



Statistical Learning in Complex Prediction Spaces: What Do We Know?

Shivani Agarwal

Department of Computer Science & Automation
Indian Institute of Science


January 2015






‘Simple’ prediction
problems


Example: Email Spam Filter



Example: Email Spam Filter




Example: Email Spam Filter




Example: Email Spam Filter




Example: Email Spam Filter




Example: Email Spam Filter




Example: Tumor Classification



Example: Tumor Classification



Example: Tumor Classification




Example: Tumor Classification


TUMOR
CLASSIFICATION
MODEL

Example: Tumor Classification




Example: Tumor Classification






Complex prediction
problems


Example: Speech Recognition



“How are
you?”




“I would like
to order
some pizza.”




“Today is a
gorgeous day
for playing
tennis.”

• • •


Example: Speech Recognition




Example: Speech Recognition




Example: Speech Recognition



Example: Speech Recognition



Example: Speech Recognition



Example: Image Segmentation



[Images taken from Li et al, ICME 2013]


Example: Image Segmentation



Learning
Algorithm

[Images taken from Li et al, ICME 2013]


Example: Image Segmentation



Example: Image Segmentation

IMAGE
SEGMENTATION
MODEL

Example: Image Segmentation




New image

Example: Image Segmentation




New image




→ Predicted
segmentation


Example: Document Ranking



Example: Document Ranking



Example: Document Ranking




Example: Document Ranking



DOCUMENT
RANKING
MODEL

Example: Document Ranking




Example: Document Ranking



Learning



Learning



Supervised Learning



$$S = ((x_1, y_1), \dots, (x_m, y_m))$$

$$\in (X \times Y)^m$$

Supervised Learning



$$S = ((x_1, y_1), \dots, (x_m, y_m))$$

$$\in (X \times Y)^m$$

Instance space

Supervised Learning



$$S = ((x_1, y_1), \dots, (x_m, y_m))$$

$$\in (X \times Y)^m$$

Instance space

Label space

Supervised Learning



$$S = ((x_1, y_1), \dots, (x_m, y_m))$$


$$h_S : X \rightarrow \hat{Y}$$

$$\in (X \times Y)^m$$

Instance space

Label space

Supervised Learning



Binary Classification



$$S = ((x_1, y_1), \dots, (x_m, y_m))$$

$$\in (X \times Y)^m$$

$$h_S : X \rightarrow \hat{Y}$$

$$Y = \hat{Y} = \{\pm 1\}$$

Prediction in Complex Spaces




$$S = ((x_1, y_1), \dots, (x_m, y_m))$$

$$\in (X \times Y)^m$$

$$h_S : X \rightarrow \hat{Y}$$

\hat{Y} = sets of complex structures
(e.g. sequences, graphs, permutations)

Prediction in Complex Spaces



Loss Function

Performance in supervised learning is often measured via a (label-dependent) **loss function**:

$$\ell : Y \times \hat{Y} \rightarrow R_+$$

Loss Function

Performance in supervised learning is often measured via a (label-dependent) **loss function**:

$$\ell : Y \times \hat{Y} \rightarrow R_+$$

$\ell(y, \hat{y})$ = 'loss' incurred on predicting \hat{y}
when true label is y

Loss Matrix

$$|Y| = n, \quad |\hat{Y}| = k$$

$$\mathbf{L} = \begin{bmatrix} & 1 & 2 & \dots & k \\ 1 & \ell(1, 1) & \ell(1, 2) & \dots & \ell(1, k) \\ 2 & \ell(2, 1) & \ell(2, 2) & \dots & \ell(2, k) \\ \vdots & \vdots & \ddots & & \vdots \\ n & \ell(n, 1) & \ell(n, 2) & \dots & \ell(n, k) \end{bmatrix}$$

Example: Binary 0-1 Classification

$$Y = \hat{Y} = \{\pm 1\}$$

$$n = k = 2$$

$$\mathbf{L}^{0-1} = \begin{bmatrix} -1 & +1 \\ 0 & 1 \\ +1 & 0 \end{bmatrix}$$

Example: Multiclass 0-1 Classification

$$Y = \hat{Y} = [n]$$

$$n = k > 2$$

$$\mathbf{L}^{0-1} = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & \dots & n \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ \vdots \\ n \end{matrix} & \left[\begin{matrix} 0 & 1 & 1 & \dots & 1 \\ 1 & 0 & 1 & \dots & 1 \\ 1 & 1 & 0 & \dots & 1 \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & 1 & 1 & \dots & 0 \end{matrix} \right] \end{matrix}$$

