# Learning in high dimension 

Jean-Philippe Vert


## The « $n \ll p$ » problem



## How to learn with $n \ll p$ ?

1. Simplify data: pairwise comparisons
2. Add prior knowledge: structured feature selection

## How to learn with $n \ll p$ ?

## 1. Simplify data: pairwise comparisons

2. Add prior knowledge: structured feature selection

## Top Scoring Pairs (TSP)

(a) TSP


$$
\text { IF SPTAN1 => CD33* THEN ALL, ELSE AML. } \quad \Delta=0.9787
$$

(b) $k$-TSP

IF SPTAN1 => CD33* THEN ALL, ELSE AML.
IF TCF3* > APLP2 THEN ALL, ELSE AML.

$$
\text { IF ATP2A3* }=>\text { CST3* THEN ALL, ELSE AML. }
$$

IF DGKD > MGST1 THEN ALL, ELSE AML.
IF CCND3* => NPC2 THEN ALL, ELSE AML.
IF TOP2B* > PLCB2 THEN ALL, ELSE AML.
IF Macmarcks => CTSD* THEN ALL, ELSE AML.

$$
\begin{aligned}
& \Delta=0.9787 \\
& \Delta=0.9787 \\
& \Delta=0.9574 \\
& \Delta=0.9387 \\
& \Delta=0.9387 \\
& \Delta=0.9387 \\
& \Delta=0.9387 \\
& \Delta=0.9362 \\
& \Delta=0.9200
\end{aligned}
$$

(Geman et al., 2004; Tan et al., 2005; Leek, 2009; ...)

## Generalization of TSP



## Practical problem



# Storing O(p^2) bits per sample 

Training a linear model in O(p^2) dimensions

## A trick


+kernel trick = we can train linear models efficiently

## Experiment

## Datasets

| Dataset | No. of features | No. of samples (training/test) |  |
| :---: | :---: | :---: | :---: |
|  |  | $C_{1}$ | $C_{2}$ |
| Breast Cancer 1 | 23624 | $44 / 7$ (Non-relapse) | $32 / 12$ (Relapse) |
| Breast Cancer 2 | 22283 | 142 (Non-relapse) | 56 (Relapse) |
| Breast Cancer 3 | 22283 | 71 (Poor Prognosis) | 138 (Good Prognosis) |
| Colon Tumor | 2000 | 40 (Tumor) | 22 (Normal) |
| Lung Cancer 1 | 7129 | 24 (Poor Prognosis) | 62 (Good Prognosis) |
| Lung Cancer 2 | 12533 | $16 / 134$ (ADCA) | $16 / 15$ (MPM) |
| Medulloblastoma | 7129 | 39 (Failure) | 21 (Survivor) |
| Ovarian Cancer | 15154 | 162 (Cancer) | 91 (Normal) |
| Prostate Cancer 1 | 12600 | $50 / 9$ (Normal) | $52 / 25$ (Tumor) |
| Prostate Cancer 2 | 12600 | 13 (Non-relapse) | 8 (Relapse) |

## Methods

- Kernel machines Support Vector Machines (SVM) and Kernel Fisher Discriminant (KFD) with Kendall kernel, linear kernel, Gaussian RBF kernel, polynomial kernel.
- Top Scoring Pairs (TSP) classifiers [Tan et al., 2005].
- Hybrid scheme of SVM + TSP feature selection algorithm.


## Results



## Summary

- Robust representation as $O\left(p^{\wedge} 2\right)$ bits
- Computationally efficient (Kendall kernel)
- Good accuracy
- Extension to missing values OK
- Extension to « fuzzy comparison » OK
- Open questions:
- robustness across technologies (Patil et al., 2015) ?
- correction for batch / structure?


## How to learn with $n \ll p$ ?

1. Simplify data: pairwise comparisons
2. Add prior knowledge: structured feature selection

## Feature Selection


$p$ features

« Molecular signature»

Also relevant for

- isoform identification from RNA-seq data (IsoLasso, FlipFlop etc...)
- gene network inference (GENIE3, TIGRESS, etc...)


