# Regularized Least Squares AND <br> Support Vector Machines 

Francesca Odone and Lorenzo Rosasco

RegML 2014

## About this class

Goal To introduce two main examples of Tikhonov regularization algorithms, deriving and comparing their computational properties.

## BASICS: DATA

- Training set: $S=\left\{\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)\right\}$, $x_{i} \in \mathbb{R}^{d}, i=1, \ldots, n$
- Inputs: $X=\left\{x_{1}, \ldots, x_{n}\right\}$.
- Labels: $Y=\left\{y_{1}, \ldots, y_{n}\right\}$.


## Basics: RKHS, Kernel

- RKHS $\mathcal{H}$ with a positive semidefinite kernel function $K$ :

$$
\begin{aligned}
\text { linear: } & K\left(x_{i}, x_{j}\right)=x_{i}^{\top} x_{j} \\
\text { polynomial: } & K\left(x_{i}, x_{j}\right)=\left(x_{i}^{\top} x_{j}+1\right)^{d} \\
\text { gaussian: } & K\left(x_{i}, x_{j}\right)=\exp \left(-\frac{\left\|x_{i}-x_{j}\right\|^{2}}{\sigma^{2}}\right)
\end{aligned}
$$

- Define the kernel matrix $\mathbf{K}$ to satisfy $\mathbf{K}_{i j}=K\left(x_{i}, x_{j}\right)$.
- The kernel function with one argument fixed is $K_{x}=K(x, \cdot)$.
- Given an arbitrary input $X_{*}, \mathbf{K}_{x_{*}}$ is a vector whose $i$ th entry is $K\left(x_{i}, x_{*}\right)$.


## Tikhonov Regularization

We are interested into studying Tikhonov Regularization

$$
\underset{f \in \mathcal{H}}{\operatorname{argmin}}\left\{\sum_{i=1}^{n} V\left(y_{i}, f\left(x_{i}\right)\right)+\lambda\|f\|_{\mathcal{H}}^{2}\right\} .
$$

## Representer Theorem

The representer theorem guarantees that the solution can be written as

$$
f=\sum_{j=1}^{n} c_{j} K_{X_{j}}
$$

for some $c=\left(c_{1}, \ldots, c_{n}\right) \in \mathbb{R}^{n}$.
So $\mathbf{K} \boldsymbol{c}$ is a vector whose $i$ th element is $f\left(x_{i}\right)$ :

$$
f\left(x_{i}\right)=\sum_{j=1}^{n} c_{j} K_{x_{i}}\left(x_{j}\right)=\sum_{j=1}^{n} c_{j} \mathbf{K}_{i j}
$$

and $\|f\|_{\mathcal{H}}^{2}=c^{\top} \mathbf{K} c$.

## RKHS Norm and Representer Theorem

Since $f=\sum_{j=1}^{n} c_{j} K_{x_{j}}$, then

$$
\begin{aligned}
\|f\|_{\mathcal{H}}^{2} & =\langle f, f\rangle_{\mathcal{H}} \\
& =\left\langle\sum_{i=1}^{n} c_{i} K_{x_{i}}, \sum_{j=1}^{n} c_{j} K_{x_{j}}\right\rangle_{\mathcal{H}} \\
& =\sum_{i=1}^{n} \sum_{j=1}^{n} c_{i} c_{j}\left\langle K_{x_{i}}, K_{x_{j}}\right\rangle_{\mathcal{H}} \\
& =\sum_{i=1}^{n} \sum_{j=1}^{n} c_{i} c_{j} K\left(x_{i}, x_{j}\right)=c^{t} \mathrm{~K} c
\end{aligned}
$$

- RLS
- dual problem
- regularization path
- linear case
- SVM
- dual problem
- linear case
- historical derivation


## The RLS PROBLEM

Goal: Find the function $f \in \mathcal{H}$ that minimizes the weighted sum of the square loss and the RKHS norm

