## Supervised classification for structured data: Applications in bio- and chemoinformatics

Jean-Philippe Vert<br>Jean-Philippe.Vert@mines-paristech.fr<br>Mines ParisTech / Curie Institute / Inserm

The third school on The Analysis of Patterns, Pula Science Park, Italy, June 1, 2009

## Virtual screening for drug discovery



NCI AIDS screen results (from http://cactus.nci.nih.gov).

## Image retrieval and classification



From Harchaoui and Bach (2007).

## Cancer diagnosis



## Cancer prognosis







## Pattern recognition, aka supervised classification



## Pattern recognition, aka supervised classification



## Pattern recognition, aka supervised classification



## Pattern recognition, aka supervised classification



## Pattern recognition, aka supervised classification



## Formalization

## The problem

- Given a set of training instances $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)$, where $x_{i} \in \mathcal{X}$ are data and $y_{i} \in \mathcal{Y}$ are continuous or discrete variables of interest,
- Estimate a function

$$
y=f(x)
$$

where $x$ is any new data to be labeled.

- $f$ should be accurate and intepretable.


## Linear classifiers

## The model

- Each sample $x \in \mathcal{X}$ is represented by a vector of features (or descriptors, or patterns):

$$
\Phi(x)=\left(\Phi_{1}(x), \ldots, \Phi_{p}(x)\right)
$$

- Based on the training set we estimate a linear function:

$$
f_{\beta}(x)=\sum_{i=1}^{p} \beta_{i} \Phi_{i}(x)=\beta^{\top} \Phi(x)
$$

## Two (related) questions

- How to design the features $\Phi(x)$ ?
- How to estimate the model $\beta$ ?


## Linear classifiers

## The model

- Each sample $x \in \mathcal{X}$ is represented by a vector of features (or descriptors, or patterns):

$$
\Phi(x)=\left(\Phi_{1}(x), \ldots, \Phi_{p}(x)\right)
$$

- Based on the training set we estimate a linear function:

$$
f_{\beta}(x)=\sum_{i=1}^{p} \beta_{i} \Phi_{i}(x)=\beta^{\top} \Phi(x)
$$

## Two (related) questions

- How to design the features $\Phi(x)$ ?
- How to estimate the model $\beta$ ?


## Outline

(9) Explicit computation of features: the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification

4 Conclusion

## Outline

(1) Explicit computation of features: the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification

4. Conclusion

## Motivation



NCI AIDS screen results (from http://cactus.nci.nih.gov).

## The approach

(1) Represent explicitly each graph $x$ by a vector of fixed dimension $\Phi(x) \in \mathbb{R}^{p}$.
(2) Use an algorithm for regression or pattern recognition in $\mathbb{R}^{p}$.


## The approach

- Represent explicitly each graph $x$ by a vector of fixed dimension $\Phi(x) \in \mathbb{R}^{p}$.


## (3) Use an algorithm for regression or pattern recognition in $\mathbb{R}^{p}$.



## The approach

- Represent explicitly each graph $x$ by a vector of fixed dimension $\Phi(x) \in \mathbb{R}^{p}$.
(2) Use an algorithm for regression or pattern recognition in $\mathbb{R}^{p}$.



## Example

## 2D structural keys in chemoinformatics

- Index a molecule by a binary fingerprint defined by a limited set of pre-defined stuctures

- Use a machine learning algorithms such as SVM, NN, PLS, decision tree, ...


## Challenge: which descriptors (patterns)?



- Expressiveness: they should retain as much information as possible from the graph
- Computation : they should be fast to compute
- Large dimension of the vector representation: memory storage, speed, statistical issues


## Indexing by substructures



- Often we believe that the presence substructures are important predictive patterns
- Hence it makes sense to represent a graph by features that indicate the presence (or the number of occurrences) of particular substructures
- However, detecting the presence of particular substructures may be computationally challenging...


## Subgraphs

## Definition

A subgraph of a graph $(V, E)$ is a connected graph $\left(V^{\prime}, E^{\prime}\right)$ with $V^{\prime} \subset V$ and $E^{\prime} \subset E$.




## Indexing by all subgraphs?



## Theorem

Computing all subgraph occurrences is NP-hard.

## Proof.

- The linear graph of size $n$ is a subgraph of a graph $X$ with $n$ vertices iff $X$ has an Hamiltonian path
- The decision problem whether a graph has a Hamiltonian path is NP-complete.


## Indexing by all subgraphs?



## Theorem

Computing all subgraph occurrences is NP-hard.

## Proof.

- The linear graph of size $n$ is a subgraph of a graph $X$ with $n$ vertices iff $X$ has an Hamiltonian path
- The decision problem whether a graph has a Hamiltonian path is NP-complete.


## Indexing by all subgraphs?



## Theorem

Computing all subgraph occurrences is NP-hard.

## Proof.

- The linear graph of size $n$ is a subgraph of a graph $X$ with $n$ vertices iff $X$ has an Hamiltonian path
- The decision problem whether a graph has a Hamiltonian path is NP-complete.


## Paths

## Definition

- A path of a graph $(V, E)$ is sequence of distinct vertices $v_{1}, \ldots, v_{n} \in V\left(i \neq j \Longrightarrow v_{i} \neq v_{j}\right)$ such that $\left(v_{i}, v_{i+1}\right) \in E$ for $i=1, \ldots, n-1$.
- Equivalently the paths are the linear subgraphs.




## Indexing by all paths?



## Theorem

Computing all path occurrences is NP-hard.

## Proof.

## Same as for subgraphs.

## Indexing by all paths?



## Theorem

Computing all path occurrences is NP-hard.

## Proof. <br> Same as for subgraphs.

## Indexing by all paths?



## Theorem

Computing all path occurrences is NP-hard.

## Proof.

## Same as for subgraphs.

## Indexing by what?

## Substructure selection

We can imagine more limited sets of substuctures that lead to more computationnally efficient indexing (non-exhaustive list)

- substructures selected by domain knowledge (MDL fingerprint)
- all path up to length $k$ (Openeye fingerprint, Nicholls 2005)
- all shortest paths (Borgwardt and Kriegel, 2005)
- all subgraphs up to $k$ vertices (graphlet kernel, Sherashidze et al., 2009)
- all frequent subgraphs in the database (Helma et al., 2004)


## Example : Indexing by all shortest paths



## Properties (Borgwardt and Kriegel, 2005)

- There are $O\left(n^{2}\right)$ shortest naths.
- The vector of counts can be computed in $O\left(n^{4}\right)$ with the Floyd-Warshall algorithm.


## Example : Indexing by all shortest paths



## Properties (Borgwardt and Kriegel, 2005)

- There are $O\left(n^{2}\right)$ shortest paths.
- The vector of counts can be computed in $O\left(n^{4}\right)$ with the Floyd-Warshall algorithm.


