

SNPs

individuals

...	AG	CT	GT	GG	CT	CC	CC	CC	AG	AG	AG	AG	AG	AA	CT	AA	GG	GG	CC	GG	AG	CG	AC	CC	AA	CC	AA	GG	TT	AG	CT	CG	CG	CG	AT	CT	CT	AG	CT	AG	GG	GT	GA	AG	...
...	GG	TT	TT	GG	TT	CC	CC	CC	GG	AA	AG	AG	AG	AA	CT	AA	GG	GG	CC	GG	AA	GG	AA	CC	AA	CC	AA	GG	TT	AA	TT	GG	GG	GG	TT	TT	CC	GG	TT	GG	GG	TT	GG	AA	...
...	GG	TT	TT	GG	TT	CC	CC	CC	GG	AA	AG	AG	AA	AG	CT	AA	GG	GG	CC	AG	AG	CG	AC	CC	AA	CC	AA	GG	TT	AG	CT	CG	CG	CG	AT	CT	CT	AG	CT	AG	GG	GT	GA	AG	...
...	GG	TT	TT	GG	TT	CC	CC	CC	GG	AA	AG	AG	AG	AA	CC	GG	AA	CC	CC	AG	GG	CC	AC	CC	AA	CG	AA	GG	TT	AG	CT	CG	CG	CG	AT	CT	CT	AG	CT	AG	GT	GT	GA	AG	...
...	GG	TT	TT	GG	TT	CC	CC	CC	GG	AA	GG	GG	GG	AA	CT	AA	GG	GG	CT	GG	AA	CC	AC	CG	AA	CC	AA	GG	TT	GG	CC	CG	CG	CG	AT	CT	CT	AG	CT	AG	GG	TT	GG	AA	...
...	GG	TT	TT	GG	TT	CC	CC	CG	CC	AG	AG	AG	AG	AA	CT	AA	GG	GG	CT	GG	AG	CC	CC	CG	AA	CC	AA	GT	TT	AG	CT	CG	CG	CG	AT	CT	CT	AG	CT	AG	GG	TT	GG	AA	...
...	GG	TT	TT	GG	TT	CC	CC	CC	GG	AA	AG	AG	AG	AA	TT	AA	GG	GG	CC	AG	AG	CG	AA	CC	AA	CG	AA	GG	TT	AA	TT	GG	GG	GG	TT	TT	CC	GG	TT	GG	GT	TT	GG	AA	...



Focus at a specific locus and assay the observed alleles.

SNP: exactly two alternate alleles appear.

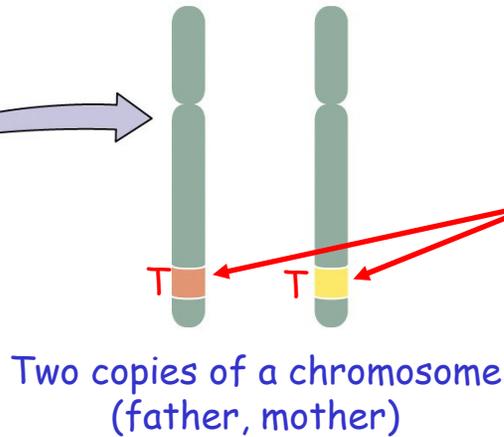
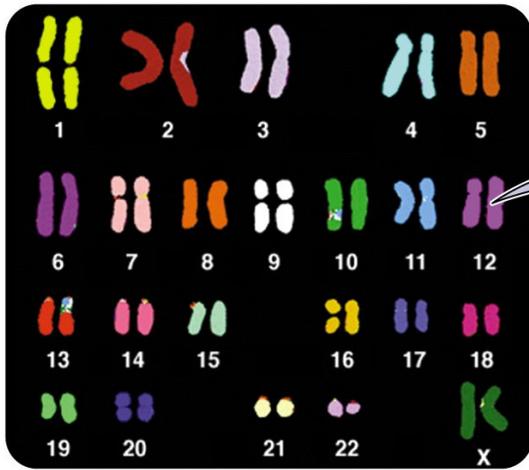
An individual could be:

- Heterozygotic (in our study, CT = TC)
- Homozygotic at the first allele, e.g., C

SNPs

individuals

... AG CT GT GG CT CC CC CC CC AG AG AG AG AG AA CT AA GG GG CC GG AG CG AC CC AA CC AA GG TT AG CT CG CG CG AT CT CT AG CT AG GG GT GA AG ...  
 ... GG TT TT GG TT CC CC CC CC GG AA AG AG AG AA CT AA GG GG CC GG AA GG AA CC AA CC AA GG TT AA TT GG GG GG TT TT CC GG TT GG GG TT GG AA ...  
 ... GG TT TT GG TT CC CC CC CC GG AA AG AG AA AG CT AA GG GG CC AG AG CG AC CC AA CC AA GG TT AG CT CG CG CG AT CT CT AG CT AG GG GT GA AG ...  
 ... GG TT TT GG TT CC CC CC CC GG AA AG AG AG AA CC GG AA CC CC AG GG CC AC CC AA CG AA GG TT AG CT CG CG CG AT CT CT AG CT AG GT GT GA AG ...  
 ... GG TT TT GG TT CC CC CC CC GG AA GG GG GG AA CT AA GG GG CT GG AA CC AC CG AA CC AA GG TT GG **CC** CG CG CG AT CT CT AG CT AG GG TT GG AA ...  
 ... GG TT TT GG TT CC CC CG CC AG AG AG AG AG AA CT AA GG GG CT GG AG CC CC CG AA CC AA GT TT AG CT CG CG CG AT CT CT AG CT AG GG TT GG AA ...  
 ... GG TT TT GG TT CC CC CC CC GG AA AG AG AG AA TT AA GG GG CC AG AG CG AA CC AA CG AA GG TT AA TT GG GG GG TT TT CC GG TT GG GT TT GG AA ...



Focus at a specific locus and assay the observed alleles.

SNP: exactly two alternate alleles appear.