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

## RLS and Representer Theorem

Using the representer theorem the RLS problem is:

$$
\underset{c \in \mathbb{R}^{n}}{\operatorname{argmin}} \frac{1}{2 n}\|\mathbf{Y}-\mathbf{K} c\|_{2}^{2}+\frac{\lambda}{2} c^{\top} \mathbf{K} c
$$

The above functional is differentiable, we can find the minimum setting the gradient w.r.t $c$ to 0 :

## RLS and Representer Theorem

Using the representer theorem the RLS problem is:

$$
\underset{c \in \mathbb{R}^{n}}{\operatorname{argmin}} \frac{1}{2 n}\|\mathbf{Y}-\mathbf{K} c\|_{2}^{2}+\frac{\lambda}{2} c^{\top} \mathbf{K} c
$$

The above functional is differentiable, we can find the minimum setting the gradient w.r.t $c$ to 0 :

$$
\begin{aligned}
-\mathbf{K}(\mathbf{Y}-\mathbf{K} c)+\lambda n \mathbf{K} c & =0 \\
(\mathbf{K}+\lambda n I) c & =\mathbf{Y} \\
c & =(\mathbf{K}+\lambda n I)^{-1} \mathbf{Y}
\end{aligned}
$$

We find $c$ by solving a system of linear equations.

$$
(\mathbf{K}+\lambda n l) c=\mathbf{Y} .
$$

- The matrix $\mathbf{K}+\lambda n l$ is symmetric positive definite (with $\lambda>0$ ), so the appropriate algorithm is Cholesky factorization.
- In Matlab, the operator $\backslash$ seems to be using Cholesky, so you can just write $c=(\mathbf{K}+1$ ambda $\star \mathrm{n} * I) \backslash Y$;
- To be safe (or in Octave)

$$
R=\operatorname{chol}(\mathbf{K}+l \operatorname{ambda} * \mathrm{n} * I) ; c=\left(R \backslash\left(R^{\prime} \backslash Y\right)\right) ;
$$

The above algorithm has complexity $O\left(n^{3}\right)$.

## The RLS Solution, Comments

$$
\boldsymbol{c}=(\mathbf{K}+\lambda n l)^{-1} \mathbf{Y}
$$

The prediction at a new input $x_{*}$ is:

$$
\begin{aligned}
f\left(x_{*}\right) & =\sum_{j=1}^{n} c_{j} \mathbf{K}_{x_{j}}\left(x_{*}\right) \\
& =\mathbf{K}_{x_{*}} c \\
& =\mathbf{K}_{x_{*}} \mathbf{G}^{-1} \mathbf{Y}
\end{aligned}
$$

where $\mathbf{G}=\mathbf{K}+\lambda n l$.
Note that the above operation is $O\left(n^{2}\right)$.

## RLS Regularization Path

Typically we have to choose $\lambda$ and hence to compute the solutions corresponding to different values of $\lambda$.

- Is there a more efficent method than solving $c(\lambda)=(\mathbf{K}+\lambda n I)^{-1} \mathbf{Y}$ anew for each $\lambda$ ?


## RLS Regularization Path

Typically we have to choose $\lambda$ and hence to compute the solutions corresponding to different values of $\lambda$.

- Is there a more efficent method than solving $c(\lambda)=(\mathbf{K}+\lambda n I)^{-1} \mathbf{Y}$ anew for each $\lambda$ ?
- Form the eigendecomposition $\mathbf{K}=\mathbf{Q} \wedge \mathbf{Q}^{T}$, where $\wedge$ is diagonal with $\Lambda_{i i} \geq 0$ and $\mathbf{Q Q}^{T}=I$.
- Then

$$
\begin{aligned}
\mathbf{G} & =\mathbf{K}+\lambda n l \\
& =\mathbf{Q} \wedge \mathbf{Q}^{T}+\lambda n l \\
& =\mathbf{Q}(\Lambda+\lambda n I) \mathbf{Q}^{T},
\end{aligned}
$$

which implies that $\mathbf{G}^{-1}=\mathbf{Q}(\Lambda+\lambda n I)^{-1} \mathbf{Q}^{T}$.