Example: Multiclass 0-1 Classification

$$Y = \hat{Y} = [n]$$

$$n = k > 2$$

n = 5

$$\mathbf{L}^{0-1} = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{matrix} & \left[\begin{matrix} 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 & 0 \end{matrix} \right] \end{matrix}$$

Example: Sequence Prediction with Hamming Loss

$$Y = \hat{Y} = \{0,1\}^r$$

$$n = k = 2^r$$

$$r = 3$$

$$\mathbf{L}^{\text{Ham}} = \begin{matrix} & \begin{matrix} 000 & 001 & 010 & 011 & 100 & 101 & 110 & 111 \end{matrix} \\ \begin{matrix} 000 \\ 001 \\ 010 \\ 011 \\ 100 \\ 101 \\ 110 \\ 111 \end{matrix} & \left[\begin{matrix} 0 & 1 & 1 & 2 & 1 & 2 & 2 & 3 \\ 1 & 0 & 2 & 1 & 2 & 1 & 3 & 2 \\ 1 & 2 & 0 & 1 & 2 & 3 & 1 & 2 \\ 2 & 1 & 1 & 0 & 3 & 2 & 2 & 1 \\ 1 & 2 & 2 & 3 & 0 & 1 & 1 & 2 \\ 2 & 1 & 3 & 2 & 1 & 0 & 2 & 1 \\ 2 & 3 & 1 & 2 & 1 & 2 & 0 & 1 \\ 3 & 2 & 2 & 1 & 2 & 1 & 1 & 0 \end{matrix} \right] \end{matrix}$$

Example: Document Ranking with Pairwise Disagreement Loss

$$Y = \{0,1\}^r, \hat{Y} = S_r$$

$$n = 2^r, k = r!$$

$$r = 3$$

$$\begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} \begin{pmatrix} 1 \\ 3 \\ 2 \end{pmatrix} \begin{pmatrix} 2 \\ 3 \\ 1 \end{pmatrix} \begin{pmatrix} 2 \\ 1 \\ 3 \end{pmatrix} \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix} \begin{pmatrix} 3 \\ 2 \\ 1 \end{pmatrix}$$

$$\mathbf{L}^{\text{PD}} = \begin{array}{c} \text{A 3x6 matrix of binary values (0s and 1s)} \\ \text{with the following pattern:} \\ \begin{array}{cccccc} \text{Row 1:} & \text{X X X} & \text{X X X} & 0 & 0 & 0 \\ \text{Row 2:} & \text{X X X} & \text{X X X} & 2 & 1 & 1 \\ \text{Row 3:} & \text{X X X} & \text{X X X} & 1 & 2 & 2 \\ \text{Row 4:} & \text{X X X} & \text{X X X} & 2 & 2 & 1 \\ \text{Row 5:} & \text{X X X} & \text{X X X} & 0 & 0 & 2 \\ \text{Row 6:} & \text{X X X} & \text{X X X} & 1 & 0 & 1 \\ \text{Row 7:} & \text{X X X} & \text{X X X} & 0 & 1 & 2 \\ \text{Row 8:} & \text{X X X} & \text{X X X} & 0 & 0 & 0 \end{array} \end{array} \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 2 & 1 & 1 & 2 & 0 & 0 \\ 1 & 2 & 0 & 0 & 2 & 1 \\ 2 & 2 & 0 & 1 & 1 & 0 \\ 0 & 0 & 2 & 1 & 1 & 2 \\ 1 & 0 & 2 & 2 & 0 & 1 \\ 0 & 1 & 1 & 0 & 2 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Formal Setup

Instance space X

Label space $Y = \{1, \dots, n\} = [n]$

Prediction space $\hat{Y} = \{1, \dots, k\} = [k]$

Loss matrix $\mathbf{L} \in R_+^{n \times k}$

Formal Setup

Instance space X

Label space $Y = \{1, \dots, n\} = [n]$

Prediction space $\hat{Y} = \{1, \dots, k\} = [k]$

Loss matrix $\mathbf{L} \in R_+^{n \times k}$

Goal: Given training sample

$S = ((x_1, y_1), \dots, (x_m, y_m)) \in (X \times [n])^m,$

learn prediction model $h_S : X \rightarrow [k]$

What is a Good Prediction Model

$$h : X \rightarrow [k]?$$

Should minimize target loss **on new instances**

What is a Good Prediction Model

$$h : X \rightarrow [k]?$$

Should minimize target loss **on new instances**

Assume all instance-label pairs drawn i.i.d. from a probability distribution D on $X \times [n]$.