An individual could be:

- Heterozygotic (in our study, CT = TC)
- Homozygotic at the first allele, e.g., C
- Homozygotic at the second allele, e.g., T

### SNPs

individuals

... AG CT GT GG CT CC CC CC CC AG AG AG AG AG AA CT AA GG GG CC GG AG CG AC CC AA CC AA GG TT AG CT CG CG CG AT CT CT AG CT AG GG GT GA AG ...  
 ... GG TT TT GG TT CC CC CC CC GG AA AG AG AG AA CT AA GG GG CC GG AA GG AA CC AA CC AA GG TT AA TT GG GG GG TT TT CC GG TT GG GG TT GG AA ...  
 ... GG TT TT GG TT CC CC CC CC GG AA AG AG AA AG CT AA GG GG CC AG AG CG AC CC AA CC AA GG TT AG CT CG CG CG AT CT CT AG CT AG GG GT GA AG ...  
 ... GG TT TT GG TT CC CC CC CC GG AA AG AG AG AA CC GG AA CC CC AG GG CC AC CC AA CG AA GG TT AG CT CG CG CG AT CT CT AG CT AG GT GT GA AG ...  
 ... GG TT TT GG TT CC CC CC CC GG AA GG GG GG AA CT AA GG GG CT GG AA CC AC CG AA CC AA GG TT GG CC CG CG CG AT CT CT AG CT AG GG TT GG AA ...  
 ... GG TT TT GG TT CC CC CG CC AG AG AG AG AG AA CT AA GG GG CT GG AG CC CC CG AA CC AA GT TT AG CT CG CG CG AT CT CT AG CT AG GG TT GG AA ...  
 ... GG TT TT GG TT CC CC CC CC GG AA AG AG AG AA TT AA GG GG CC AG AG CG AA CC AA CG AA GG TT AA **TT** GG GG TT TT CC GG TT GG GT TT GG AA ...

# Encoding the SNP data ...

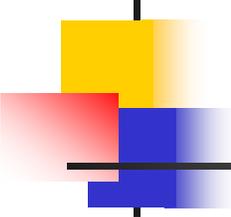
... as an  $m \times n$  matrix  $A$ :

- Exactly two "known" nucleotides (out of A,G,C,T) appear in each column.
- Two alleles might be both equal to the first one (+1), both equal to the second one (-1), or different (0).

		SNPs																									
individuals	{	0 0 0 1 0 -1 1 1 1 0 0 0 0 0 1 0 1 -1 -1 1 -1 0 0 0 1 1 1 1 -1 -1 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0																									
		-1 -1 -1 1 -1 -1 1 1 1 -1 1 0 0 0 1 0 1 -1 -1 1 -1 1 1 1 1 1 -1 -1 1 -1 -1 -1 -1 1 -1 -1 -1 1 -1 -1 1																									
		-1 -1 -1 1 -1 -1 1 1 1 -1 1 0 0 1 0 0 1 -1 -1 1 0 0 0 0 1 1 1 1 -1 -1 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0																									
		-1 -1 -1 1 -1 -1 1 1 1 -1 1 0 0 0 1 1 -1 1 1 1 0 -1 1 0 1 1 0 1 -1 -1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0																									
		-1 -1 -1 1 -1 -1 1 1 1 -1 1 -1 -1 -1 1 0 1 -1 -1 0 -1 1 1 0 0 1 1 1 -1 -1 -1 1 0 0 0 0 0 0 0 0 0 1 -1 -1 1																									
		-1 -1 -1 1 -1 -1 1 0 1 0 0 0 0 0 1 0 1 -1 -1 0 -1 0 1 -1 0 1 1 1 -1 -1 0 0 0 0 0 0 0 0 0 0 0 0 0 1 -1 -1 1																									
		-1 -1 -1 1 -1 -1 1 1 1 -1 1 0 0 0 1 -1 1 -1 -1 1 0 0 0 1 1 1 0 1 -1 -1 1 -1 -1 -1 -1 -1 1 -1 -1 -1 0 -1 -1 1																									
		}																									

Notes:

- Redundancy in rows and columns  $\Leftrightarrow$  Redundancy in SNPs and people.
- SVD has been used (Lin and Altman),
- but, then must *get actual-SNPs/people from eigen-SNPs/people.*



# The SNP data we considered

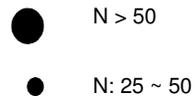
---

## Yale dataset

- Samples from 2000 individuals from 38 different populations.
- Four genomic regions (PAH, SORCS3, HOXB,17q25), a total of  $\approx 250$  SNPs.

## HapMap dataset

- Samples from 270 individuals from 4 different populations (YRI, CEU, CHB, JPT)
- Four genomic regions (PAH, SORCS3, HOXB,17q25), a total of  $\approx 1336$  SNPs.