## RLS Regularization Path Cont’d

- $O\left(n^{3}\right)$ time to solve one (dense) linear system, or to compute the eigendecomposition (constant is maybe $4 x$ worse). Given $\mathbf{Q}$ and $\wedge$, we can find $c(\lambda)$ in $O\left(n^{2}\right)$ time:

$$
c(\lambda)=\mathbf{Q}(\Lambda+\lambda n l)^{-1} \mathbf{Q}^{\top} \mathbf{Y}
$$

noting that $(\Lambda+\lambda n l)$ is diagonal.

- Finding $c(\lambda)$ for many $\lambda$ 's is (essentially) free!


## PARAMETER CHOICE

- idea: try different $\lambda$ and see which one performs best
- How to try them? A simple choice is to use a validation set of data
- If we have "enough" training data we may sample out a training and a validation set.
- Otherwise a common practice is K-fold Cross Validation (KCV):
(1) Divide data into $K$ sets of equal size: $S_{1}, \ldots, S_{k}$
(2) For each $i$ train on the other $K-1$ sets and test on the $i$ th set
- If $K=n$ we get the leave-one-out strategy (LOO)


## PARAMETER CHOICE

- Notice that some data should always be kept aside to be used as test set, to test the generalization performance of the system after parameter tuning took place


Entire set of data

## The Linear Case

- The linear kernel is $K\left(x_{i}, x_{j}\right)=x_{i}^{T} x_{j}$.
- The linear kernel offers many advantages for computation.
- Key idea: we get a decomposition of the kernel matrix for free: $\mathbf{K}=\mathbf{X X}{ }^{\top}$
- where $\mathbf{X}=\left[x_{1}^{\top}, \ldots, x_{n}^{\top}\right]$ is the data matrix $n \times d$
- In the linear case, we will see that we have two different computation options.


## LINEAR KERNEL, LINEAR FUNCTION

With a linear kernel, the function we are learning is linear as well:

$$
\begin{aligned}
f\left(x_{*}\right) & =\mathbf{K}_{x_{*}} c \\
& =x_{*}^{T} \mathbf{X}^{T} c \\
& =x_{*}^{T} w,
\end{aligned}
$$

where we define $w$ to be $\mathbf{X}^{T} c$.

## LINEAR KERNEL CONT.

For the linear kernel,

$$
\begin{aligned}
& \min _{c \in \mathbb{R}^{n}} \frac{1}{2 n}\|\mathbf{Y}-\mathbf{K} c\|_{2}^{2}+\frac{\lambda}{2} c^{T} \mathbf{K} c \\
= & \min _{c \in \mathbb{R}^{n}} \frac{1}{2 n}\left\|\mathbf{Y}-\mathbf{X X}^{T} c\right\|_{2}^{2}+\frac{\lambda}{2} c^{T} \mathbf{X} \mathbf{X}^{T} c \\
= & \min _{w \in \mathbb{R}^{d}} \frac{1}{2 n}\|\mathbf{Y}-\mathbf{X} w\|_{2}^{2}+\frac{\lambda}{2}\|w\|_{2}^{2} .
\end{aligned}
$$

Taking the gradient with respect to $w$ and setting it to zero

$$
\mathbf{X}^{T} \mathbf{X} w-\mathbf{X}^{T} \mathbf{Y}+\lambda n w=0
$$

we get

$$
w=\left(\mathbf{X}^{T} \mathbf{X}+\lambda n l\right)^{-1} \mathbf{X}^{T} \mathbf{Y}
$$

$$
w=\left(\mathbf{X}^{\top} \mathbf{X}+\lambda n l\right)^{-1} \mathbf{X}^{\top} \mathbf{Y} .
$$

Choleski decomposition allows us to solve the above problem in $O\left(d^{3}\right)$ for any fixed $\lambda$.

