# The convex geometry of inverse problems 

Benjamin Recht
Department of Computer Sciences University of Wisconsin-Madison

Joint work with
Venkat Chandrasekaran
Pablo Parrilo
Alan Willsky


## Linear Inverse Problems

- Find me a solution of

$$
y=\Phi x
$$

- $\Phi \mathrm{m} \times \mathrm{n}, \mathrm{m}<\mathrm{n}$
- Of the infinite collection of solutions, which one should we pick?
- Leverage structure:
- How do we design algorithms to solve underdetermined systems problems with priors?


## Sparsity

- 1-sparse vectors of Euclidean norm 1


$$
\|x\|_{1}=\sum_{i=1}^{n}\left|x_{i}\right|
$$

$$
\begin{array}{ll}
\operatorname{minimize} & \|x\|_{1} \\
\text { subject to } & \Phi x=y
\end{array}
$$



Compressed Sensing: Candes, Romberg, Tao, Donoho, Tanner, Etc...

## Rank

- $2 \times 2$ matrices
- plotted in 3d

$$
\left[\begin{array}{ll}
x & y \\
y & z
\end{array}\right]
$$



$$
\|X\|_{\star}=\sum_{i} \sigma_{i}(X)
$$

- $2 \times 2$ matrices
- plotted in 3d

$$
\begin{aligned}
& \left\|\left[\begin{array}{ll}
x & y \\
y & z
\end{array}\right]\right\|_{*} \leq 1 \\
& \|X\|_{*}=\sum_{i} \sigma_{i}(X)
\end{aligned}
$$

Nuclear Norm Heuristic


R, Fazel, and Parillo 2007
Rank Minimization/Matrix Completion

## Integer Programming

- Integer solutions: all components of $x$ are $\pm 1$


$$
\|x\|_{\infty}=\max _{i}\left|x_{i}\right|
$$

# minimize $\quad\|x\|_{\infty}$ <br> subject to $\quad \Phi x=y$ 



Donoho and Tanner 2008 Mangasarian and Recht. 2009.

## Parsimonious Models



- Search for best linear combination of fewest atoms
- "rank" = fewest atoms needed to describe the model



## Union of Subspaces



- X has structured sparsity: linear combination of elements from a set of subspaces $\left\{\mathrm{U}_{\mathrm{g}}\right\}$.
- Atomic set: unit norm vectors living in one of the $U_{g}$

$$
\|x\|_{\mathcal{G}}=\inf \left\{\sum_{g \in G}\left\|w_{g}\right\|: x=\sum_{g \in G} w_{g}, w_{g} \in U_{g}\right\}
$$

- Proposed by Jacob, Obozinski and Vert (2009).


## Permutation Matrices

- X a sum of a few permutation matrices
- Examples: Multiobject Tracking (Huang et al), Ranked elections (Jagabathula, Shah)
- Convex hull of the permutation matrices: Birkhoff Polytope of doubly stochastic matrices
- Permutahedra: convex hull of permutations of a fixed vector.

$$
[1,2,3,4]
$$




- Moments: convex hull of of $\left[1, t, t^{2}, t^{3}, t^{4}, \ldots\right]$, $t \in T$, some basic set.
- System Identification, Image Processing, Numerical Integration, Statistical Inference
- Solve with semidefinite programming
- Cut-matrices: sums of rank-one sign matrices.
- Collaborative Filtering, Clustering in Genetic Networks, Combinatorial Approximation Algorithms
- Approximate with semidefinite programming
- Low-rank Tensors: sums of rank-one tensors
- Computer Vision, Image Processing, Hyperspectral Imaging, Neuroscience
- Approximate with alternating leastsquares


## Atomic Norms

- Given a basic set of atoms, $\mathcal{A}$, define the function $\|x\|_{\mathcal{A}}=\inf \{t>0: x \in t \operatorname{conv}(\mathcal{A})\}$
- When $\mathcal{A}$ is centrosymmetric, we get a norm

$$
\begin{aligned}
& \|x\|_{\mathcal{A}}=\inf \left\{\sum_{a \in \mathcal{A}}\left|c_{a}\right|: x=\sum_{a \in \mathcal{A}} c_{a} a\right\} \\
& \text { IDEA: } \begin{array}{l}
\text { minimize } \quad\|z\|_{\mathcal{A}} \\
\text { subject to } \quad \Phi z=y
\end{array}
\end{aligned}
$$

- When does this work?
- How do we solve the optimization problem?