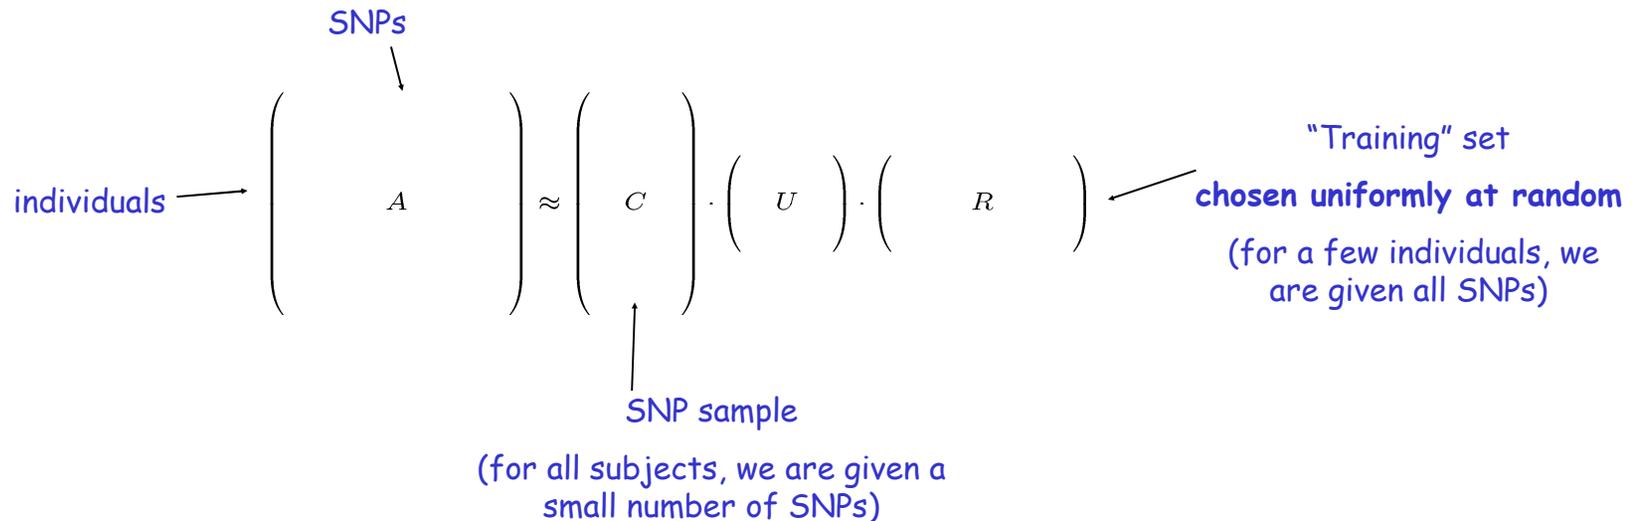


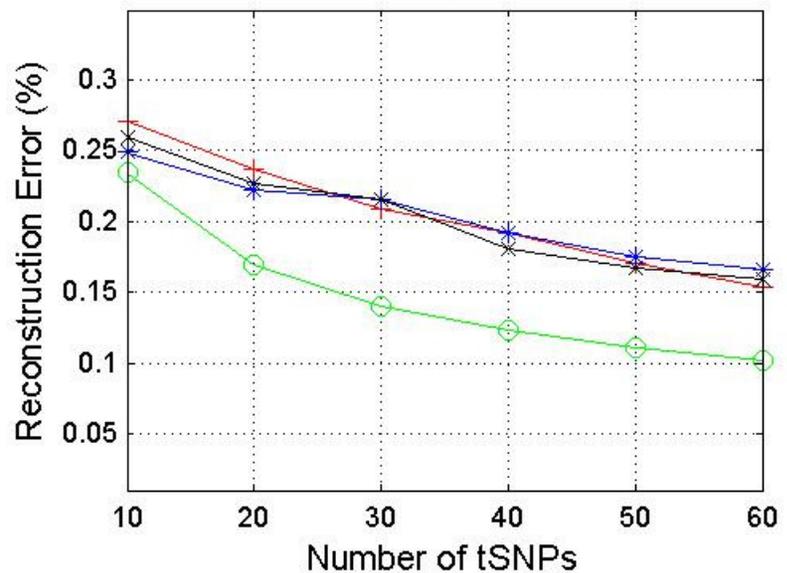
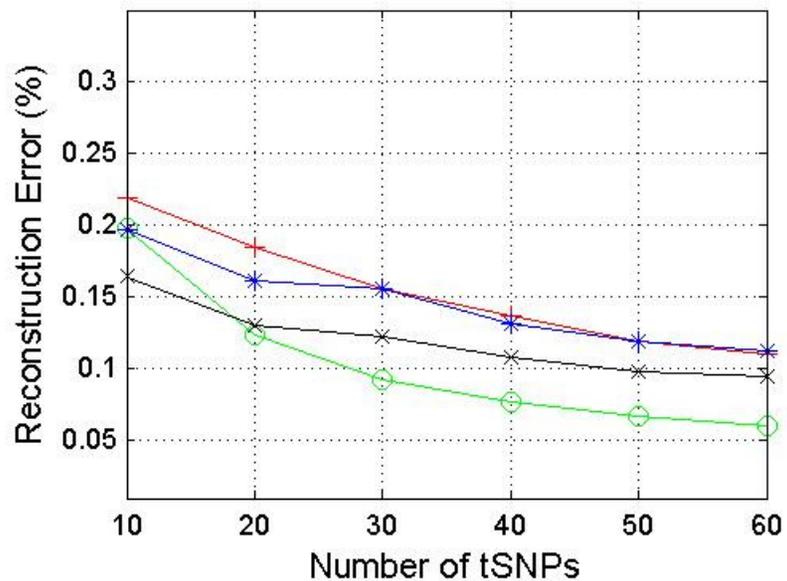
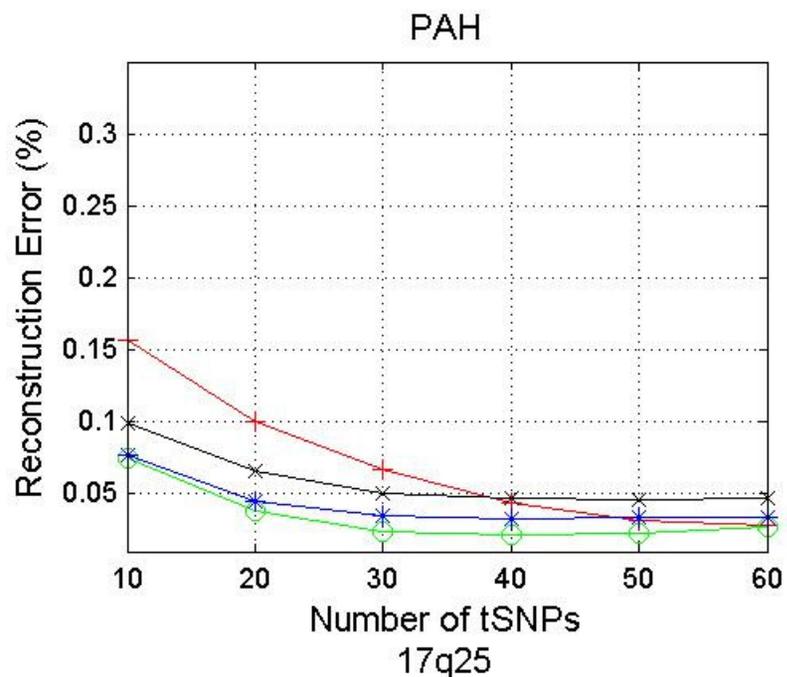
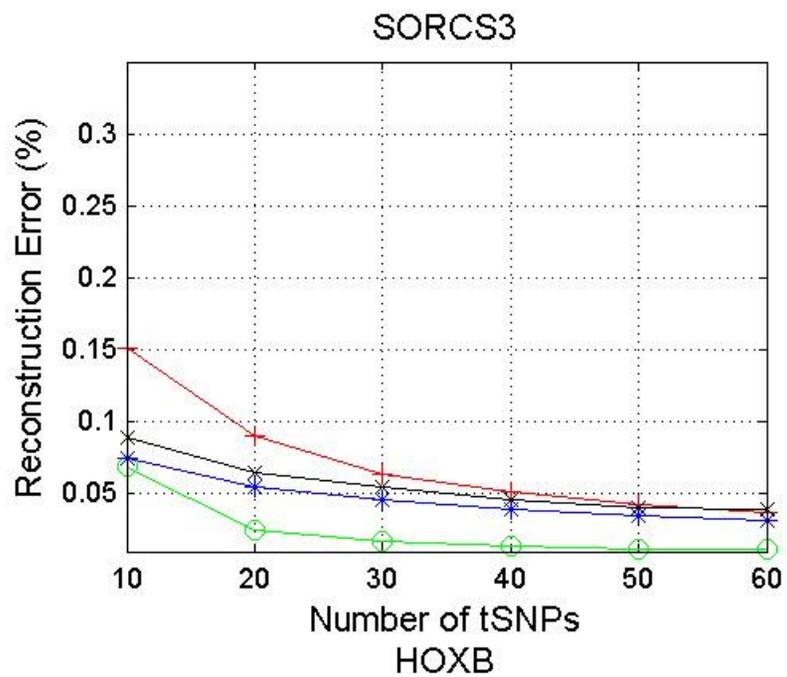
# Predicting SNPs within a population