- We can work with the covariance matrix $\mathbf{X}^{T} \mathbf{X} \in \mathbb{R}^{d \times d}$.
- The algorithm is identical to solving a general RLS problem replacing the kernel matrix by $\mathbf{X}^{\top} \mathbf{X}$ and the labels vector by $\mathbf{X}^{\top} y$.
We can classify new points in $O(d)$ time, using $w$, rather than having to compute a weighted sum of $n$ kernel products (which will usually cost $O$ (nd) time).


## Regularization Path Via SVD

To compute solutions corresponding to multiple values of $\lambda$ we can again consider an eigendecomposition/svd.

- We need $O(n d)$ memory to store the data in the first place. The SVD also requires $O(n d)$ memory, and $O\left(n d^{2}\right)$ time. Compared to the nonlinear case, we have replaced an $O(n)$ with an $O(d)$, in both time and memory. If $n \gg d$, this can represent a huge savings.
- When can we solve one RLS problem? (I.e. what are the bottlenecks?)
- When can we solve one RLS problem? (I.e. what are the bottlenecks?)
- We need to form K, which takes $O\left(n^{2} d\right)$ time and $O\left(n^{2}\right)$ memory. We need to perform a Cholesky factorization or an eigendecomposition of $\mathbf{K}$, which takes $O\left(n^{3}\right)$ time.
- In the linear case we have replaced an $O(n)$ with an $O(d)$, in both time and memory. If $n \gg d$, this can represent a huge savings.
- Usually, we run out of memory before we run out of time.
- The practical limit on today's workstations is (more-or-less) 10,000 points (using Matlab).
- RLS
- dual problem
- regularization path
- linear case
- SVM
- dual problem
- linear case
- historical derivation


## The Hinge Loss

The support vector machine (SVM) for classification arises considering the hinge loss

$$
V(f(x), y) \equiv(1-y f(x))_{+},
$$

where $(s)_{+} \equiv \max (s, 0)$.


## SVM Standard Notation

With the hinge loss, our regularization problem becomes

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

With the hinge loss, our regularization problem becomes

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

In most of the SVM literature, the problem is written as

$$
\underset{t \in \mathcal{H}}{\operatorname{argmin}} C \sum_{i=1}^{n}\left(1-y_{i} f\left(x_{i}\right)\right)_{+}+\frac{1}{2}\|f\|_{\mathcal{H}}^{2} .
$$

The formulations are equivalent setting $C=\frac{1}{2 \lambda n}$.
This problem is non-differentiable (because of the "kink" in $V$ ).

## Slack Variables Formulation

We rewrite the functional using slack variables $\xi_{i}$.

$$
\begin{array}{rcl}
\underset{f \in \mathcal{H}}{\operatorname{argmin}} & C \sum_{i=1}^{n} \xi_{i}+\frac{1}{2}\|f\|_{\mathcal{H}}^{2} & \\
\text { subject to : } & \xi_{i} \geq 1-y_{i} f\left(x_{i}\right) & i=1, \ldots, n \\
& \xi_{i} \geq 0 & i=1, \ldots, n
\end{array}
$$

## Slack Variables Formulation

We rewrite the functional using slack variables $\xi_{i}$.

$$
\underset{f \in \mathcal{H}}{\operatorname{argmin}} C \sum_{i=1}^{n} \xi_{i}+\frac{1}{2}\|f\|_{\mathcal{H}}^{2}
$$

$$
\text { subject to : } \quad \xi_{i} \geq 1-y_{i} f\left(x_{i}\right) \quad i=1, \ldots, n
$$

$$
\xi_{i} \geq 0 \quad i=1, \ldots, n
$$

Applying the representer theorem we get a constrained quadratic programming problem:

$$
\begin{array}{rcl}
\underset{c \in \mathbb{R}^{n}, \xi \in \mathbb{R}^{n}}{\operatorname{argmin}} & C \sum_{i=1}^{n} \xi_{i}+\frac{1}{2} c^{T} \mathbf{K} c & \\
\text { subject to : } & \xi_{i} \geq 1-y_{i} \sum_{j=1}^{n} c_{j} K\left(x_{i}, x_{j}\right) & i=1, \ldots, n \\
\xi_{i} \geq 0 & i=1, \ldots, n
\end{array}
$$