## Tangent Cones

- Set of directions that decrease the norm from $x$ form a cone:

$$
\mathcal{T}_{\mathcal{A}}(x)=\left\{d:\|x+\alpha d\|_{\mathcal{A}} \leq\|x\|_{\mathcal{A}} \text { for some } \alpha>0\right\}
$$



- $x$ is the unique minimizer if the intersection of this cone with the null space of $\Phi$ equals $\{0\}$


## Gaussian Widths

- When does a random subspace, $U$, intersect a convex cone $C$ at the origin?
- Gordon 88: with high probability if

$$
\operatorname{codim}(U) \geq w(C)^{2}
$$

- Where $w(C)=\mathbb{E}\left[\max _{x \in C \cap \mathbb{S}^{n-1}}\langle x, g\rangle\right]$ is the
Gaussian width.

$$
g \sim \mathcal{N}\left(0, I_{n}\right)
$$

- Corollary: For inverse problems: if $\Phi$ is a random Gaussian matrix with $m$ rows, need $m \geq w\left(\mathcal{T}_{\mathcal{A}}(x)\right)^{2}$ for recovery of $x$.


## Robust Recovery

- Suppose we observe $y=\Phi x+w \quad\|w\|_{2} \leq \delta$

$$
\begin{array}{ll}
\operatorname{minimize} & \|z\|_{\mathcal{A}} \\
\text { subject to } & \|\Phi z-y\| \leq \delta
\end{array}
$$

- If $\hat{x}$ is an optimal solution, then $\|x-\hat{x}\| \leq \frac{2 \delta}{\epsilon}$
provided that

$$
m \geq \frac{c_{0} w\left(\mathcal{T}_{\mathcal{A}}(x)\right)^{2}}{(1-\epsilon)^{2}}
$$

## Calculating Widths

- Hypercube:

$$
m \geq n / 2
$$

- Sparse Vectors, n vector, sparsity $\mathrm{s}<0.25 \mathrm{n}$

$$
m \geq 2 s\left(\log \left(\frac{n-s}{s}\right)+1\right)
$$

- Block sparse, M groups (possibly overlapping), maximum group size $B, k$ active groups

$$
m \geq 2 k(\log (M-k)+B)+k
$$



- Low-rank matrices: $n_{1} \times n_{2},\left(n_{1}<n_{2}\right)$, rank $r$

$$
m \geq 3 r\left(n_{1}+n_{2}-r\right)
$$

## General Cones

- Theorem: Let $C$ be a nonempty cone with polar cone $C^{*}$. Suppose C* subtends normalized solid angle $\mu$. Then

$$
w(C) \leq 3 \sqrt{\log \left(\frac{4}{\mu}\right)}
$$

- Corollary: For a vertex-transitive (i.e., "symmetric") polytope with p vertices, O(log p) Gaussian measurements are sufficient to recover a vertex via convex optimization.
- For $n \times n$ permutation matrix: $m=O(n \log n)$
- For $n \times n$ cut matrix: $m=O(n)$


## Algorithms

$$
\operatorname{minimize}_{z} \quad\|\Phi z-y\|_{2}^{2}+\mu\|z\|_{\mathcal{A}}
$$

- Naturally amenable to projected gradient algorithm:

$$
\text { "shrinkage" } \begin{array}{cc}
z_{k+1}=\Pi_{\eta \mu}\left(z_{k}-\eta \Phi^{*} r_{k}\right) \\
\Pi_{\tau}(z)=\arg \min _{u} \frac{1}{2}\|z-u\|^{2}+\tau\|u\|_{\mathcal{A}}
\end{array}
$$

- Relaxations: $\mathcal{A}_{1} \subset \mathcal{A}_{2} \Longrightarrow\|x\|_{\mathcal{A}_{2}} \leq\|x\|_{\mathcal{A}_{1}}$