What is a Good Prediction Model $h : X \rightarrow [k]$?

Should minimize target loss **on new instances**

Assume all instance-label pairs drawn i.i.d. from a probability distribution D on $X \times [n]$.

$$\text{er}_D^L[h] = E_{(x,y) \sim D}[\ell(y, h(x))]$$

Generalization L-error (or L-risk) of h (w.r.t. D)

Bayes Error and Regret

Bayes L-error for D :

$$\text{er}_D^{\mathbf{L},*} = \inf_{h:X \rightarrow [k]} \text{er}_D^{\mathbf{L}}[h]$$

Bayes Error and Regret

Bayes L-error for D :

$$\text{er}_D^{\mathbf{L},*} = \inf_{h:X \rightarrow [k]} \text{er}_D^{\mathbf{L}}[h]$$

L-regret of h w.r.t. D :


$$\text{regret}_D^{\mathbf{L}}[h] = \text{er}_D^{\mathbf{L}}[h] - \text{er}_D^{\mathbf{L},*}$$

What is a Good Learning Algorithm?




What is a Good Learning Algorithm?



Data Space



Model Space




What is a Good Learning Algorithm?




What is a Good Learning Algorithm?



Data Space



Model Space



What is a Good Learning Algorithm?



Statistical Consistency



Bayes L-consistent w.r.t. D if for $S \sim D^m$:


$$\text{regret}_D^L[h_S] \xrightarrow{P} 0 \quad \text{as} \quad m \rightarrow \infty.$$

Statistical Consistency




Universally Bayes L-consistent if

Bayes L-consistent w.r.t. *all* distributions D




**How do we design a statistically
consistent learning algorithm for a
given loss matrix L ?**

Road Map




Road Map




Empirical Risk Minimization (ERM)

Let H be some class of functions from X to $\{\pm 1\}$.




Empirical Risk Minimization (ERM)

Let H be some class of functions from X to $\{\pm 1\}$.



Empirical Risk Minimization (ERM)


Let H be some class of functions from X to $\{\pm 1\}$.



- ✓ For suitable H , universally 0-1 consistent in H ;
suitable extensions can be made universally Bayes 0-1 consistent


Empirical Risk Minimization (ERM)

Let H be some class of functions from X to $\{\pm 1\}$.




- ✓ For suitable H , universally 0-1 consistent in H ;
suitable extensions can be made universally Bayes 0-1 consistent
- ✗ Computationally hard!

Surrogate Risk Minimization




Surrogate Risk Minimization

$$\psi : \{\pm 1\} \times R \rightarrow R_+$$



Surrogate Risk Minimization

$$\psi : \{\pm 1\} \times \mathbb{R} \rightarrow \mathbb{R}_+$$



Surrogate Risk Minimization


$$\min_f \sum_{i=1}^m \psi(y_i, f(x_i))$$

Functions mapping
 X to R

Surrogate optimization
problem (convex for
suitable surrogate loss)

Surrogate Risk Minimization


Map back (continuous)
surrogate predictions to
(discrete) target
prediction space



Surrogate Risk Minimization

Let $\psi : \{\pm 1\} \times R \rightarrow R_+$.


Let F be some class of functions from X to R .



Surrogate Risk Minimization

Let $\psi : \{\pm 1\} \times R \rightarrow R_+$.


Let F be some class of functions from X to R .



Surrogate Risk Minimization

Let $\psi : \{\pm 1\} \times R \rightarrow R_+$.


Let F be some class of functions from X to R .



Surrogate Risk Minimization


Let $\psi : \{\pm 1\} \times R \rightarrow R_+$.

Let F be some class of functions from X to R .




✓ For convex ψ and suitable F , computationally efficient!


Example: Support Vector Machines



Example: Logistic Regression




Example: AdaBoost



Surrogate Risk Minimization

Let $\psi : \{\pm 1\} \times R \rightarrow R_+$.

Let F be some class of functions from X to R .




✓ For convex ψ and suitable F , computationally efficient!

Surrogate Risk Minimization

Let $\psi : \{\pm 1\} \times R \rightarrow R_+$.