## Example : Indexing by all subgraphs up to $k$ vertices



## Properties (Shervashidze et al., 2009)

- Naive enumeration scales as $O\left(n^{k}\right)$
- Enumeration of connected graphlets in $O\left(n d^{k-1}\right)$ for graphs with degree $\leq d$ and $k \leq 5$.
- Randomly sample subaraphs if enumeration is infeasible.


## Example : Indexing by all subgraphs up to $k$ vertices



## Properties (Shervashidze et al., 2009)

- Naive enumeration scales as $O\left(n^{k}\right)$.
- Enumeration of connected graphlets in $O\left(n d^{k-1}\right)$ for graphs with degree $\leq d$ and $k \leq 5$.
- Randomly sample subgraphs if enumeration is infeasible.


## Summary

- Explicit computation of substructure occurrences can be computationnally prohibitive (subgraph, paths)
- Several ideas to reduce the set of substructures considered
- In practice, NP-hardness may not be so prohibitive (e.g., graphs with small degrees), the strategy followed should depend on the data considered.


## Outline

(1) Explicit computation of features: the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification

4. Conclusion

## Outline

(1) Explicit computation of features: the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification

4 Conclusion

## Positive definite kernels

## Definition

- Let $\Phi(x)$ be a vector representation of the data $x$
- The kernel between two graphs is defined by:

$$
K\left(x, x^{\prime}\right)=\Phi(x)^{\top} \Phi\left(x^{\prime}\right)
$$



## The kernel trick

## The trick

- Many linear algorithms for regression or pattern recognition can be expressed only in terms of inner products between vectors
- Computing the kernel is often more efficient than computing $\Phi(x)$, especially in high or infinite dimensions!
- Perhaps we can consider more features with kernels than with explicit feature computation?


## Learning linear classifiers with kernels

## Training the model

- Minimize an empirical risk on the training samples:

$$
\min _{\beta \in \mathbb{R}^{p+1}} R_{e m p}(\beta)=\frac{1}{n} \sum_{i=1}^{n} I\left(\beta^{\top} \Phi\left(x_{i}\right), y_{i}\right)
$$

- ... subject to a constraint on $\beta$ :

$$
\|\beta\| \leq C
$$



## Making kernels

## Two important strategies (not the only ones!)

- Feature design :

$$
K\left(x, x^{\prime}\right)=\Phi(x)^{\top} \Phi\left(x^{\prime}\right)
$$

We illustrate this idea with graph kernels.

- Regularization design :

$$
\|\beta\| \leq C
$$

We illustrate this idea with kernels for microarray data.

## Outline

(1) Explicit computation of features: the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification
(4) Conclusion


## The idea

## (1) Represent implicitly each graph $x$ by a vector $\Phi(x) \in \mathcal{H}$ through the kernel

$$
K\left(x, x^{\prime}\right)=\phi(x)^{\top} \Phi\left(x^{\prime}\right)
$$

## (2) Use a kernel method for classification in $\mathcal{H}$.



## The idea

(1) Represent implicitly each graph $x$ by a vector $\Phi(x) \in \mathcal{H}$ through the kernel

$$
K\left(x, x^{\prime}\right)=\Phi(x)^{\top} \Phi\left(x^{\prime}\right) .
$$

## (3) Use a kernel method for classification in $\mathcal{H}$.



## The idea

(1) Represent implicitly each graph $x$ by a vector $\Phi(x) \in \mathcal{H}$ through the kernel

$$
K\left(x, x^{\prime}\right)=\Phi(x)^{\top} \Phi\left(x^{\prime}\right) .
$$

(2) Use a kernel method for classification in $\mathcal{H}$.


## Expressiveness vs Complexity

Definition: Complete graph kernels
A graph kernel is complete if it separates non-isomorphic graphs, i.e.:

$$
\forall G_{1}, G_{2} \in \mathcal{X}, \quad d_{K}\left(G_{1}, G_{2}\right)=0 \Longrightarrow G_{1} \simeq G_{2}
$$

Equivalently, $\Phi\left(G_{1}\right) \neq \Phi\left(G_{1}\right)$ if $G_{1}$ and $G_{2}$ are not isomorphic.
Expressiveness vs Complexity trade-off

- If a graph kernel is not complete, then there is no hope to learn all possible functions over $\mathcal{X}$ : the kernel is not expressive enough.
- On the other hand, kernel computation must be tractable, i.e., no more than polynomial (with small degree) for practical
applications.
- Can we define tractable and expressive graph kernels?


## Expressiveness vs Complexity

## Definition: Complete graph kernels

A graph kernel is complete if it separates non-isomorphic graphs, i.e.:

$$
\forall G_{1}, G_{2} \in \mathcal{X}, \quad d_{K}\left(G_{1}, G_{2}\right)=0 \Longrightarrow G_{1} \simeq G_{2}
$$

Equivalently, $\Phi\left(G_{1}\right) \neq \Phi\left(G_{1}\right)$ if $G_{1}$ and $G_{2}$ are not isomorphic.

## Expressiveness vs Complexity trade-off

- If a graph kernel is not complete, then there is no hope to learn all possible functions over $\mathcal{X}$ : the kernel is not expressive enough.
- On the other hand, kernel computation must be tractable, i.e., no more than polynomial (with small degree) for practical applications.
- Can we define tractable and expressive graph kernels?


## Complexity of complete kernels

## Proposition (Gärtner et al., 2003)

Computing any complete graph kernel is at least as hard as the graph isomorphism problem.

## Proof

- For any kernel $K$ the complexity of computing $d_{K}$ is the same as the complexity of computing $K$, because:

$$
d_{K}\left(G_{1}, G_{2}\right)^{2}=K\left(G_{1}, G_{1}\right)+K\left(G_{2}, G_{2}\right)-2 K\left(G_{1}, G_{2}\right)
$$

- If $K$ is a complete graph kernel, then computing $d_{K}$ solves the graph isomorphism problem $\left(d_{K}\left(G_{1}, G_{2}\right)=0\right.$ iff $\left.G_{1} \simeq G_{2}\right) . \square$


## Complexity of complete kernels

## Proposition (Gärtner et al., 2003)

Computing any complete graph kernel is at least as hard as the graph isomorphism problem.

## Proof

- For any kernel $K$ the complexity of computing $d_{K}$ is the same as the complexity of computing $K$, because:

$$
d_{K}\left(G_{1}, G_{2}\right)^{2}=K\left(G_{1}, G_{1}\right)+K\left(G_{2}, G_{2}\right)-2 K\left(G_{1}, G_{2}\right)
$$

- If $K$ is a complete graph kernel, then computing $d_{K}$ solves the graph isomorphism problem $\left(d_{K}\left(G_{1}, G_{2}\right)=0\right.$ iff $\left.G_{1} \simeq G_{2}\right)$.