NB! tangent cone gets wider

- Hierarchy of relaxations based on $\theta$-Bodies yield progressively tighter bounds on the atomic norm


## Atomic Norm Decompositions

- Propose a natural convex heuristic for enforcing prior information in inverse problems
- Bounds for the linear case: heuristic succeeds for most sufficiently large sets of measurements
- Stability without restricted isometries
- Standard program for computing these bounds: distance to normal cones
- Algorithms and approximation schemes for computationally difficult priors


## Extensions...

- Width Calculations for more general structures
- Recovery bounds for structured measurement matrices (application specific)
- Incorporating stochastic noise models
- Understanding of the loss due to convex relaxation and norm approximation
- Scaling generalized shrinkage algorithms to massive data sets


## JeLLYFISH

- SGD for Matrix Factorizations.
with Christopher Ré
Example: minimize $\quad \sum_{(u, v) \in E}\left(X_{u v}-M_{u v}\right)^{2}+\mu\|\mathbf{X}\|_{*}$
- Idea: approximate $\mathbf{X} \approx \mathbf{L R}^{T}$ $\operatorname{minimize}_{(\mathbf{L}, \mathbf{R})} \sum_{(u, v) \in E}\left\{\left(\mathbf{L}_{u} \mathbf{R}_{v}^{T}-Z_{u v}\right)^{2}+\mu_{u}\left\|\mathbf{L}_{u}\right\|_{F}^{2}+\mu_{v}\left\|\mathbf{R}_{v}\right\|_{F}^{2}\right\}$
- Step 1: Pick ( $u, v$ ) and compute residual:

$$
e=\left(\mathbf{L}_{u} \mathbf{R}_{v}^{T}-Z_{u v}\right)
$$

- Step 2: Take a gradient step:

$$
\left[\begin{array}{l}
\mathbf{L}_{u} \\
\mathbf{R}_{v}
\end{array}\right] \leftarrow\left[\begin{array}{c}
\left(1-\gamma \mu_{u}\right) \mathbf{L}_{u}-\gamma e \mathbf{R}_{v} \\
\left(1-\gamma \mu_{v}\right) \mathbf{R}_{v}-\gamma e \mathbf{L}_{u}
\end{array}\right]
$$

## JELLYFISH

Observation: With replacement sample=poor locality Idea: Bias sample to improve locality.

Algorithm: Shuffle the data.

1. Process $\left\{L_{1} R_{1}, L_{2} R_{2}, L_{3} R_{3}\right\}$ in parallel
2. Process $\left\{L_{1} R_{2}, L_{2} R_{3}, L_{3} R_{1}\right\}$ in parallel
3. Process $\left\{L_{1} R_{3}, L_{2} R_{1}, L_{3} R_{2}\right\}$ in parallel

Big win: No locks!
(model access)

## Example Optimization

## Example Optimization

- Shuffle all the rows and columns


## Example Optimization



- Shuffle all the rows and columns


## Example Optimization



- Shuffle all the rows and columns
- Apply block partitioning


## Example Optimization

- Shuffle all the rows and columns
- Apply block partitioning


## Example Optimization

- Shuffle all the rows and columns
- Apply block partitioning
- Train on each block independently


## Example Optimization



- Shuffle all the rows and columns
- Apply block partitioning
- Train on each block independently


## Example Optimization

- Shuffle all the rows and columns
- Apply block partitioning
- Train on each block independently


## Example Optimization

- Shuffle all the rows and columns
- Apply block partitioning
- Train on each block independently
- Repeat...


## Example Optimization



- Shuffle all the rows and columns
- Apply block partitioning
- Train on each block independently
- Repeat...


## Example Optimization



- Shuffle all the rows and columns
- Apply block partitioning
- Train on each block independently
- Repeat...


## Example Optimization



- Shuffle all the rows and columns
- Apply block partitioning
- Train on each block independently
- Repeat...


## Example Optimization



- Shuffle all the rows and columns
- Apply block partitioning
- Train on each block independently
- Repeat...


## Example Optimization



- Shuffle all the rows and columns
- Apply block partitioning
- Train on each block independently
- Repeat...
- Solves Netflix prize in under 1 minute on a 40 core machine
- Over 100x faster than standard solvers