Let F be some class of functions from X to R .



- ✓ For convex ψ and suitable F , computationally efficient!
- ✗ For suitable F , universally ψ -consistent in F ;
suitable extensions can be made universally Bayes ψ -consistent

Classification-Calibrated Surrogates

Theorem. If ψ is ‘classification-calibrated’, then

Bayes ψ -consistency \implies Bayes 0-1 consistency
(after applying **sign**)

[Bartlett et al., 2006]

Classification-Calibrated Surrogates

Theorem. If ψ is ‘classification-calibrated’, then

Bayes ψ -consistency \implies Bayes 0-1 consistency
(after applying **sign**)

✓ Hinge

[Bartlett et al., 2006]

Classification-Calibrated Surrogates

Theorem. If ψ is ‘classification-calibrated’, then

Bayes ψ -consistency \implies Bayes 0-1 consistency
(after applying **sign**)

✓ Hinge

✓ Logistic

[Bartlett et al., 2006]

Classification-Calibrated Surrogates

Theorem. If ψ is ‘classification-calibrated’, then


Bayes ψ -consistency \implies Bayes 0-1 consistency
(after applying **sign**)

✓ Hinge

✓ Logistic


✓ Exponential

Road Map




Empirical Risk Minimization (ERM)

Let H be some class of functions from X to $[k]$.




Empirical Risk Minimization (ERM)

Let H be some class of functions from X to $[k]$.




Empirical Risk Minimization (ERM)

Let H be some class of functions from X to $[k]$.




- ✓ For suitable H , universally **L**-consistent in H ;
suitable extensions can be made universally Bayes **L**-consistent
- ✗ Computationally hard!

Surrogate Risk Minimization




Surrogate Risk Minimization

$$\psi : [n] \times \mathbb{R}^d \rightarrow \mathbb{R}_+$$



Surrogate Risk Minimization

$$\psi : [n] \times \mathbb{R}^d \rightarrow \mathbb{R}_+$$




Surrogate Risk Minimization

$$\min_{\mathbf{f}} \sum_{i=1}^m \psi(y_i, \mathbf{f}(x_i))$$

Functions mapping
 X to R^d

Surrogate optimization
problem (convex for
suitable surrogate loss)


Surrogate Risk Minimization



Surrogate Risk Minimization

Let $\psi : [n] \times \mathbb{R}^d \rightarrow \mathbb{R}_+$, $\text{pred} : \mathbb{R}^d \rightarrow [k]$.


Let F be some class of functions from X to \mathbb{R}^d .



Surrogate Risk Minimization

Let $\psi : [n] \times \mathbb{R}^d \rightarrow \mathbb{R}_+$, $\text{pred} : \mathbb{R}^d \rightarrow [k]$.


Let F be some class of functions from X to \mathbb{R}^d .



Surrogate Risk Minimization

Let $\psi : [n] \times \mathbb{R}^d \rightarrow \mathbb{R}_+$, $\text{pred} : \mathbb{R}^d \rightarrow [k]$.


Let F be some class of functions from X to \mathbb{R}^d .



Surrogate Risk Minimization

Let $\psi : [n] \times \mathbb{R}^d \rightarrow \mathbb{R}_+$, $\text{pred} : \mathbb{R}^d \rightarrow [k]$.

Let F be some class of functions from X to \mathbb{R}^d .



- ✓ For convex ψ and suitable F , computationally efficient!
- ✗ For suitable F , universally ψ -consistent in F ;
suitable extensions can be made universally Bayes ψ -consistent

L-Calibrated Surrogates

Theorem. If ψ is ‘L-calibrated’, then

Bayes ψ -consistency \Rightarrow Bayes L-consistency
(after applying some **pred**)

[Zhang, 2004; Tewari & Bartlett, 2007; Ramaswamy & Agarwal, 2012]



**How do we design convex
calibrated surrogates for a given
loss matrix L ?**

Recent Work on Convex Calibrated Surrogates for Specific Target Losses L

Multiclass 0-1 Loss

Zhang, 2004; Tewari & Bartlett, 2007

Various Document (Subset) Ranking Losses

Cossack & Zhang, 2008; Xia et al, 2008; Duchi et al, 2010;
Ravikumar et al, 2011; Buffoni et al, 2011; Lan et al, 2012;
Calauzenes et al, 2012


Multilabel Losses

Gao & Zhou, 2011; Dembczynski et al, 2011

Our Work

Unified Approach for Designing
Convex Calibrated Surrogates for
General Loss Matrices

Our Work



When is a surrogate
calibrated for \mathbf{L} ?

such for Designing
Convex Calibrated Surrogates for
General Loss Matrices

Our Work

When is a surrogate
calibrated for \mathbf{L} ?