## Subgraph kernel

## Definition

- Let $\left(\lambda_{G}\right)_{G \in \mathcal{X}}$ a set or nonnegative real-valued weights
- For any graph $G \in \mathcal{X}$, let

$$
\forall H \in \mathcal{X}, \quad \Phi_{H}(G)=\mid\left\{G^{\prime} \text { is a subgraph of } G: G^{\prime} \simeq H\right\} \mid .
$$

- The subgraph kernel between any two graphs $G_{1}$ and $G_{2} \in \mathcal{X}$ is defined by:

$$
K_{\text {subgraph }}\left(G_{1}, G_{2}\right)=\sum_{H \in \mathcal{X}} \lambda_{H} \Phi_{H}\left(G_{1}\right) \Phi_{H}\left(G_{2}\right)
$$



## Subgraph kernel complexity

Proposition (Gärtner et al., 2003)
Computing the subgraph kernel is NP-hard.
Proof (1/2)

- Let $P_{n}$ be the path graph with $n$ vertices.
- Subgraphs of $P_{n}$ are path graphs:

$$
\Phi\left(P_{n}\right)=n e_{P_{1}}+(n-1) e_{P_{2}}+\ldots+e_{P_{n}}
$$

- The vectors $\Phi\left(P_{1}\right), \ldots, \Phi\left(P_{n}\right)$ are linearly independent, therefore:

where the coefficients $\alpha_{i}$ can be found in polynomial time (solving a $n \times n$ triangular system).


## Subgraph kernel complexity

## Proposition (Gärtner et al., 2003)

Computing the subgraph kernel is NP-hard.

## Proof (1/2)

- Let $P_{n}$ be the path graph with $n$ vertices.
- Subgraphs of $P_{n}$ are path graphs:

$$
\Phi\left(P_{n}\right)=n e_{P_{1}}+(n-1) e_{P_{2}}+\ldots+e_{P_{n}} .
$$

- The vectors $\Phi\left(P_{1}\right), \ldots, \Phi\left(P_{n}\right)$ are linearly independent, therefore:

$$
e_{P_{n}}=\sum_{i=1}^{n} \alpha_{i} \Phi\left(P_{i}\right)
$$

where the coefficients $\alpha_{i}$ can be found in polynomial time (solving a $n \times n$ triangular system).

## Subgraph kernel complexity

## Proposition (Gärtner et al., 2003)

Computing the subgraph kernel is NP-hard.

## Proof (2/2)

- If $G$ is a graph with $n$ vertices, then it has a path that visits each node exactly once (Hamiltonian path) if and only if $\Phi(G)^{\top} e_{n}>0$, i.e.,

$$
\Phi(G)^{\top}\left(\sum_{i=1}^{n} \alpha_{i} \Phi\left(P_{i}\right)\right)=\sum_{i=1}^{n} \alpha_{i} K_{\text {subgraph }}\left(G, P_{i}\right)>0
$$

- The decision problem whether a graph has a Hamiltonian path is NP-complete.


## Path kernel



## Definition

The path kernel is the subgraph kernel restricted to paths, i.e.,

$$
K_{\text {path }}\left(G_{1}, G_{2}\right)=\sum_{H \in \mathcal{P}} \lambda_{H} \Phi_{H}\left(G_{1}\right) \Phi_{H}\left(G_{2}\right)
$$

where $\mathcal{P} \subset \mathcal{X}$ is the set of path graphs.
Proposition (Gärtner et al., 2003)
Computing the path kernel is NP-hard.

## Path kernel



## Definition

The path kernel is the subgraph kernel restricted to paths, i.e.,

$$
K_{\text {path }}\left(G_{1}, G_{2}\right)=\sum_{H \in \mathcal{P}} \lambda_{H} \Phi_{H}\left(G_{1}\right) \Phi_{H}\left(G_{2}\right)
$$

where $\mathcal{P} \subset \mathcal{X}$ is the set of path graphs.

## Proposition (Gärtner et al., 2003)

Computing the path kernel is NP-hard.

## Summary

## Expressiveness vs Complexity trade-off

- It is intractable to compute complete graph kernels.
- It is intractable to compute the subgraph kernels.
- Restricting subgraphs to be linear does not help: it is also intractable to compute the path kernel.
- One approach to define polynomial time computable graph kernels is to have the feature space be made up of graphs homomorphic to subgraphs, e.g., to consider walks instead of paths.


## Walks

## Definition

- A walk of a graph $(V, E)$ is sequence of $v_{1}, \ldots, v_{n} \in V$ such that $\left(v_{i}, v_{i+1}\right) \in E$ for $i=1, \ldots, n-1$.
- We note $\mathcal{W}_{n}(G)$ the set of walks with $n$ vertices of the graph $G$, and $\mathcal{W}(G)$ the set of all walks.



## Walks $\neq$ paths



## Walk kernel

## Definition

- Let $\mathcal{S}_{n}$ denote the set of all possible label sequences of walks of length $n$ (including vertices and edges labels), and $\mathcal{S}=\cup_{n \geq 1} \mathcal{S}_{n}$.
- For any graph $\mathcal{X}$ let a weight $\lambda_{G}(w)$ be associated to each walk $w \in \mathcal{W}(G)$.
- Let the feature vector $\Phi(G)=\left(\Phi_{s}(G)\right)_{s \in \mathcal{S}}$ be defined by:

$$
\Phi_{s}(G)=\sum_{w \in \mathcal{W}(G)} \lambda_{G}(w) 1(s \text { is the label sequence of } w) .
$$

- A walk kernel is a graph kernel defined by:


## Walk kernel

## Definition

- Let $\mathcal{S}_{n}$ denote the set of all possible label sequences of walks of length $n$ (including vertices and edges labels), and $\mathcal{S}=\cup_{n \geq 1} \mathcal{S}_{n}$.
- For any graph $\mathcal{X}$ let a weight $\lambda_{G}(w)$ be associated to each walk $w \in \mathcal{W}(G)$.
- Let the feature vector $\Phi(G)=\left(\Phi_{s}(G)\right)_{s \in \mathcal{S}}$ be defined by:

$$
\Phi_{s}(G)=\sum_{w \in \mathcal{W}(G)} \lambda_{G}(w) \mathbf{1}(s \text { is the label sequence of } w) .
$$

- A walk kernel is a graph kernel defined by:

$$
K_{\text {walk }}\left(G_{1}, G_{2}\right)=\sum_{s \in \mathcal{S}} \Phi_{s}\left(G_{1}\right) \Phi_{s}\left(G_{2}\right)
$$

## Walk kernel examples

## Examples

- The $n$ th-order walk kernel is the walk kernel with $\lambda_{G}(w)=1$ if the length of $w$ is $n, 0$ otherwise. It compares two graphs through their common walks of length $n$.


## - The random walk kernel is obtained with $\lambda_{G}(w)=P_{G}(w)$, where $P_{G}$ is a Markov random walk on $G$. In that case we have:

## $K\left(G_{1}, G_{2}\right)=P\left(\operatorname{label}\left(W_{1}\right)=\operatorname{label}\left(W_{2}\right)\right)$

where $W_{1}$ and $W_{2}$ are two independant random walks on $G_{1}$ and $G_{2}$, respectively (Kashima et al., 2003).