## Early disappointments...

## Gene expression profiling predicts clinical outcome of breast cancer

Laura J. van ’t Veer ${ }^{\star} \dagger$, Hongyue Dai $\dagger \ddagger$, Marc J. van de Vijver ${ }^{\star} \dagger$, Yudong D. He $\ddagger$, Augustinus A. M. Hart*, Mao Mao $\ddagger$, Hans L. Peterse ${ }^{\star}$, Karin van der Kooy*, Matthew J. Marton $\ddagger$, Anke T. Witteveen*, George J. Schreiber $\ddagger$, Ron M. Kerkhoven ${ }^{\star}$, Chris Roberts $\ddagger$, Peter S. Linsley $\ddagger$, René Bernards ${ }^{*}$ \& Stephen H. Friend $\ddagger$

70 genes (Nature, 2002)

Gene-expression profiles to predict distant metastasis of lymph-node-negative primary breast cancer

## Not because of feature selection method


(Haury et al., 2011)

## What's wrong?

## Increasing n helps




Can we try to « decrease p »?
Add prior knowledge,
Structured feature selection

## Sparsity with the LASSO

- Linear model
$f(x)=w 1 x 1+w 2 x 2+\ldots+w P x P$
- Sparse when wK=0 for many K's
- Learn a sparse model by minimize Error(w)
such that
w is in the grey box 0
- O is convex -> efficient algorith
- O has edges -> sparsity


# Structured sparsity with atomic norms 

1) Choose a set of ATOMS
(Chandrasekaran et al., 2012, ...)

# Structured sparsity with atomic norms 

1) Choose a set of ATOMS
2) Take the convex hull $O$

(Chandrasekaran et al., 2012, ...)

# Structured sparsity with atomic norms 

1) Choose a set of ATOMS
2) Take the convex hull
3) Minimize Error(w) such that w is in the convex hull

The solution is a sparse model over the ATOMS!
(Chandrasekaran et al., 2012, ...)

## Quizz: where are the atoms?


$\|w\|_{2}$
Ridge
$\|w\|_{1}$
Lasso
$\sqrt{w_{1}^{2}+w_{2}^{2}}+\left|w_{3}\right|$
Group Lasso

Trace norm

## Graph Lasso



To select features that tend to be connected over a given network
(Jacob et al., 2009)

# Breast cancer prognosis signature with Lasso (accuracy=61\%) 



## Breast cancer prognosis signature with Graph Lasso (accuracy=64\%)



## Joint isoform detection from multiple RNA-Seq samples



Elsa Bernard


Laurent Jacob


Julien Mairal


Eric Viara


## Learning sparse models with disjoint support?

## Motivation

- Multiclass or multi-task classification problems
- Eg: cascade of classifiers




## An atomic norm


(Vervier et al., 2014)

## Application: Microbial identification from MALDI-TOF MS spectra



## Learning low-rank matrices with sparse factors?



## An atomic norm



## An atomic norm

$$
X=\sum_{i=1}^{r} u_{i} v_{i}^{\top}
$$

Theorem
Learning with this norm is «statistically optimal» to infer sparse low-rank matrices

But<br>Convex but NP-hard


E. Richard

(Richard et al., 2014)

## Preliminary results on sparse PCA







| Sample covariance | Trace | $\ell_{1}$ | Trace $+\ell_{1}$ | Sequential | $\Omega_{k, \succeq}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $4.20 \pm 0.02$ | $0.98 \pm 0.01$ | $2.07 \pm 0.01$ | $0.96 \pm 0.01$ | $0.93 \pm 0.08$ | $\mathbf{0 . 5 9} \pm \mathbf{0 . 0 3}$ |

(Richard et al., 2014)

## Summary

- Include prior knowledge: « sparse on some dictionary »
- Convex, (usually) computationally efficient
- Leads to interpretable model
- Good framework for data integration


## Thanks



## Future

- Find representations simple (for statistical reasons), robust to artefacts (batch, technology, ...)
- $\mathrm{n} \ll \mathrm{p}$ still far from solved