## How to Solve?

$$
\begin{array}{ccl}
\underset{c \in \mathbb{R}^{n}, \xi \in \mathbb{R}^{n}}{\operatorname{argmin}} & C \sum_{i=1}^{n} \xi_{i}+\frac{1}{2} c^{\top} \mathbf{K} c & \\
\text { subject to : } & \xi_{i} \geq 1-y_{i}\left(\sum_{j=1}^{n} c_{j} K\left(x_{i}, x_{j}\right)\right) & i=1, \ldots, n \\
& \xi_{i} \geq 0 & i=1, \ldots, n
\end{array}
$$

- This is a constrained optimization problem. The general approach:
- Form the primal problem - we did this.
- Lagrangian from primal - just like Lagrange multipliers.
- Dual - one dual variable associated to each primal constraint in the Lagrangian.


## The Primal and Dual Problems

$$
\begin{array}{cc}
\operatorname{argmin}_{c \in \mathbb{R}^{n}, \xi \in \mathbb{R}^{n}}^{\operatorname{ar}} \quad C \sum_{i=1}^{n} \xi_{i}+\frac{1}{2} c^{T} \mathbf{K} c & \\
\text { subject to : } \xi_{i} \geq 1-y_{i}\left(\sum_{j=1}^{n} c_{j} K\left(x_{i}, x_{j}\right)\right) & i=1, \ldots, n \\
& \xi_{i} \geq 0 \\
& i=1, \ldots, n \\
\max _{\alpha \in \mathbb{R}^{n}} \sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \alpha^{T}(\operatorname{diag} \mathbf{Y}) \mathbf{K}(\operatorname{diag} \mathbf{Y}) \alpha \\
& 0 \leq \alpha_{i} \leq C
\end{array}
$$

The dual problem is easier to solve: simple box constraints.

- Basic idea: solve the dual problem to find the optimal $\alpha$ 's, and use them to find $c$

$$
c_{i}=\alpha_{i} y_{i}
$$

- The dual problem is easier to solve than the primal problem. It has simple box constraints and a single equality constraint, and the problem can be decomposed into a sequence of smaller problems.


## Optimality Conditions

All optimal solutions $(c, \xi)$ to the primal problem must satisfy the following conditions for some ( $\alpha, \zeta$ ):

$$
\begin{array}{rlrl}
\frac{\partial L}{\partial c_{i}}=\sum_{j=1}^{n} c_{j} K\left(x_{i}, x_{j}\right)-\sum_{j=1}^{n} y_{i} \alpha_{j} K\left(x_{i}, x_{j}\right) & =0 & & i=1, \ldots, n \\
\frac{\partial L}{\partial \xi_{i}}=C-\alpha_{i}-\zeta_{i} & =0 & & i=1, \ldots, n \\
y_{i}\left(\sum_{j=1}^{n} y_{j} \alpha_{j} K\left(x_{i}, x_{j}\right)\right)-1+\xi_{i} \geq 0 & & i=1, \ldots, n \\
\alpha_{i}\left[y_{i}\left(\sum_{j=1}^{n} y_{j} \alpha_{j} K\left(x_{i}, x_{j}\right)\right)-1+\xi_{i}\right] & =0 & & i=1, \ldots, n \\
\zeta_{i} \xi_{i} & =0 & & i=1, \ldots, n \\
\xi_{i}, \alpha_{i}, \zeta_{i} & \geq 0 & i=1, \ldots, n
\end{array}
$$

## Optimality Conditions

- They are also known as the Karush-Kuhn-Tucker (KKT) conditions.
- These optimality conditions are both necessary and sufficient for optimality: $(c, \xi, \alpha, \zeta)$ satisfy all of the conditions if and only if they are optimal for both the primal and the dual.