Split each population: **training and test sets.**

Goal: Given SNP information for **all individuals in the training set** AND for a small number of SNPs for **all individuals** (tagging-SNPs), *predict all unassayed SNPs.*

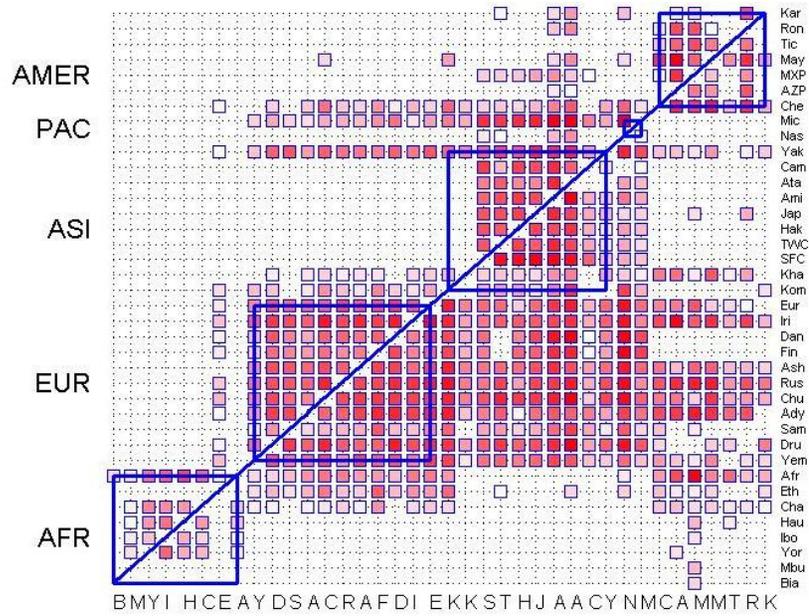
Note: **Tagging-SNPs** are selected using *only* the training set.





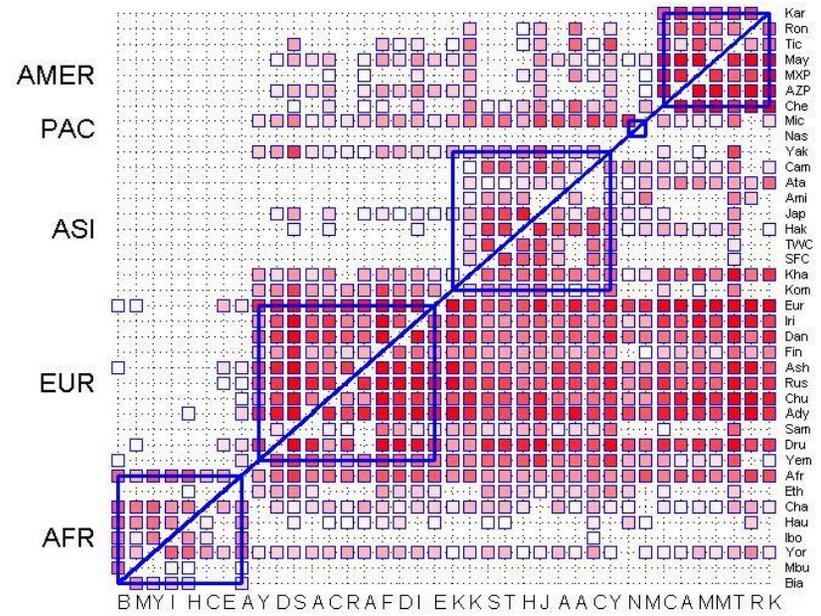


PAH



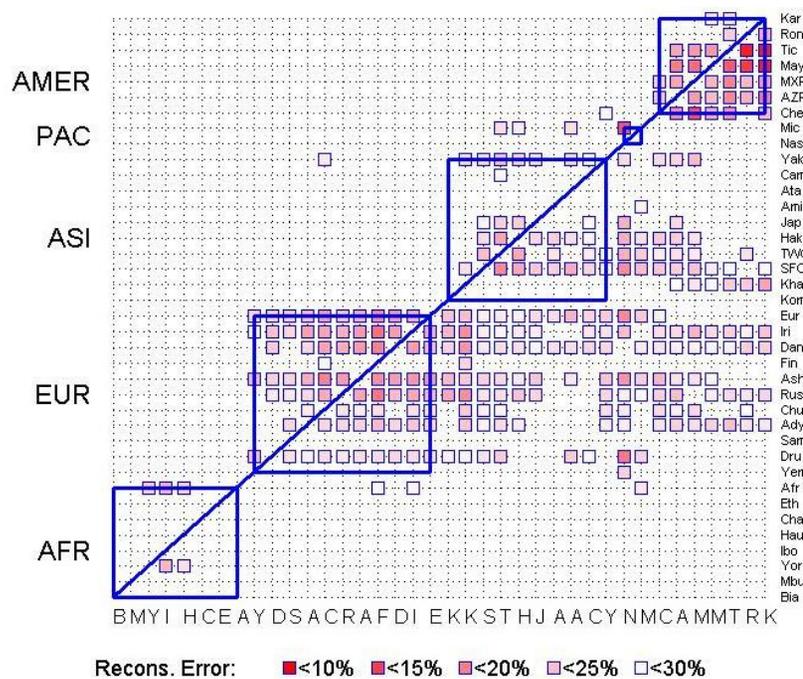
Recons. Error: ■ <10% ■ <15% ■ <20% ■ <25% □ <30%