- The geometric walk kernel is obtained (when it converges) with
$\lambda_{G}(w)=\beta^{\operatorname{length}(w)}$, for $\beta>0$. In that case the feature space is of infinite dimension (Gärtner et al., 2003).


## Walk kernel examples

## Examples

- The $n$ th-order walk kernel is the walk kernel with $\lambda_{G}(w)=1$ if the length of $w$ is $n, 0$ otherwise. It compares two graphs through their common walks of length $n$.
- The random walk kernel is obtained with $\lambda_{G}(w)=P_{G}(w)$, where $P_{G}$ is a Markov random walk on $G$. In that case we have:

$$
K\left(G_{1}, G_{2}\right)=P\left(\text { label }\left(W_{1}\right)=\operatorname{label}\left(W_{2}\right)\right),
$$

where $W_{1}$ and $W_{2}$ are two independant random walks on $G_{1}$ and $G_{2}$, respectively (Kashima et al., 2003).


## Walk kernel examples

## Examples

- The $n$ th-order walk kernel is the walk kernel with $\lambda_{G}(w)=1$ if the length of $w$ is $n, 0$ otherwise. It compares two graphs through their common walks of length $n$.
- The random walk kernel is obtained with $\lambda_{G}(w)=P_{G}(w)$, where $P_{G}$ is a Markov random walk on $G$. In that case we have:

$$
K\left(G_{1}, G_{2}\right)=P\left(\text { label }\left(W_{1}\right)=\operatorname{label}\left(W_{2}\right)\right),
$$

where $W_{1}$ and $W_{2}$ are two independant random walks on $G_{1}$ and $G_{2}$, respectively (Kashima et al., 2003).

- The geometric walk kernel is obtained (when it converges) with $\lambda_{G}(w)=\beta^{\operatorname{length}(w)}$, for $\beta>0$. In that case the feature space is of infinite dimension (Gärtner et al., 2003).


## Computation of walk kernels

## Proposition

These three kernels (nth-order, random and geometric walk kernels) can be computed efficiently in polynomial time.

## Product graph

## Definition

Let $G_{1}=\left(V_{1}, E_{1}\right)$ and $G_{2}=\left(V_{2}, E_{2}\right)$ be two graphs with labeled vertices. The product graph $G=G_{1} \times G_{2}$ is the graph $G=(V, E)$ with:
(1) $V=\left\{\left(v_{1}, v_{2}\right) \in V_{1} \times V_{2}: v_{1}\right.$ and $v_{2}$ have the same label $\}$,
(2) $E=$

$$
\left\{\left(\left(v_{1}, v_{2}\right),\left(v_{1}^{\prime}, v_{2}^{\prime}\right)\right) \in V \times V:\left(v_{1}, v_{1}^{\prime}\right) \in E_{1} \text { and }\left(v_{2}, v_{2}^{\prime}\right) \in E_{2}\right\} .
$$



G1


G2


G1 $\times$ G2

## Walk kernel and product graph

## Lemma

There is a bijection between:
(1) The pairs of walks $w_{1} \in \mathcal{W}_{n}\left(G_{1}\right)$ and $w_{2} \in \mathcal{W}_{n}\left(G_{2}\right)$ with the same label sequences,
(2) The walks on the product graph $w \in \mathcal{W}_{n}\left(G_{1} \times G_{2}\right)$.

## Corolary



## Walk kernel and product graph

## Lemma

There is a bijection between:
(1) The pairs of walks $w_{1} \in \mathcal{W}_{n}\left(G_{1}\right)$ and $w_{2} \in \mathcal{W}_{n}\left(G_{2}\right)$ with the same label sequences,
(2) The walks on the product graph $w \in \mathcal{W}_{n}\left(G_{1} \times G_{2}\right)$.

## Corollary

$$
\begin{aligned}
K_{\text {walk }}\left(G_{1}, G_{2}\right) & =\sum_{s \in \mathcal{S}} \Phi_{s}\left(G_{1}\right) \Phi_{s}\left(G_{2}\right) \\
& =\sum_{\left(w_{1}, w_{2}\right) \in \mathcal{W}\left(G_{1}\right) \times \mathcal{W}\left(G_{1}\right)} \lambda_{G_{1}}\left(w_{1}\right) \lambda_{G_{2}}\left(w_{2}\right) \mathbf{1}\left(I\left(w_{1}\right)=I\left(w_{2}\right)\right) \\
& =\sum_{w \in \mathcal{W}\left(G_{1} \times G_{2}\right)} \lambda_{G_{1} \times G_{2}}(w) .
\end{aligned}
$$

## Computation of the nth-order walk kernel

- For the $n$ th-order walk kernel we have $\lambda_{G_{1} \times G_{2}}(w)=1$ if the length of $w$ is $n, 0$ otherwise.
- Therefore:

$$
K_{n t h-\operatorname{order}}\left(G_{1}, G_{2}\right)=\sum_{w \in \mathcal{W}_{n}\left(G_{1} \times G_{2}\right)} 1
$$

- Let $A$ be the adjacency matrix of $G_{1} \times G_{2}$. Then we get:

$$
K_{\text {nth-order }}\left(G_{1}, G_{2}\right)=\sum_{i, j}\left[A^{n}\right]_{i, j}=1^{\top} A^{n} 1
$$

- Computation in $O\left(n\left|G_{1}\right|\left|G_{2}\right| d_{1} d_{2}\right)$, where $d_{i}$ is the maximum degree of $G_{i}$.


## Computation of random and geometric walk kernels

- In both cases $\lambda_{G}(w)$ for a walk $w=v_{1} \ldots v_{n}$ can be decomposed as:

$$
\lambda_{G}\left(v_{1} \ldots v_{n}\right)=\lambda^{i}\left(v_{1}\right) \prod_{i=2}^{n} \lambda^{t}\left(v_{i-1}, v_{i}\right)
$$

- Let $\Lambda_{i}$ be the vector of $\lambda^{i}(v)$ and $\Lambda_{t}$ be the matrix of $\lambda^{t}\left(v, v^{\prime}\right)$ :

$$
\begin{aligned}
K_{\text {walk }}\left(G_{1}, G_{2}\right) & =\sum_{n=1}^{\infty} \sum_{w \in \mathcal{W}_{n}\left(G_{1} \times G_{2}\right)} \lambda^{i}\left(v_{1}\right) \prod_{i=2}^{n} \lambda^{t}\left(v_{i-1}, v_{i}\right) \\
& =\sum_{n=0}^{\infty} \Lambda_{i} \Lambda_{t}^{n} \mathbf{1} \\
& =\Lambda_{i}\left(I-\Lambda_{t}\right)^{-1} 1
\end{aligned}
$$

- Computation in $O\left(\left|G_{1}\right|^{3}\left|G_{2}\right|^{3}\right)$


## Extensions 1: label enrichment

## Atom relabebling with the Morgan index



- Compromise between fingerprints and structural keys features.
- Other relabeling schemes are possible (graph coloring).
- Faster computation with more labels (less matches implies a smaller product graph).