What is the smallest
dimension d that supports
a convex calibrated
surrogate for \mathbf{L} ?

Our Work

When is a surrogate
calibrated for \mathbf{L} ?

What is the smallest
dimension d that supports
a convex calibrated
surrogate for \mathbf{L} ?

Can we design explicit
low-dimensional
surrogates for \mathbf{L} ?


Our Work

When is a surrogate
calibrated for L ?


What is the smallest
dimension d that supports
a convex calibrated
surrogate for L ?

Can we design explicit
low-dimensional
surrogates for L ?

Trigger Probability Sets of Loss L



Positive Normal Sets of Surrogate ψ




Necessary Condition for Calibration

Theorem. If ψ is \mathbf{L} -calibrated, then every positive normal set of ψ must be contained in some trigger probability set of \mathbf{L} .

[Ramaswamy & Agarwal, 2012; 2014]



Necessary Condition for Calibration

Proof by picture.





[Ramaswamy & Agarwal, 2012; 2014]

Example



[Ramaswamy & Agarwal, 2012; 2014]

Example





[Ramaswamy & Agarwal, 2012; 2014]

Sufficient Condition for Calibration

Theorem. If there is a finite collection of positive normal sets of ψ that are each contained in some trigger probability set of \mathbf{L} and that collectively cover the simplex, then ψ is \mathbf{L} -calibrated.



[Ramaswamy & Agarwal, 2012; 2014]

Example



[Ramaswamy & Agarwal, 2012; 2014]

Example



[Ramaswamy & Agarwal, 2012; 2014]

Our Work


When is a surrogate
calibrated for L ?

What is the smallest
dimension d that supports
a convex calibrated
surrogate for L ?

Can we design explicit
low-dimensional
surrogates for L ?

Convex Calibration Dimension


$\text{CCdim}(\mathbf{L}) = \text{smallest dimension } d$
that supports a convex
calibrated surrogate for \mathbf{L}



[Ramaswamy & Agarwal, 2012; 2014]

Convex Calibration Dimension

$\text{CCdim}(\mathbf{L}) = \text{smallest dimension } d$
that supports a convex
calibrated surrogate for \mathbf{L}



$\text{CCdim}(\mathbf{L}) \leq n-1$

[Ramaswamy & Agarwal, 2012; 2014]

Upper Bound on Convex Calibration Dimension

Theorem.

$$\text{CCdim}(\mathbf{L}) \leq \text{rank}(\mathbf{L})$$

[Ramaswamy & Agarwal, 2012; 2014]

Lower Bound on Convex Calibration Dimension

Theorem. For losses \mathbf{L} whose columns can be obtained from one another by permuting entries,

$$\text{CCdim}(\mathbf{L}) \geq \text{rank}(\mathbf{L}) - 2$$

[Ramaswamy & Agarwal, 2012; 2014]

Example: Multiclass 0-1 Classification

$$Y = \hat{Y} = [n]$$

$$n = k > 2$$

$$\mathbf{L}^{0-1} = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & \dots & n \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ \vdots \\ n \end{matrix} & \left[\begin{matrix} 0 & 1 & 1 & \dots & 1 \\ 1 & 0 & 1 & \dots & 1 \\ 1 & 1 & 0 & \dots & 1 \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & 1 & 1 & \dots & 0 \end{matrix} \right] \end{matrix}$$