SORCS3

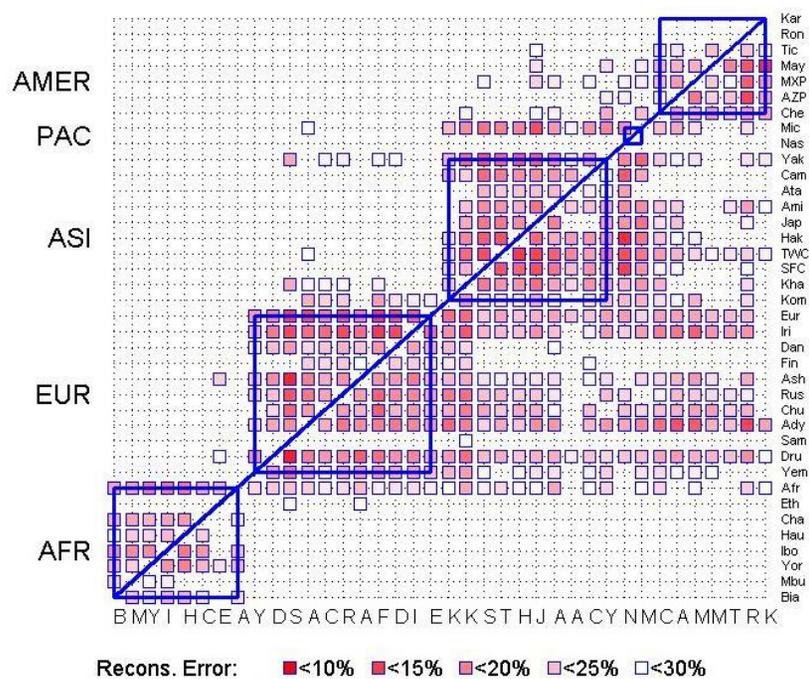


Recons. Error: ■ <10% ■ <15% ■ <20% ■ <25% □ <30%

17q25



HOXB

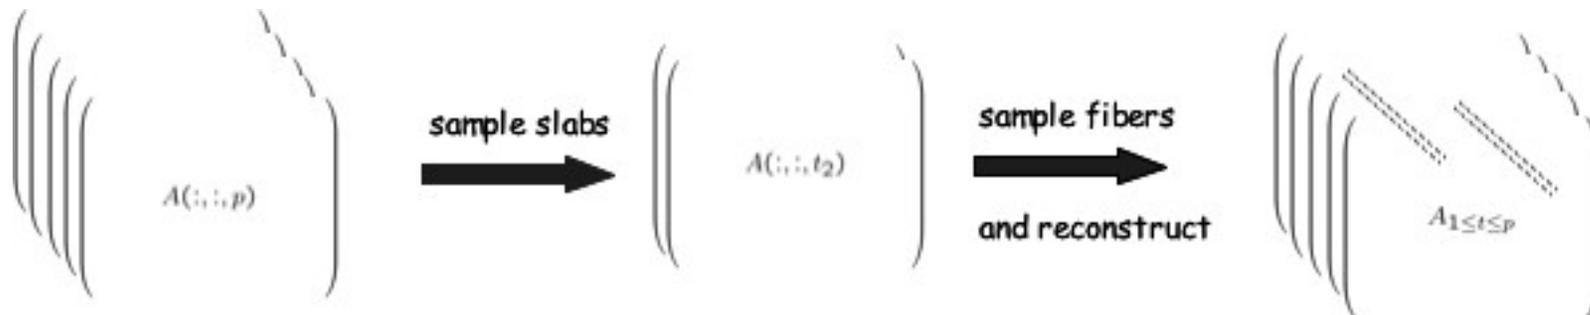
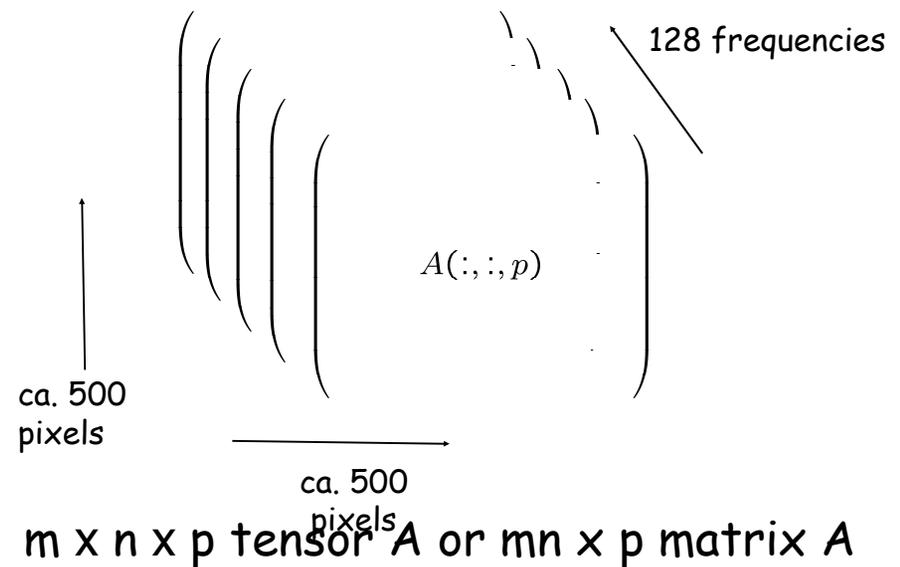