## Extension 2: Non-tottering walk kernel

## Tottering walks

A tottering walk is a walk $w=v_{1} \ldots v_{n}$ with $v_{i}=v_{i+2}$ for some $i$.


Tottering

- Tottering walks seem irrelevant for many applications
- Focusing on non-tottering walks is a way to get closer to the path kernel (e.g., equivalent on trees).


## Computation of the non-tottering walk kernel (Mahé et al., 2005)

- Second-order Markov random walk to prevent tottering walks
- Written as a first-order Markov random walk on an augmented graph
- Normal walk kernel on the augmented graph (which is always a directed graph).



## Extension 3: Subtree kernels



## Example: Tree-like fragments of molecules



## Computation of the subtree kernel

- Like the walk kernel, amounts to compute the (weighted) number of subtrees in the product graph.
- Recursion: if $\mathcal{T}(v, n)$ denotes the weighted number of subtrees of depth $n$ rooted at the vertex $v$, then:

$$
\mathcal{T}(v, n+1)=\sum_{R \subset \mathcal{N}(v)} \prod_{v^{\prime} \in R} \lambda_{t}\left(v, v^{\prime}\right) \mathcal{T}\left(v^{\prime}, n\right),
$$

where $\mathcal{N}(v)$ is the set of neighbors of $v$.

- Can be combined with the non-tottering graph transformation as preprocessing to obtain the non-tottering subtree kernel.


## Application in chemoinformatics (Mahé et al., 2004)

## MUTAG dataset

- aromatic/hetero-aromatic compounds
- high mutagenic activity /no mutagenic activity, assayed in Salmonella typhimurium.
- 188 compouunds: 125 + / 63 -


## Results

10-fold cross-validation accuracy

| Method | Accuracy |
| :--- | :---: |
| Progol1 | $81.4 \%$ |
| 2D kernel | $91.2 \%$ |

## 2D Subtree vs walk kernels



Screening of inhibitors for 60 cancer cell lines.

## Image classification (Harchaoui and Bach, 2007)

## COREL14 dataset

- 1400 natural images in 14 classes
- Compare kernel between histograms (H), walk kernel (W), subtree kernel (TW), weighted subtree kernel (wTW), and a combination (M).

Performance comparison on Corel14


## Summary: graph kernels

## What we saw

- Kernels do not allow to overcome the NP-hardness of subgraph patterns
- They allow to work with approximate subgraphs (walks, subtrees), in infinite dimension, thanks to the kernel trick
- However: using kernels makes it difficult to come back to patterns after the learning stage


## Outline

(1) Explicit computation of features: the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification

4 Conclusion

## Microarrays measure gene expression

Make cDNA reverse transcript
Label cDNAs with fluorescent dyes

Control Experimental


Principle of cDNA microarray assay for gene expression (after Gibson \& Muse 2002)

Red = "up-regulation"
Green = "down-regulation"
Black = constitutive expression


## Cancer classification from microarray data



## Gene networks



## Gene networks and expression data

## Motivation

- Basic biological functions usually involve the coordinated action of several proteins:
- Formation of protein complexes
- Activation of metabolic, signalling or regulatory pathways
- Many pathways and protein-protein interactions are already known
- Hypothesis: the weights of the classifier should be "coherent" with respect to this prior knowledge



## An idea


(1) Use the gene network to extract the "important information" in gene expression profiles by Fourier analysis on the graph
(2) Learn a linear classifier on the smooth components

## Graph Laplacian

## Definition

The Laplacian of the graph is the matrix $L=D-A$.

$$
L=D-A=\left(\begin{array}{ccccc}
1 & 0 & -1 & 0 & 0 \\
0 & 1 & -1 & 0 & 0 \\
-1 & -1 & 3 & -1 & 0 \\
0 & 0 & -1 & 2 & -1 \\
0 & 0 & 0 & 1 & 1
\end{array}\right)
$$

## Fourier basis

- $L$ is positive semidefinite
- The eigenvectors $e_{1}, \ldots, e_{n}$ of $L$ with eigenvalues
$0=\lambda_{1} \leq \ldots \leq \lambda_{n}$ form a basis called Fourier basis
- For any $f: V \rightarrow \mathbb{R}$, the Fourier transform of $f$ is the vector $\hat{f} \in \mathbb{R}^{n}$ defined by:

$$
\hat{f}_{i}=f^{\top} e_{i}, \quad i=1, \ldots, n
$$

- The inverse Fourier formula holds:

$$
f=\sum_{i=1}^{n} \hat{f}_{i} e_{i}
$$

## Fourier basis



## Fourier basis



## Smoothing operator

## Definition

- Let $\phi: \mathbb{R}^{+} \rightarrow \mathbb{R}^{+}$be non-increasing.
- A smoothing operator $S_{\phi}$ transform a function $f: V \rightarrow \mathbb{R}$ into a smoothed version:

$$
S_{\phi}(f)=\sum_{i=1}^{n} \hat{f}_{i} \phi\left(\lambda_{i}\right) e_{i}
$$

## Smoothing operators

## Examples

- Identity operator $\left(S_{\phi}(f)=f\right)$ :

$$
\phi(\lambda)=1, \quad \forall \lambda
$$

- Low-pass filter:
- Attenuation of high frequencies:


## Smoothing operators

## Examples

- Identity operator $\left(S_{\phi}(f)=f\right)$ :

$$
\phi(\lambda)=1, \quad \forall \lambda
$$

- Low-pass filter:

$$
\phi(\lambda)= \begin{cases}1 & \text { if } \lambda \leq \lambda^{*} \\ 0 & \text { otherwise }\end{cases}
$$

- Attenuation of high frequencies:


## Smoothing operators

## Examples

- Identity operator $\left(S_{\phi}(f)=f\right)$ :

$$
\phi(\lambda)=1, \quad \forall \lambda
$$

- Low-pass filter:

$$
\phi(\lambda)= \begin{cases}1 & \text { if } \lambda \leq \lambda^{*} \\ 0 & \text { otherwise }\end{cases}
$$

- Attenuation of high frequencies:

$$
\phi(\lambda)=\exp (-\beta \lambda)
$$

## Supervised classification and regression

## Working with smoothed profiles

- Classical methods for linear classification and regression with a ridge penalty solve:

$$
\min _{\beta \in \mathbb{R}^{p}} \frac{1}{n} \sum_{i=1}^{n} I\left(\beta^{\top} f_{i}, y_{i}\right)+\lambda \beta^{\top} \beta .
$$

- Applying these algorithms on the smooth profiles means solving:

$$
\min _{\beta \in \mathbb{R}^{p}} \frac{1}{n} \sum_{i=1}^{n} I\left(\beta^{\top} S_{\phi}\left(f_{i}\right), y_{i}\right)+\lambda \beta^{\top} \beta .
$$

## Link with shrinkage estimator

## Lemma

This is equivalent to:

$$
\min _{v \in \mathbb{R}^{p}} \frac{1}{n} \sum_{i=1}^{n} I\left(v^{\top} f_{i}, y_{i}\right)+\lambda \sum_{i=1}^{p} \frac{\hat{v}_{i}^{2}}{\phi\left(\lambda_{i}\right)},
$$

hence the linear classifier $v$ is smooth.