## Optimaliti conditions

## Interpreting the solution

Solution

$$
f(x)=\sum_{i=1}^{n} y_{i} \alpha_{i} K\left(x, x_{i}\right)
$$

From the KKT conditions we can derive the following:

$$
\begin{aligned}
\alpha_{i}=0 & \Longrightarrow y_{i} f\left(x_{i}\right) \geq 1 \\
0<\alpha_{i}<C & \Longrightarrow y_{i} f\left(x_{i}\right)=1 \\
\alpha_{i}=C & \Longrightarrow y_{i} f\left(x_{i}\right) \leq 1 \\
\alpha_{i}=0 & \Longleftrightarrow y_{i} f\left(x_{i}\right)>1 \\
\alpha_{i}=C & \Longleftrightarrow y_{i} f\left(x_{i}\right)<1
\end{aligned}
$$

## The Geometric Approach

- The "traditional" approach to describe SVM is to start with the concepts of separating hyperplanes and margin.
- The theory is usually developed in a linear space, beginning with the idea of a perceptron, a linear hyperplane that separates the positive and the negative examples.
- Defining the margin as the distance from the hyperplane to the nearest example, the basic observation is that intuitively, we expect a hyperplane with larger margin to generalize better than one with smaller margin.

Large and Small Margin Hyperplanes


(a)

(b)

## GEOMETRICAL MARGIN

## SEPARABLE CASE

For simplicity we consider the linear separable case


- Consider the decision surface

$$
D=\left\{x: w^{\top} x=0\right\}
$$

- Given a point $x_{i}$ its projection on the decision surface is $x_{i}^{\prime}=x_{i}-\beta \frac{w}{\|w\|}$.

$$
w^{\top} x_{i}-\beta \frac{w}{\|w\|}=0 \text { iff } \beta=y_{i} \frac{w^{\top}}{\|w\|} x
$$

$\beta$ is often called a geometrical margin which is scale invariant.

## MAXIMIZING THE MARGIN

## SEPARABLE CASE



- $\beta_{w}=\min _{i=1 \ldots n} \beta_{i}$
- 

$$
\begin{aligned}
\max _{w \in \mathbb{R}^{d}} & \beta_{w} \\
\text { subject to } & \beta_{w} \geq 0 \\
& \|w\|=1
\end{aligned}
$$

## MAXIMIZING THE MARGIN

## SEPARABLE CASE



- $\beta_{w}=\min _{i=1 \ldots n} \beta_{i}$
max
$w \in R^{d}$
subject to $\quad y_{i} \frac{w^{\top}}{\|w\|} x_{i} \geq \beta_{w}$

$$
\|w\|=1, \beta_{w} \geq 0
$$

## MAXIMIZING THE MARGIN

## SEPARABLE CASE

- we consider $\alpha=\beta_{w}\|\boldsymbol{w}\|$,
- because of the scale invariance we may set $\alpha=1$, thus we obtain

| $\max _{w \in R^{d}}$ | $\frac{1}{\\|w\\|}$ |
| ---: | :--- |
| subject to | $y_{i} w^{\top} x_{i} \geq 1$ |

- or equivalently

| $\min _{w \in R^{d}}$ | $\frac{1}{2}\\|w\\|^{2}$ |
| :---: | :--- |
| subject to | $y_{i} w^{\top} x_{i} \geq 1$ |

## MAXIMIZING THE MARGIN

Non-separable means there are points on the wrong side of the margin, i.e.

$$
\exists i \text { s.t. } y_{i} w^{\top} x_{i}<1
$$

We add slack variables to account for the wrongness:

$$
\begin{array}{rc}
\underset{\xi_{i}, w}{\operatorname{argmin}} & \sum_{i=1}^{n} \xi_{i}+\frac{1}{2}\|w\|^{2} \\
\text { s.t. } & y_{i} w^{\top} x_{i} \geq 1-\xi_{i}, \forall i
\end{array}
$$

## Geometric Interpretation of Reduced Optimality Conditions



## ADDING A BIAS TERM



- The original SVM formulation includes a bias term, so that $f(x)=w^{\top} x+b$
- This amounts at adding a further constraint $\sum_{i=1}^{n} y_{i} \alpha_{i} x_{i}=0$


## SVM - Summary

- The SVM is a Tikhonov regularization problem, with the hinge loss.
- Solving the SVM means solving a constrained quadratic program, roughly $O\left(n^{3}\right)$
- It's better to work with the dual program.
- Solutions can be sparse - few non-zero coefficients, this can have impact for memory and computational requirements.
- The non-zero coefficients correspond to points not classified correctly enough - a.k.a. "support vectors."