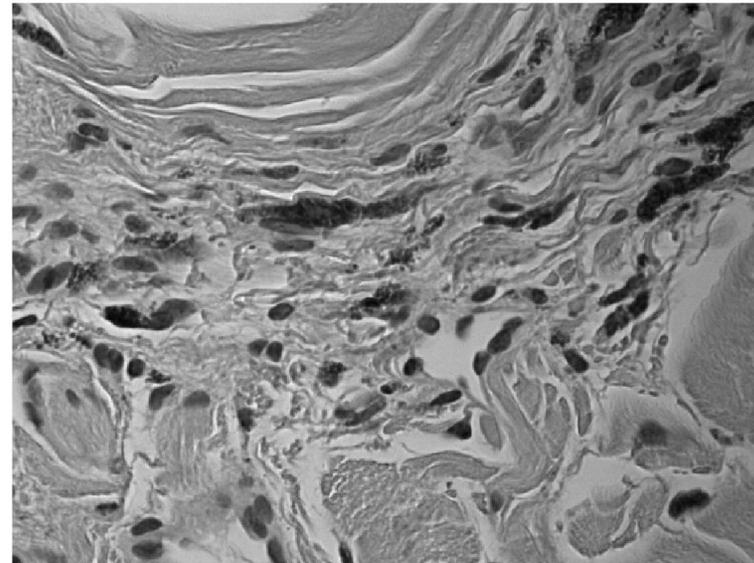
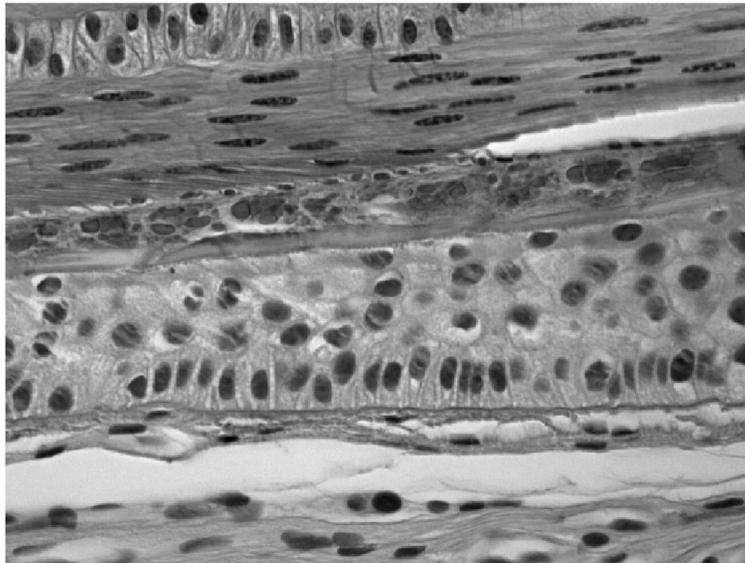
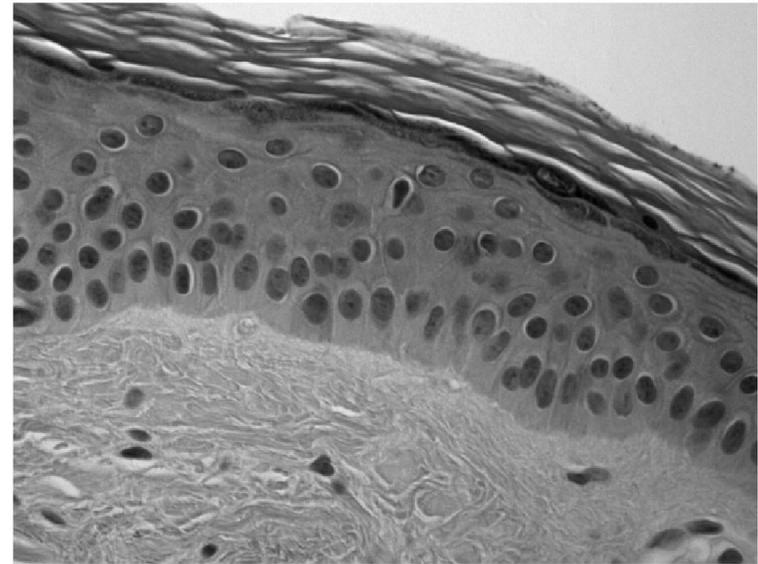
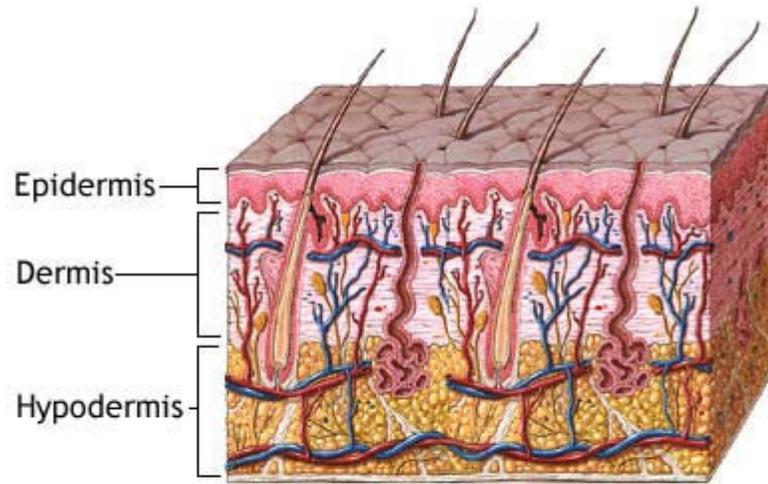
# CUR Data Application: Hyperspectral Image Analysis

(Joint work with M. Maggioni and R. Coifman lab at Yale University)

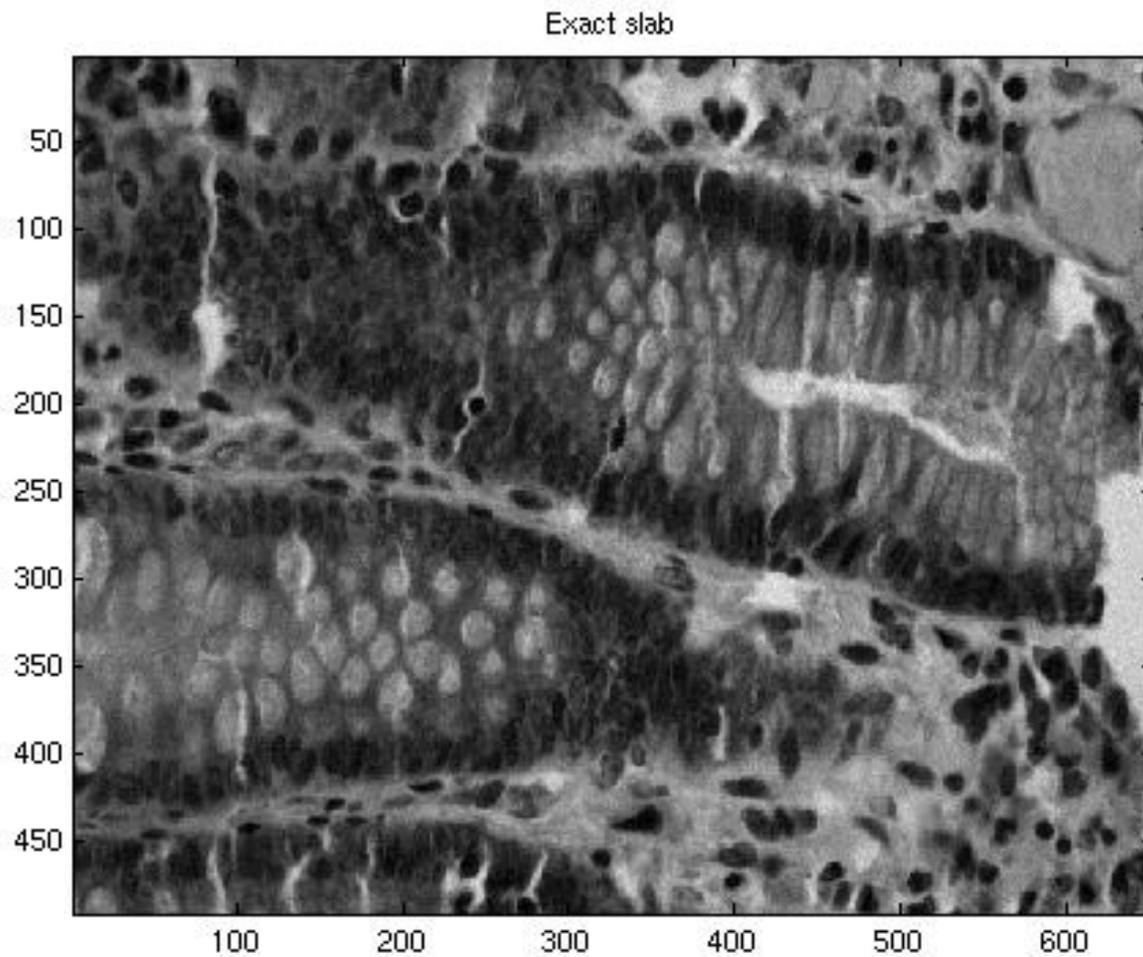
**The Data:** Images of a **single object** (e.g., earth or colon cells) at **many consecutive frequencies**.

