# RegML 2016 <br> Class 6 <br> Structured sparsity 

Lorenzo Rosasco
UNIGE-MIT-IIT

June 30, 2016

## Exploiting structure

Building blocks of a function can be more structure than single variables

## Sparsity



Variables divided in non-overlapping groups

## Group sparsity



- $f(x)=\sum_{j=1}^{d} w_{j} x_{j}$
- $w=(\underbrace{w_{1}, \ldots}_{w(1)}, \ldots, \underbrace{\left., \ldots, w_{d}\right)}_{w(G)})$
- each group $\mathcal{G}_{g}$ has size $\left|\mathcal{G}_{g}\right|$, so $w(g) \in \mathbb{R}^{\left|\mathcal{G}_{g}\right|}$


## Group sparsity regularization

Regularization exploiting structure

$$
R_{\text {group }}(w)=\sum_{g=1}^{G}\|w(g)\|=\sum_{g=1}^{G} \sqrt{\sum_{j=1}^{\left|\mathcal{G}_{g}\right|}(w(g))_{j}^{2}}
$$

## Group sparsity regularization

Regularization exploiting structure

$$
R_{\text {group }}(w)=\sum_{g=1}^{G}\|w(g)\|=\sum_{g=1}^{G} \sqrt{\sum_{j=1}^{\left|\mathcal{G}_{g}\right|}(w(g))_{j}^{2}}
$$

Compare to

$$
\sum_{g=1}^{G}\|w(g)\|^{2}=\sum_{g=1}^{G} \sum_{j=1}^{\left|\mathcal{G}_{g}\right|}(w(g))_{j}^{2}
$$

## Group sparsity regularization

Regularization exploiting structure

$$
R_{\text {group }}(w)=\sum_{g=1}^{G}\|w(g)\|=\sum_{g=1}^{G} \sqrt{\sum_{j=1}^{\left|\mathcal{G}_{g}\right|}(w(g))_{j}^{2}}
$$

Compare to

$$
\sum_{g=1}^{G}\|w(g)\|^{2}=\sum_{g=1}^{G} \sum_{j=1}^{\left|\mathcal{G}_{g}\right|}(w(g))_{j}^{2}
$$

or

$$
\sum_{g=1}^{G}\|w(g)\|^{2}=\sum_{g=1}^{G} \sum_{j=1}^{\left|\mathcal{G}_{g}\right|}\left|(w(g))_{j}\right|
$$

## $\ell_{1}-\ell_{2}$ norm

We take the $\ell_{2}$ norm of all the groups

$$
(\|w(1)\|, \ldots,\|w(G)\|)
$$

and then the $\ell_{1}$ norm of the above vector

$$
\sum_{g=1}^{G}\|w(g)\|
$$

## Groups lasso

$$
\min _{w} \frac{1}{n}\|\hat{X} w-\hat{y}\|^{2}+\lambda \sum_{g=1}^{G}\|w(g)\|
$$

- reduces to the Lasso if groups have cardinality one


## Computations

$$
\min _{w} \frac{1}{n}\|\hat{X} w-\hat{y}\|^{2}+\lambda \underbrace{\sum_{g=1}^{G}\|w(g)\|}_{\text {non differentiable }}
$$

Convex, non-smooth, but composite structure

$$
w_{t+1}=\operatorname{Prox}_{\gamma \lambda R_{\text {group }}}\left(w_{t}-\gamma \frac{2}{n} \hat{X}^{\top}\left(\hat{X} w_{t}-\hat{y}\right)\right)
$$

## Block thresholding

It can be shown that

$$
\operatorname{Prox}_{\lambda R_{\text {group }}}(w)=\left(\operatorname{Prox}_{\lambda\|\cdot\|}(w(1)), \ldots, \operatorname{Prox}_{\lambda\|\cdot\|}(w(G))\right.
$$

$\left(\operatorname{Prox}_{\lambda\|\cdot\|}(w(g))\right)^{j}= \begin{cases}w(g)^{j}-\lambda \frac{w(g)^{j}}{\|w(g)\|} & \|w(g)\|>\lambda \\ 0 & \|w(g)\| \leq \lambda\end{cases}$

- Entire groups of coefficients set to zero!
- Reduces to softhresholding if groups have cardinality one


## Other norms

$\ell_{1}-\ell_{p}$ norms

$$
R(w)=\sum_{g=1}^{G}\|w(g)\|_{p}=\sum_{g=1}^{G}\left(\sum_{j=1}^{\left|\mathcal{G}_{g}\right|}(w(g))_{j}^{p}\right)^{\frac{1}{p}}
$$