## Proof

- Let $v=\sum_{i=1}^{n} \phi\left(\lambda_{i}\right) e_{i} e_{i}^{\top} \beta$, then
- Then $\hat{v}_{i}=\phi\left(\lambda_{i}\right) \hat{\beta}_{i}$ and $\beta^{\top} \beta=\sum_{i=1}^{n} \frac{\hat{v}_{i}^{2}}{\phi\left(\lambda_{i}\right)^{2}}$


## Link with shrinkage estimator

## Lemma

This is equivalent to:

$$
\min _{v \in \mathbb{R}^{p}} \frac{1}{n} \sum_{i=1}^{n} I\left(v^{\top} f_{i}, y_{i}\right)+\lambda \sum_{i=1}^{p} \frac{\hat{v}_{i}^{2}}{\phi\left(\lambda_{i}\right)},
$$

hence the linear classifier $v$ is smooth.

## Proof

- Let $v=\sum_{i=1}^{n} \phi\left(\lambda_{i}\right) e_{i} \boldsymbol{e}_{i}^{\top} \beta$, then

$$
\beta^{\top} S_{\phi}\left(f_{i}\right)=\beta^{\top} \sum_{i=1}^{n} \hat{f}_{i} \phi\left(\lambda_{i}\right) e_{i}=f^{\top} v
$$

- Then $\hat{v}_{i}=\phi\left(\lambda_{i}\right) \hat{\beta}_{i}$ and $\beta^{\top} \beta=\sum_{i=1}^{n} \frac{\hat{v}_{i}^{2}}{\phi\left(\lambda_{i}\right)^{2}}$.


## Kernel methods

## Smoothing kernel

Kernel methods (SVM, kernel ridge regression..) only need the inner product between smooth profiles:

$$
\begin{aligned}
K(f, g) & =S_{\phi}(f)^{\top} S_{\phi}(g) \\
& =\sum_{i=1}^{n} \hat{f}_{i} \hat{g}_{i} \phi\left(\lambda_{i}\right)^{2} \\
& =f^{\top}\left(\sum_{i=1}^{n} \phi\left(\lambda_{i}\right)^{2} e_{i} e_{i}^{\top}\right) g \\
& =f^{\top} K_{\phi} g,
\end{aligned}
$$

with

$$
K_{\phi}=\sum_{i=1}^{n} \phi\left(\lambda_{i}\right)^{2} e_{i} e_{i}^{\top}
$$

## Examples

- For $\phi(\lambda)=\exp (-t \lambda)$, we recover the diffusion kernel:

$$
K_{\phi}=\exp _{M}(-2 t L) .
$$

- For $\phi(\lambda)=1 / \sqrt{1+\lambda}$, we obtain
and the penalization is:



## Examples

- For $\phi(\lambda)=\exp (-t \lambda)$, we recover the diffusion kernel:

$$
K_{\phi}=\exp _{M}(-2 t L)
$$

- For $\phi(\lambda)=1 / \sqrt{1+\lambda}$, we obtain

$$
K_{\phi}=(L+I)^{-1}
$$

and the penalization is:

$$
\sum_{i=1}^{n} \frac{\hat{v}_{i}^{2}}{\phi\left(\lambda_{i}\right)}=v^{\top}(L+I) v=\|v\|_{2}^{2}+\sum_{i \sim j}\left(v_{i}-v_{j}\right)^{2}
$$

## Data

## Expression

- Study the effect of low irradiation doses on the yeast
- 12 non irradiated vs 6 irradiated
- Which pathways are involved in the response at the transcriptomic level?


## Graph

- KEGG database of metabolic pathways
- Two genes are connected is they code for enzymes that catalyze successive reactions in a pathway (metabolic gene network).
- 737 genes, 4694 vertices.


## Classification performance



## Classifier



## Classifier


a)


## Summary

- With kernels we are able to soft constrain the shape of the classifier through regularization, e.g.:

$$
\min _{v \in \mathbb{R}^{p}} R_{e m p}(v)+\lambda \sum_{i=1}^{p} \frac{\hat{v}_{i}^{2}}{\phi\left(\lambda_{i}\right)},
$$

- This is related to priors in Bayesian learning
- The resulting classifier is interpretable, even without selection of a specific list of features.


## Outline

(1) Explicit computation of features: the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification

4 Conclusion

## Linear classifiers



## Training the model

- Minimize an empirical risk on the training samples:

$$
\min _{\beta \in \mathbb{R}^{p+1}} R_{\text {emp }}(\beta)=\frac{1}{n} \sum_{i=1}^{n} l\left(f_{\beta}\left(x_{i}\right), y_{i}\right)
$$

- ... subject to some constraint on $\beta$, e.g.:

$$
\Omega(\beta) \leq C .
$$

## Linear classifiers



## Training the model

- Minimize an empirical risk on the training samples:

$$
\min _{\beta \in \mathbb{R}^{p+1}} R_{e m p}(\beta)=\frac{1}{n} \sum_{i=1}^{n} l\left(f_{\beta}\left(x_{i}\right), y_{i}\right)
$$

- ... subject to some constraint on $\beta$, e.g.:

$$
\Omega(\beta) \leq C
$$

## Example : Norm Constraints

## The approach

A common method in statistics to learn with few samples in high dimension is to constrain the Euclidean norm of $\beta$

$$
\Omega_{\text {ridge }}(\beta)=\|\beta\|_{2}^{2}=\sum_{i=1}^{p} \beta_{i}^{2},
$$

(ridge regression, support vector machines, kernel methods...)