**The Goal:** **Lossy compression, data reconstruction, and classification** using a small number of samples (images and/or pixels).



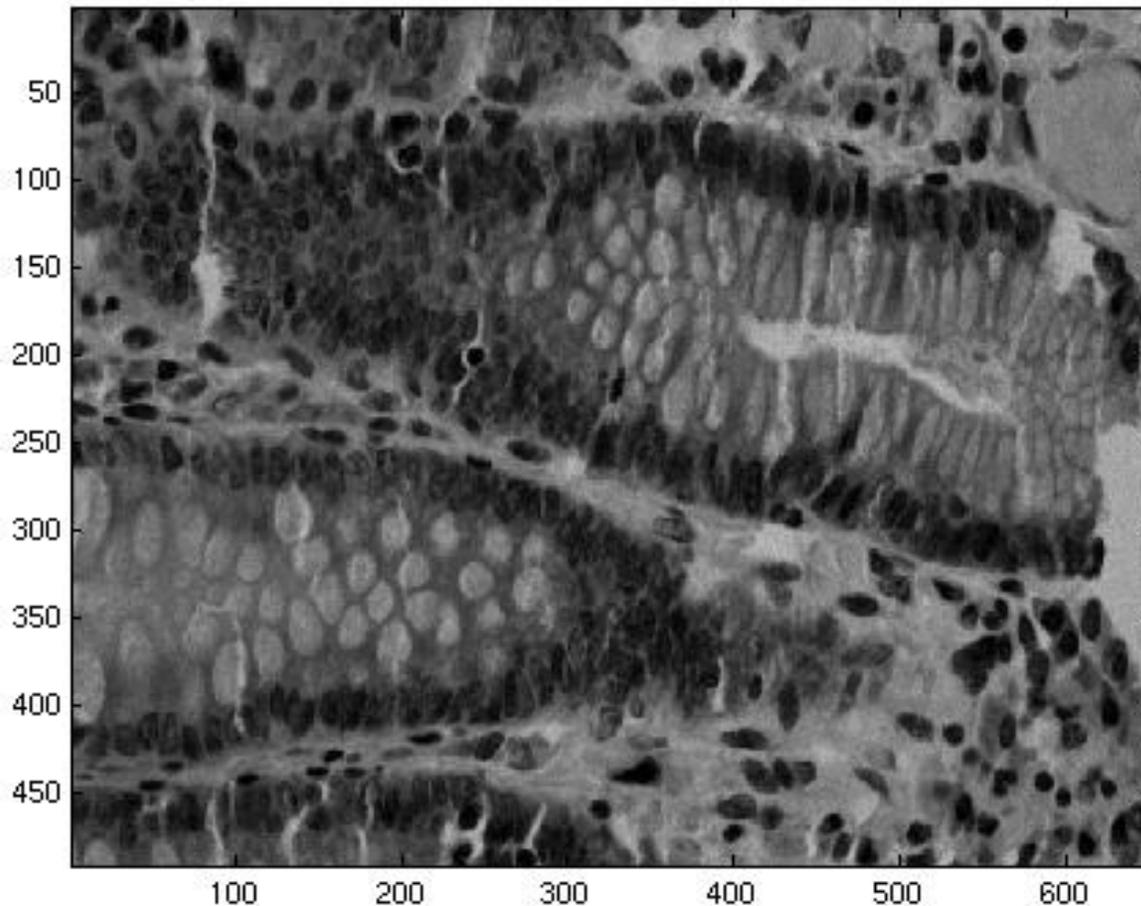


Look at the exact (65-th) slab.



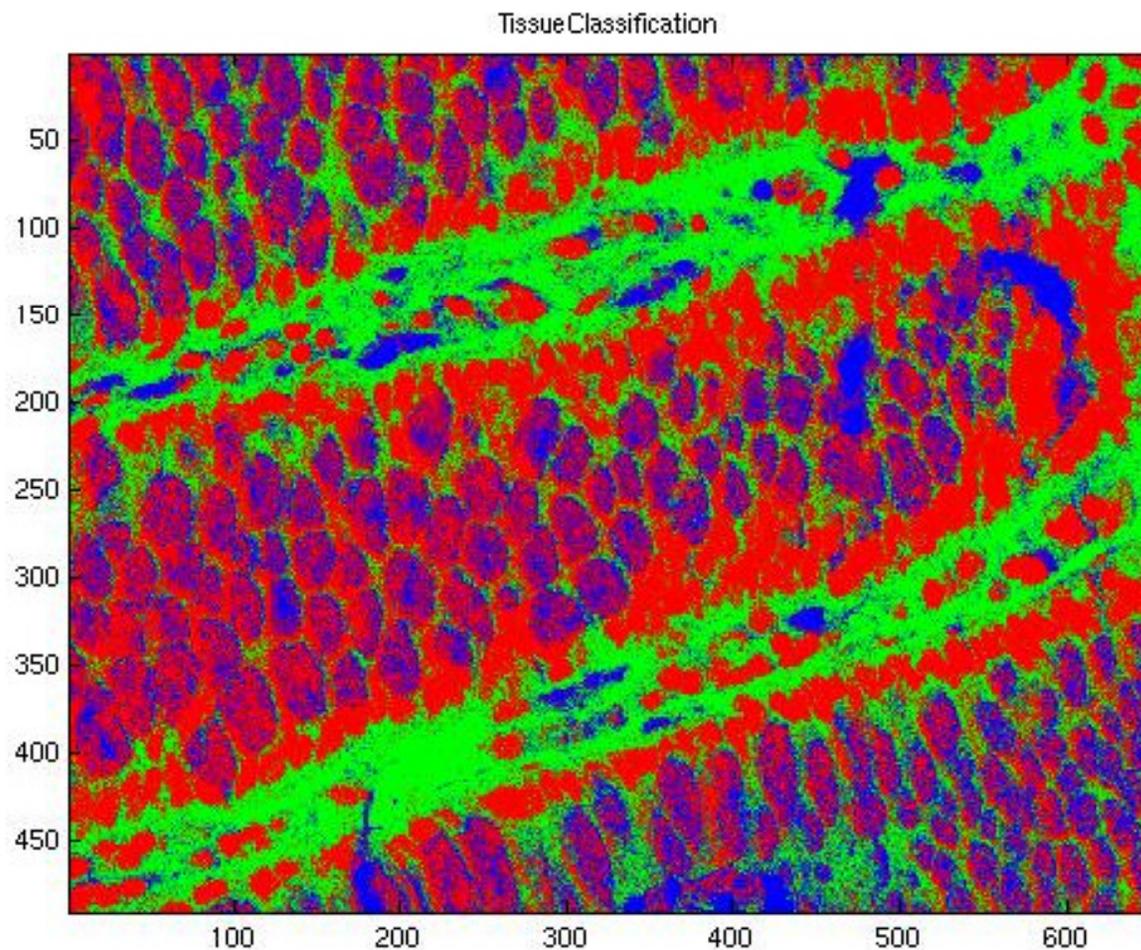
# The (65-th) slab approximately reconstructed