## Overlapping groups



Variables divided in possibly overlapping groups

## Regularization with overlapping groups



Group Lasso

$$
R_{\mathrm{GL}}(w)=\sum_{g=1}^{G}\|w(g)\|
$$

## Regularization with overlapping groups



Group Lasso

$$
R_{\mathrm{GL}}(w)=\sum_{g=1}^{G}\|w(g)\|
$$

$\rightarrow$ The selected variables are union of group complements

## Regularization with overlapping groups



Let $\bar{w}(g) \in \mathbb{R}^{d}$ be equal to $w(g)$ on group $\mathcal{G}_{g}$ and zero otherwise

## Regularization with overlapping groups



Let $\bar{w}(g) \in \mathbb{R}^{d}$ be equal to $w(g)$ on group $\mathcal{G}_{g}$ and zero otherwise Group Lasso with overlap

$$
R_{\mathrm{GLO}}(w)=\inf \left\{\sum_{g=1}^{G}\|w(g)\| \mid w(1), \ldots, w(g) \text { s.t. } w=\sum_{g=1}^{G} \bar{w}(g)\right\}
$$

## Regularization with overlapping groups



Let $\bar{w}(g) \in \mathbb{R}^{d}$ be equal to $w(g)$ on group $\mathcal{G}_{g}$ and zero otherwise Group Lasso with overlap

$$
R_{\mathrm{GLO}}(w)=\inf \left\{\sum_{g=1}^{G}\|w(g)\| \mid w(1), \ldots, w(g) \text { s.t. } w=\sum_{g=1}^{G} \bar{w}(g)\right\}
$$

- Multiple ways to write $w=\sum_{g=1}^{G} \bar{w}(g)$


## Regularization with overlapping groups



Let $\bar{w}(g) \in \mathbb{R}^{d}$ be equal to $w(g)$ on group $\mathcal{G}_{g}$ and zero otherwise Group Lasso with overlap

$$
R_{\mathrm{GLO}}(w)=\inf \left\{\sum_{g=1}^{G}\|w(g)\| \mid w(1), \ldots, w(g) \text { s.t. } w=\sum_{g=1}^{G} \bar{w}(g)\right\}
$$

- Multiple ways to write $w=\sum_{g=1}^{G} \bar{w}(g)$
- Selected variables are groups!


## An equivalence

It holds

$$
\min _{w} \frac{1}{n}\|\hat{X} w-\hat{y}\|^{2}+\lambda R_{\mathrm{GLO}}(w) \Leftrightarrow \min _{\tilde{w}} \frac{1}{n}\|\tilde{X} \tilde{w}-\hat{y}\|^{2}+\lambda \sum_{g=1}^{G}\|w(g)\|
$$

- $\tilde{X}$ is the matrix obtained by replicating columns/variables
- $\tilde{w}=(w(1), \ldots, w(G))$, vector with (nonoverlapping!) groups


## An equivalence (cont.)

Indeed

$$
\begin{array}{r}
\min _{w} \frac{1}{n}\|\hat{X} w-\hat{y}\|^{2}+\lambda \sum_{\substack{w(1), \ldots, w(g) \\
\text { s.t. } \sum_{g=1}^{G} \bar{w}(g)=w}} \sum_{g=1}^{G}\|w(g)\|= \\
\inf _{\substack{w(1), \ldots, w(g) \\
\text { s.t. } \sum_{g=1}^{G} \bar{w}(g)=w}} \frac{1}{n}\|\hat{X} w-\hat{y}\|^{2}+\lambda \sum_{g=1}^{G}\|w(g)\|= \\
\inf _{w(1), \ldots, w(g)} \frac{1}{n}\left\|\hat{X}\left(\sum_{g=1}^{G} \bar{w}(g)\right)-\hat{y}\right\|^{2}+\lambda \sum_{g=1}^{G}\|w(g)\|= \\
\inf _{w(1), \ldots, w(g)} \frac{1}{n}\left\|\sum_{g=1}^{G} \hat{X}_{\mid \mathcal{G}_{g}} w(g)-\hat{y}\right\|^{2}+\lambda \sum_{g=1}^{G}\|w(g)\|= \\
\min _{\tilde{w}} \frac{1}{n}\|\tilde{X} \tilde{w}-\hat{y}\|^{2}+\lambda \sum_{g=1}^{G}\|w(g)\|
\end{array}
$$

## Computations

- Can use block thresholding with replicated variables $\Longrightarrow$ potentially wasteful
- The proximal operator for $R_{\mathrm{GLO}}$ can be computed efficiently but not in closed form


## More structure

Structured overlapping groups

- trees
- DAG

Structure can be exploited in computations...