## Pros

- Good performance in classification


## Cons

- Limited interpretation (small weights)
- No prior biological knowledge


## Example : Feature Selection

## The approach

Constrain most weights to be 0, i.e., select a few genes whose expression are sufficient for classification.

$$
\Omega_{\text {Best subset selection }}(\beta)=\|\beta\|_{0}=\sum_{i=1}^{p} 1\left(\beta_{i}>0\right) .
$$

This is usually a NP-hard problem, many greedy variants have been proposed (filter methods, wrapper methods)

## Pros

- Good performance
- Biomarker selection
- Interpretability


## Cons

- NP-hard
- Gene selection not robust
- No use of prior knowledge


## Example : Sparsity inducing convex priors

## The approach

Constrain most weights to be 0 through a convex non-differentiable penalty:

$$
\Omega_{\mathrm{LASSO}}(\beta)=\|\beta\|_{1}=\sum_{i=1}^{p}\left|\beta_{i}\right| .
$$

- Several variants exist, e.g., elastic net penalty ( $\left.\|\beta\|_{1}+\|\beta\|_{2}\right), \ldots$ )


## Pros

- Good performance
- Biomarker selection
- Interpretability


## Cons

- Gene selection not robust
- No use of prior knowledge


## Why LASSO leads to sparse solutions

Geometric interpretation with $p=2$



## Efficienty computation of the regularization path

$$
\begin{equation*}
\min _{\beta \in \mathbb{R}^{p+1}} R^{n}\left(f_{\beta}\right)=\sum_{i=1}^{n}\left(f_{\beta}\left(\mathbf{x}_{i}\right)-\mathbf{y}_{i}\right)^{2}+\lambda \sum_{i=1}^{p}\left|\beta_{i}\right| \tag{2}
\end{equation*}
$$

- No explicit solution, but this is just a quadratic program.
- LARS (Efron et al., 2004) provides a fast algorithm to compute the solution for all $\lambda$ 's simultaneously (regularization path)



## Incorporating prior knowledge

## The idea

- If we have a specific prior knowledge about the "correct" weights, it can be included in $\Omega$ in the contraint:

Minimize $\boldsymbol{R}_{\text {emp }}(\beta)$ subject to $\Omega(\beta) \leq C$.

- If we design a convex function $\Omega$, then the algorithm boils down to a convex optimization problem (usually easy to solve).
- Similar to priors in Bayesian statistics


## Outline

(1) Explicit computation of features: the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification
(4) Conclusion


## Motivation



- Indexing by all subgraphs is appealing but intractable in practice (both explicitly and with the kernel trick)
- Can we work implicitly with this representation using sparse learning, e.g., LASSO regression or boosting?
- This may lead to both accurate predictive model and the identification of discriminative patterns.
- The iterations of LARS or boosting amount to an optimization problem over subgraphs, which may be solved efficiently using graph mining technique...


## Boosting over subgraph indexation (Kudo et al., 2004)

- Weak learner = decision stump indexed by subgraph H and $\alpha= \pm 1$ :

$$
h_{\alpha, H}(G)=\alpha \Phi_{H}(G)
$$

- Boosting: at each iteration, for a given distribution $d_{1}+\ldots+d_{n}=1$ over the training points $\left(G_{i}, y_{i}\right)$, select a weak learner (subgraph $\tilde{H}$ ) which maximizes the gain

$$
\operatorname{gain}(H, \alpha)=\sum_{i=1}^{n} y_{i} h_{\alpha, H}\left(G_{i}\right)
$$

- This can be done "efficiently" by branch-and-bound over a DFS code tree (Yan and Han, 2002).


## The DFS code tree



## Graph LASSO regularization path (Tsuda, 2007)





## Summary

- Sparse learning is practically feasible in the space of graphs indexed by all subgraphs
- Leads to subgraph selection
- Several extensions
- LASSO regularization path (Tsuda, 2007)
- gboost (Saigo et al., 2009)
- A beautiful and promising marriage between machine learning and data mining


## Outline

(1) Explicit computation of features: the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification

4. Conclusion

## Chromosomic aberrations in cancer

| 88 | 88 | 8月 |  |  | $8 \%$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | ${ }^{2}$ | 3 |  |  | 4 | 5 |
| 18 | $8 \times$ | ** | \% | \% 8 | 88 | 88 |
| - | 7 | 8 | 9 | ${ }^{10}$ | " | 12 |
| A ${ }^{\text {a }}$ | ** | * ${ }_{\text {a }}$ |  | "\% | ** | 4 |
| ${ }^{13}$ | 14 | 15 |  | 16 | 17 | 18 |
| ** | = |  |  | * | \% ${ }^{\text {x }}$ |  |
| 19 | ${ }^{20}$ |  |  | 22 | $\times$ |  |
| $\gamma$ | $)_{2}^{1}$ | 81 |  |  |  | $j_{5}^{\prime \prime}$ |
| 18 | ${ }_{\substack{881 \\ 7}}$ | 8! | 8 | 138 | $)_{11}^{8}$ | 18 12 |
| $101$ | $10$ | ${ }_{15}$ |  | 818 16 | $\begin{aligned} & 8 \\ & 17 \end{aligned}$ | 18 18 |
| 41 19 | 81 20 |  |  | ${ }_{22}$ | $\underset{\text { d }}{ }$ | $\stackrel{*}{*}^{*}$ |
| 19 | ${ }^{20}$ |  |  | 22 | $\times$ | \% |

## Comparative Genomic Hybridization (CGH)

## Motivation

- Comparative genomic hybridization (CGH) data measure the DNA copy number along the genome
- Very useful, in particular in cancer research
- Can we classify CGH arrays for diagnosis or prognosis purpose?



## Aggressive vs non-aggressive melanoma








## Classification of array CGH

## Prior knowledge

- Let $\mathbf{x}$ be a CGH profile
- We focus on linear classifiers, i.e., the sign of :

$$
f(\mathbf{x})=\mathbf{x}^{\top} \beta .
$$

- We expect $\beta$ to be
- sparse : only a few positions should be discriminative
- piecewise constant : within a region, all probes should contribute equally



## A penalty for CGH array classification

The fused LASSO penalty (Tibshirani et al., 2005)

$$
\Omega_{\text {fusedlasso }}(\beta)=\sum_{i}\left|\beta_{i}\right|+\sum_{i \sim j}\left|\beta_{i}-\beta_{j}\right| .
$$

- First term leads to sparse solutions
- Second term leads to piecewise constant solutions
- Combined with a hinge loss leads to a fused SVM (Rapaport et al., 2008);


## Application: metastasis prognosis in melanoma





## Outline

(1) Explicit computation of features: the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification
(4) Conclusion


## How to select jointly genes belonging to the same pathways?



## Selecting pre-defined groups of variables

## Group lasso (Yuan \& Lin, 2006)

If groups of covariates are likely to be selected together, the $\ell_{1} / \ell_{2}$-norm induces sparse solutions at the group level:

$$
\Omega_{\text {group }}(w)=\sum_{g}\left\|w_{g}\right\|_{2}
$$



## What if a gene belongs to several groups?

## Issue of using the group-lasso

- $\Omega_{\text {group }}(w)=\sum_{g}\left\|w_{g}\right\|_{2}$ sets groups to 0 .
- One variable is selected $\Leftrightarrow$ all the groups to which it belongs are selected.
Cell
cycle


## IGF,

G2

$$
\left\|w_{g_{1}}\right\|_{2}=\left\|w_{g_{3}}\right\|_{2}=0
$$

IGF selection $\Rightarrow$ selection of unwanted groups

Removal of any group
containing a gene $\Rightarrow$ the weight of the gene is 0 .