Approximate slab (approximate least squares fit)

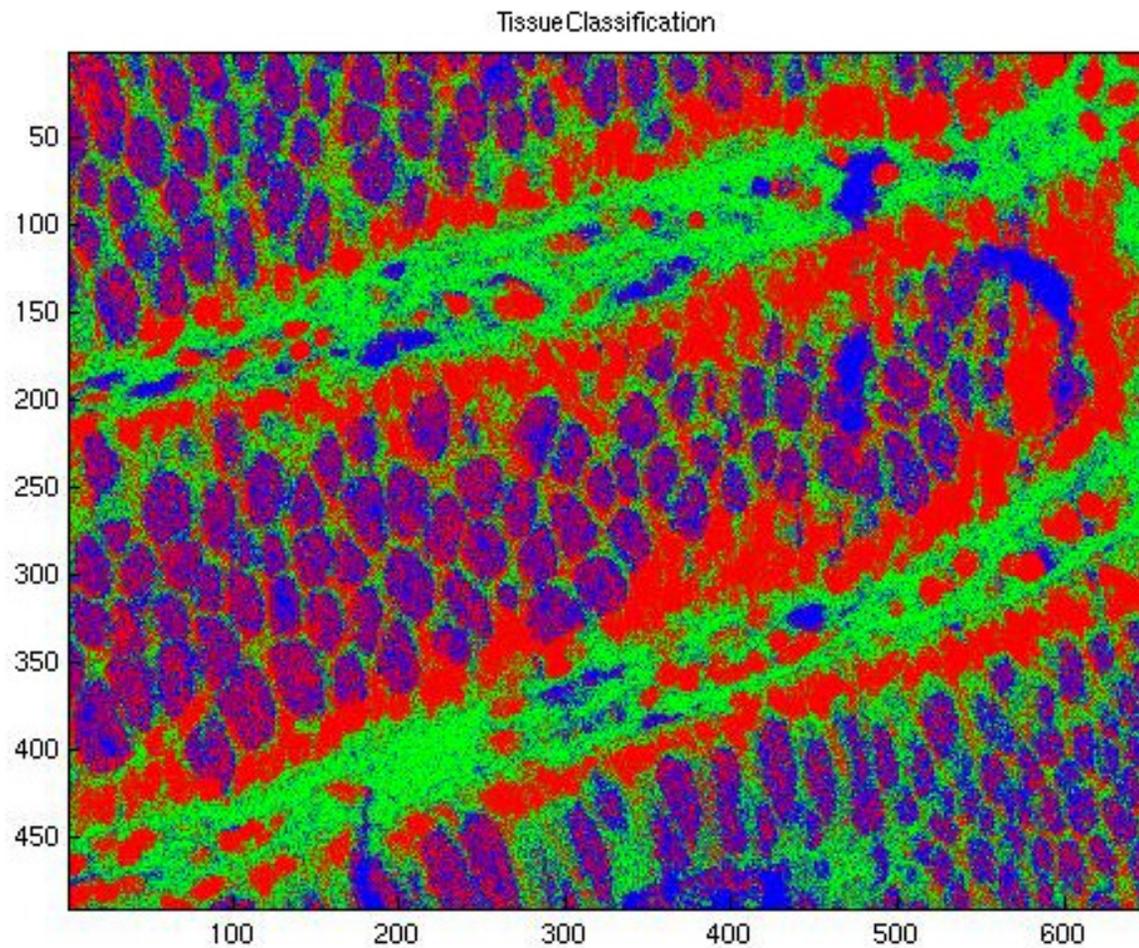


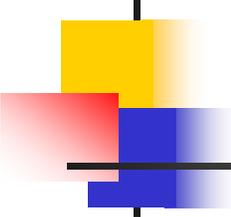
This slab was reconstructed by approximate least-squares fit to the basis from slabs 41 and 50, using 1000 (of 250K) pixels/fibers.

# Tissue Classification - Exact Data



# Tissue Classification - $N_s=12$ & $N_f=1000$





# CUR Data Application: Recommendation Systems

Problem:  $m$  customers and  $n$  products;  $A_{ij}$  is the (unknown) rating/utility of product  $j$  for customer  $i$ .

Goal: recreate  $A$  from a few samples to recommend high utility products.

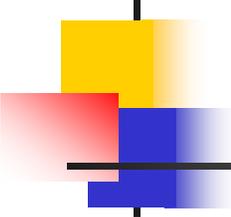
- (KRRT98): Assuming strong clustering of the products, competitive algorithms even with only 2 samples/customer.
- (AFKMS01): Assuming sampling of  $\Omega(mn)$  entries of  $A$  and a gap requirement, accurately recreate  $A$ .
- *Lots of applied work*, especially at large internet companies!

Q: Can we get competitive performance by sampling  $o(mn)$  elements?

A: Apply the CUR decomposition:

$$\begin{array}{c} \text{customers} \rightarrow \left( \begin{array}{c} \text{products} \\ A \end{array} \right) \approx \left( \begin{array}{c} \text{Customer sample} \\ \text{(purchases, small surveys)} \\ C \end{array} \right) \cdot \left( \begin{array}{c} U \end{array} \right) \cdot \left( \begin{array}{c} R \end{array} \right) \leftarrow \begin{array}{c} \text{Customer sample} \\ \text{(guinea pigs)} \end{array} \end{array}$$





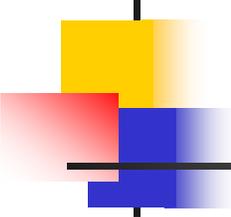
# Application to Jester Joke Recommendations

---

Use just the 14,140 “full” users who rated all 100 Jester jokes.

For each user, convert the utility vector to 100 x 100 pair-wise preference matrix.

Choose, e.g., 300 users (slabs), and a small number of comparisons (fibers).



# Conclusion

---

## CUR Low-Rank Matrix Decompositions:

- Uses *actual columns and/or rows*.
- Useful if data have *low-rank structure and other structure*.
- *Provable performance guarantees* within  $\varepsilon$  of "best."
- Performs well in practice on *genetic, medical imaging, and internet data*.