## Beyond linear models

Consider a dictionary made by union of distinct dictionaries

$$
f(x)=\sum_{g=1}^{G} \underbrace{f_{g}(x)}=\sum_{g=1}^{G} \Phi_{g}(x)^{\top} w(g),
$$

where each dictionary defines a feature map

$$
\Phi_{g}(x)=\left(\phi_{1}^{g}(x), \ldots, \phi_{p_{g}}^{g}(x)\right)
$$

Easy extension with usual change of variable...

## Representer theorems

Let

$$
f(x)=x^{\top}\left(\sum_{g=1}^{G} \bar{w}(g)\right)=\sum_{g=1}^{G} \bar{x}(g)^{\top} \bar{w}(g)=\sum_{g=1}^{G} f_{g}(x),
$$

## Representer theorems

Let

$$
f(x)=x^{\top}\left(\sum_{g=1}^{G} \bar{w}(g)\right)=\sum_{g=1}^{G} \bar{x}(g)^{\top} \bar{w}(g)=\sum_{g=1}^{G} f_{g}(x),
$$

Idea Show that

$$
\bar{w}(g)=\sum_{i=1}^{n} \bar{x}(g)_{i} c(g)_{i},
$$

i.e.

$$
f_{g}(x)=\sum_{i=1}^{n} \bar{x}(g)^{\top} \bar{x}(g)_{i} c(g)_{i}=\sum_{i=1}^{n} \underbrace{x(g)^{\top} x(g)_{i}}_{\Phi_{g}(x)^{\top} \Phi_{g}\left(x_{i}\right)=K_{g}\left(x, x_{i}\right)} c(g)_{i}
$$

## Representer theorems

Let

$$
f(x)=x^{\top}\left(\sum_{g=1}^{G} \bar{w}(g)\right)=\sum_{g=1}^{G} \bar{x}(g)^{\top} \bar{w}(g)=\sum_{g=1}^{G} f_{g}(x),
$$

Idea Show that

$$
\bar{w}(g)=\sum_{i=1}^{n} \bar{x}(g)_{i} c(g)_{i},
$$

i.e.

$$
f_{g}(x)=\sum_{i=1}^{n} \bar{x}(g)^{\top} \bar{x}(g)_{i} c(g)_{i}=\sum_{i=1}^{n} \underbrace{x(g)^{\top} x(g)_{i}}_{\Phi_{g}(x)^{\top} \Phi_{g}\left(x_{i}\right)=K_{g}\left(x, x_{i}\right)} c(g)_{i}
$$

Note that in this case

$$
\left\|f_{g}\right\|^{2}=\|w(g)\|^{2}=c(g)^{\top} \underbrace{\hat{X}(g) \hat{X}(g)^{\top}}_{\hat{K}(g)} c(g)
$$

## Coefficients update

$$
\left.c_{t+1}=\operatorname{Prox}_{\gamma \lambda R_{\text {group }}}\left(c_{t}-\gamma\left(\hat{K} c_{t}-\hat{y}\right)\right)\right)
$$

where $\hat{K}=(\hat{K}(1), \ldots, \hat{K}(G))$, and $c_{t}=\left(c_{t}(1), \ldots, c_{t}(G)\right)$

Block Thresholding It can be shown that
$\left(\operatorname{Prox}_{\lambda\|\cdot\|}(c(g))\right)^{j}= \begin{cases}c(g)^{j}-\lambda \underbrace{\frac{c(g)^{j}}{\sqrt{c(g)^{\top} \hat{K}(g) c(g)}}}_{\left\|f_{g}\right\|} & \left\|f_{g}\right\|>\lambda \\ 0 & \left\|f_{g}\right\| \leq \lambda\end{cases}$

## Non-parametric sparsity

$$
f(x)=\sum_{g=1}^{G} f_{g}(x)
$$

$$
f_{g}(x)=\sum_{i=1}^{n} x(g)^{\top} x(g)_{i}(c(g))_{i} \quad \mapsto \quad f_{g}(x)=\sum_{i=1}^{n} K_{g}\left(x, x_{i}\right)(c(g))_{i}
$$

$\left(K_{1}, \ldots, K_{G}\right)$ family of kernels

$$
\sum_{g=1}^{G}\|w(g)\| \Longrightarrow \sum_{g=1}^{G}\left\|f_{g}\right\|_{K_{g}}
$$