## Multi-OUTPUT

- In many practical problems, it is convenient to model the object of interest as a function with multiple outputs.
- In machine learning, this problem typically goes under the name of multi-output learning.
- A possible approach is to do re-write penalized empirical risk minimization

$$
\min _{f^{1}, \ldots, f^{T}} E R R\left[f^{1}, \ldots, f^{T}\right]+\lambda P E N\left(f^{1}, \ldots, f^{T}\right)
$$

Typically

- The error term is the sum of the empirical risks.
- The penalty term enforces similarity among the tasks.


## Multi-Class

## Multi-Class coding

A classical problem is multi-category classification where each input can be assigned to one of $T$ classes.

- We can consider $T$ labels $Y=\{1,2, \ldots T\}$ : this choice forces an unnatural ordering among classes


## Multi-CLASS

## MUlti-Class coding

A classical problem is multi-category classification where each input can be assigned to one of $T$ classes.

- We can consider $T$ labels $Y=\{1,2, \ldots T\}$ : this choice forces an unnatural ordering among classes
- We can define a coding, that is a one-to-one map $C: Y \rightarrow \mathcal{Y}$ where $\mathcal{Y}=\left(\ell_{1}, \ldots, \ell_{T}\right)$ are a set of coding vectors


## Multi-Class and Multi-Label

## MULTi-CLASS

In multi-category classification each input can be assigned to one of $T$ classes. We can think of encoding each class with a vector, for example: class one can be $(1,0 \ldots, 0)$, class 2 $(0,1 \ldots, 0)$ etc.

## Multi-class and Multi-Label

## MULTI-CLASS

In multi-category classification each input can be assigned to one of $T$ classes. We can think of encoding each class with a vector, for example: class one can be $(1,0 \ldots, 0)$, class 2 $(0,1 \ldots, 0)$ etc.

## Multilabel

Images contain at most $T$ objects each input image is associate to a vector

$$
(1,0,1 \ldots, 0)
$$

where $1 / 0$ indicate presence/absence of the an object.

## Multi-class RLS - ONE VS ALL

Consider the coding where class 1 is $(1,-1, \ldots,-1)$, class 2 is $(-1,1, \ldots,-1) \ldots$

## Multi-class RLS - one vs all

Consider the coding where class 1 is $(1,-1, \ldots,-1)$, class 2 is $(-1,1, \ldots,-1) \ldots$

One can easily check that the problem

$$
\min _{f_{1}, \ldots, f_{T}}\left\{\frac{1}{n} \sum_{j=1}^{T} \sum_{i=1}^{n}\left(y_{i}^{j}-f^{j}\left(x_{i}\right)\right)^{2}+\lambda \sum_{j=1}^{T}\left\|f^{j}\right\|_{K}^{2}\right.
$$

is exactly the one versus all scheme with regularized least squares.

## Multi-class RLS - SOlution

$$
(\mathbf{K}+\lambda n l) W=\mathbf{Y}
$$

with $W$ a $d \times T$ matrix and $\mathbf{Y}$ a $n \times T$ matrix whose $i$-th column contains 1 s if input belongs to class $i,-1$ otherwise.

The classification rule can be written as

$$
\begin{gathered}
c: X \rightarrow\{1, \ldots, T\} \\
c(x)=\arg \max _{t=1, \ldots, T} \sum_{j=1}^{n} W_{i}^{t} K\left(x, x_{i}\right)
\end{gathered}
$$