## Overlap norm (Jacob et al., 2009)

## An idea

Introduce latent variables $v_{g}$ :

$$
\left\{\begin{array}{l}
\min _{w, V} L(w)+\lambda \sum_{g \in \mathcal{G}}\left\|v_{g}\right\|_{2} \\
w=\sum_{g \in \mathcal{G}} v_{g} \\
\operatorname{supp}\left(v_{g}\right) \subseteq g
\end{array}\right.
$$



## Properties

- Resulting support is a union of groups in $\mathcal{G}$.
- Possible to select one variable without selecting all the groups containing it.
- Setting one $v_{g}$ to 0 doesn't necessarily set to 0 all its variables in $W$.


## A new norm

## Overlap norm

$$
\left\{\begin{array}{l}
\min _{w, v} L(w)+\lambda \sum_{g \in \mathcal{G}}\left\|v_{g}\right\|_{2} \\
w=\sum_{g \in \mathcal{G}} v_{g}=\min _{w} L(w)+\lambda \Omega_{\text {overlap }}(w) \\
\operatorname{supp}\left(v_{g}\right) \subseteq g
\end{array}\right.
$$

with

$$
\Omega_{\text {overlap }}(w) \triangleq\left\{\begin{array}{l}
\min _{v} \sum_{g \in \mathcal{G}}\left\|v_{g}\right\|_{2} \\
w=\sum_{g \in \mathcal{G}} v_{g} \\
\operatorname{supp}\left(v_{g}\right) \subseteq g
\end{array}\right.
$$

## Property

- $\Omega_{\text {overlap }}(w)$ is a norm of $w$.
- $\Omega_{\text {overlap }}($.$) associates to w$ a specific (not necessarily unique) decomposition $\left(v_{g}\right)_{g \in \mathcal{G}}$ which is the argmin of $(*)$.


## Overlap and group unity balls



Balls for $\Omega_{\text {group }}^{\mathcal{G}}(\cdot)$ (middle) and $\Omega_{\text {overlap }}^{\mathcal{G}}(\cdot)$ (right) for the groups $\mathcal{G}=\{\{1,2\},\{2,3\}\}$ where $w_{2}$ is represented as the vertical coordinate. Left: group-lasso ( $\mathcal{G}=\{\{1,2\},\{3\}\}$ ), for comparison.

## Theoretical results

Consistency in group support (Jacob et al., 2009)

- Let $\bar{w}$ be the true parameter vector.
- Assume that there exists a unique decomposition $\bar{v}_{g}$ such that $\bar{w}=\sum_{g} \bar{v}_{g}$ and $\Omega_{\text {overlap }}^{\mathcal{G}}(\bar{w})=\sum\left\|\bar{v}_{g}\right\|_{2}$.
- Consider the regularized empirical risk minimization problem $L(w)+\lambda \Omega_{\text {overlap }}^{\mathcal{G}}(w)$.


## Theoretical results

## Consistency in group support (Jacob et al., 2009)

- Let $\bar{w}$ be the true parameter vector.
- Assume that there exists a unique decomposition $\bar{v}_{g}$ such that $\bar{w}=\sum_{g} \bar{v}_{g}$ and $\Omega_{\text {overlap }}^{\mathcal{G}}(\bar{w})=\sum\left\|\bar{v}_{g}\right\|_{2}$.
- Consider the regularized empirical risk minimization problem $L(w)+\lambda \Omega_{\text {overlap }}^{\mathcal{G}}(w)$.
Then
- under appropriate mutual incoherence conditions on $X$,
- as $n \rightarrow \infty$,
- with very high probability,
the optimal solution $\hat{w}$ admits a unique decomposition $\left(\hat{v}_{g}\right)_{g \in \mathcal{G}}$ such that

$$
\left\{g \in \mathcal{G} \mid \hat{v}_{g} \neq 0\right\}=\left\{g \in \mathcal{G} \mid \bar{v}_{g} \neq 0\right\}
$$

## Experiments

## Synthetic data: overlapping groups

- 10 groups of 10 variables with 2 variables of overlap between two successive groups : $\{1, \ldots, 10\},\{9, \ldots, 18\}, \ldots,\{73, \ldots, 82\}$.
- Support: union of 4 th and 5 th groups.
- Learn from 100 training points.




Frequency of selection of each variable with the lasso (left) and $\Omega_{\text {overlap }}^{\mathcal{G}}$ (.) (middle), comparison of the RMSE of both methods (right).

## Extension : Graph lasso

## Graph lasso

- Consider groups that are subgraphs whose union would give such connected components (e.g., edges $E$ ).

- $\Omega_{\mathrm{graph}}(w)=\min _{v \in \mathcal{V}_{E}} \sum_{e \in E}\left\|v_{e}\right\| \quad$ s.t. $\sum_{e \in E} v_{e}=w, \operatorname{supp}\left(v_{e}\right)=e$.


## Graph lasso vs kernel on graph

- Graph lasso:

$$
\Omega_{\text {graph lasso }}(w)=\sum_{i \sim j} \sqrt{w_{i}^{2}+w_{j}^{2}} .
$$

constrains the sparsity, not the values

- Graph kernel

$$
\Omega_{\text {graph kernel }}(w)=\sum_{i \sim j}\left(w_{i}-w_{j}\right)^{2} .
$$

constrains the values (smoothness), not the sparsity

## Results

## Breast cancer data

- Gene expression data for 8, 141 genes in 295 breast cancer tumors.
- Canonical pathways from MSigDB containing 639 groups of genes, 637 of which involve genes from our study.

| METHOD | $\ell_{1}$ | $\Omega_{\text {OVERLAP }}^{\mathcal{G}}()$. |
| :--- | :---: | :---: |
| ERROR | $0.38 \pm 0.04$ | $0.36 \pm 0.03$ |
| $\sharp$ PATH. | $148,58,183$ | $6,5,78$ |
| PROP. PATH. | $0.32,0.14,0.41$ | $0.01,0.01,0.17$ |

- Graph on the genes.

| METHOD | $\ell_{1}$ | $\Omega_{\text {graph }}()$. |
| :--- | :---: | :---: |
| ERROR | $0.39 \pm 0.04$ | $0.36 \pm 0.01$ |
| Av. SIZE C.C. | $1.1,1,1.0$ | $1.3,1.4,1.2$ |

## Outline

(1) Explicit computation of features : the case of graph features
(2) Using kernels

- Introduction to kernels
- Graph kernels
- Kernels for gene expression data using gene networks
(3) Using sparsity-inducing shrinkage estimators
- Feature selection for all subgraph indexation
- Classification of array CGH data with piecewise-linear models
- Structured gene selection for microarray classification

4 Conclusion

## Conclusion

- Machine learning with complex and structured data becomes the rule
- We surveyed several ideas
- Feature construction
- Learning with kernels
- Learning with sparsity
- Performance and interpretability are both important
- Many promising bridges between machine learning and data mining!