## $\ell_{1}$ MKL

$$
\begin{aligned}
& \inf _{\substack{w(1), \ldots, w(g) \\
\text { s.t. } \sum_{g=1}^{G} \bar{w}(g)=w}} \frac{1}{n}\|\hat{X} w-\hat{y}\|^{2}+\lambda \sum_{g=1}^{G}\|w(g)\|= \\
& \Downarrow \\
& \min _{\substack{ \\
f_{1}, \ldots, f_{g} \\
\text { s.t. } \sum_{g=1}^{G} f_{g}=f}} \frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-f\left(x_{i}\right)\right)^{2}+\lambda \sum_{g=1}^{G}\left\|f_{g}\right\|_{K_{g}}
\end{aligned}
$$

## $\ell_{2}$ MKL

$$
\sum_{g=1}^{G}\|w(g)\|^{2} \Longrightarrow \sum_{g=1}^{G}\left\|f_{g}\right\|_{K_{g}}^{2}
$$

Corresponds to using the kernel

$$
K\left(x, x^{\prime}\right)=\sum_{g=1}^{G} K_{g}\left(x, x^{\prime}\right)
$$

## $\ell_{1}$ or $\ell_{2} \mathrm{MKL}$

- $\ell_{2}$ *much* faster
- $\ell_{1}$ could be useful is only few kernels are relevant


## Why MKL?

- Data fusion- different features
- Model selection, e.g. gaussian kernels with different widths
- Richer model- many kernels!


## MKL \& kernel learning

It can be shown that

$$
\begin{aligned}
& \min _{\substack{ \\
f_{1}, \ldots, f_{g} \\
\text { s.t. } \sum_{g=1}^{G} \\
f_{g}}} \frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-f\left(x_{i}\right)\right)^{2}+\lambda \sum_{g=1}^{G}\left\|f_{g}\right\|_{K_{g}} \\
&
\end{aligned}
$$

$\Uparrow$

$$
\min _{K \in \mathcal{K}} \min _{f \in \mathcal{H} K} \frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-f\left(x_{i}\right)\right)^{2}+\lambda\|f\|_{K}^{2}
$$

where $\mathcal{K}=\left\{K \mid K=\sum_{g} K_{g} \alpha_{g}, \quad \alpha_{g} \geq 0\right\}$

## Sparsity beyond vectors

Recall multi-variable regression

$$
\begin{gathered}
\left(x_{i}, y_{i}\right)_{i=1^{n}}, \quad x_{i} \in \mathbb{R}^{d}, \quad y_{i} \in \mathbb{R}^{T} \\
f(x)=x^{\top} \underbrace{W}_{d \times T} \\
\min _{W}\|\hat{X} W-\hat{Y}\|_{F}^{2}+\lambda \operatorname{Tr}\left(W A W^{\top}\right)
\end{gathered}
$$

## Sparse regularization

- We have seen

$$
\operatorname{Tr}\left(W W^{\top}\right)=\sum_{j=1}^{d} \sum_{t=1}^{T}\left(W_{t, j}\right)^{2}
$$

- We could consider now

$$
\sum_{j=1}^{d} \sum_{t=1}^{T}\left|W_{t, j}\right|
$$

## Spectral Norms/p-Schatten norms

- We have seen

$$
\operatorname{Tr}\left(W W^{\top}\right)=\sum_{t=1}^{\min \{d, T\}} \sigma_{i}^{2}
$$

- We could consider now

$$
R(W)=\|W\|_{*}=\sum_{t=1}^{\min \{d, T\}} \sigma_{i}, \quad \text { nuclear norm }
$$

- or

$$
R(W)=\left(\sum_{t=1}^{\min \{d, T\}}\left(\sigma_{i}\right)^{p}\right)^{1 / p}, \quad \text { p-Schatten norm }
$$

## Nuclear norm regularization

$$
\min _{W}\|\hat{X} W-\hat{Y}\|_{F}^{2}+\lambda\|W\|_{*}
$$

## Computations

$$
W_{t+1}=\operatorname{Prox}_{\gamma \lambda\|\cdot\|_{*}}\left(W_{t}-2 \gamma \hat{X}^{\top}\left(\widehat{X} W_{t}-\widehat{Y}\right)\right)
$$

Let $W=U \Sigma V^{\top}, \Sigma=\operatorname{diag}\left(\sigma_{1}, \ldots, \sigma_{p}\right)$

$$
\operatorname{Prox}_{\|\cdot\|_{*}}(W)=U \operatorname{diag}\left(\operatorname{Prox}_{\|\cdot\|_{1}}\left(\sigma_{1}, \ldots, \sigma_{p}\right)\right) V^{\top}
$$

## This class

- Structured sparsity
- MKL
- Matrix sparsity


## Next class

## Data representation!