Example: Multiclass 0-1 Classification

$$Y = \hat{Y} = [n]$$

$$n = k > 2$$

$$\text{rank}(\mathbf{L}^{0-1}) = n$$

$$\begin{bmatrix} n & 1 & 1 & 1 & \dots & 0 \end{bmatrix}$$

Example: Sequence Prediction with Hamming Loss

$$Y = \hat{Y} = \{0,1\}^r$$

$$n = k = 2^r$$

$r = 3$

$$\mathbf{L}^{\text{Ham}} = \begin{matrix} & \begin{matrix} 000 & 001 & 010 & 011 & 100 & 101 & 110 & 111 \end{matrix} \\ \begin{matrix} 000 \\ 001 \\ 010 \\ 011 \\ 100 \\ 101 \\ 110 \\ 111 \end{matrix} & \left[\begin{matrix} 0 & 1 & 1 & 2 & 1 & 2 & 2 & 3 \\ 1 & 0 & 2 & 1 & 2 & 1 & 3 & 2 \\ 1 & 2 & 0 & 1 & 2 & 3 & 1 & 2 \\ 2 & 1 & 1 & 0 & 3 & 2 & 2 & 1 \\ 1 & 2 & 2 & 3 & 0 & 1 & 1 & 2 \\ 2 & 1 & 3 & 2 & 1 & 0 & 2 & 1 \\ 2 & 3 & 1 & 2 & 1 & 2 & 0 & 1 \\ 3 & 2 & 2 & 1 & 2 & 1 & 1 & 0 \end{matrix} \right] \end{matrix}$$

Example: Sequence Prediction with Hamming Loss

$$Y = \hat{Y} = \{0,1\}^r$$

$$n = k = 2$$

$$\mathbf{L}^H$$

$$\text{rank}(\mathbf{L}^{\text{Ham}}) = r$$

101	2	1	3	2	1	0	1	1	2
110	2	3	1	2	1	2	0	0	1
111	3	2	2	1	2	1	1	1	0

Example: Document Ranking with Pairwise Disagreement Loss

$$Y = \{0,1\}^r, \hat{Y} = S_r$$

$$n = 2^r, k = r!$$

$$r = 3$$

$$\begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} \begin{pmatrix} 1 \\ 3 \\ 2 \end{pmatrix} \begin{pmatrix} 2 \\ 3 \\ 1 \end{pmatrix} \begin{pmatrix} 2 \\ 1 \\ 3 \end{pmatrix} \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix} \begin{pmatrix} 3 \\ 2 \\ 1 \end{pmatrix}$$

$$\mathbf{L}^{\text{PD}} = \begin{matrix} \text{A 3x6 matrix of binary values (0s and 1s) representing pairwise comparisons between the 6 documents. The matrix is:} \\ \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 2 & 1 & 1 & 2 & 0 & 0 \\ 1 & 2 & 0 & 0 & 2 & 1 \\ 2 & 2 & 0 & 1 & 1 & 0 \\ 0 & 0 & 2 & 1 & 1 & 2 \\ 1 & 0 & 2 & 2 & 0 & 1 \\ 0 & 1 & 1 & 0 & 2 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \end{matrix}$$

Example: Document Ranking with Pairwise Disagreement Loss

$$Y = G_r, \hat{Y} = S_r$$

$$n = |G_r|, k = r!$$

$$r = 3$$

$$\begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} \begin{pmatrix} 1 \\ 3 \\ 2 \end{pmatrix} \begin{pmatrix} 2 \\ 3 \\ 1 \end{pmatrix} \begin{pmatrix} 2 \\ 1 \\ 3 \end{pmatrix} \begin{pmatrix} 3 \\ 1 \\ 2 \end{pmatrix} \begin{pmatrix} 3 \\ 2 \\ 1 \end{pmatrix}$$

$$\mathbf{L}^{\text{PD}} = \begin{array}{c} \begin{array}{c} 1 \\ 2 \xrightarrow{\hspace{1cm}} 3 \\ \vdots \\ 1 \\ 2 \xrightarrow{\hspace{1cm}} 3 \end{array} \end{array} \begin{bmatrix} 0 & 1 & 1 & 1 & 1 & 2 \\ & \vdots & & & & \\ 0 & 1 & 2 & 1 & 2 & 3 \end{bmatrix}$$


Example: Document Ranking with Pairwise Disagreement Loss

$$Y = G_r, \hat{Y} = S$$

$$n = |G_r|$$

$$\text{rank}(L^{\text{PD}}) = \Theta(r^2)$$

$$L^{\text{PD}} =$$



Application: Stronger Versions of Recent Results on Non-Existence of Convex Calibrated Surrogates

(Duchi et al, 2010; Calauzenes et al, 2012): no convex calibrated surrogates for \mathbf{L}^{PD} in $\leq r$ dimensions

Application: Stronger Versions of Recent Results on Non-Existence of Convex Calibrated Surrogates

(Duchi et al, 2010; Calauzenes et al, 2012): no convex calibrated surrogates for \mathbf{L}^{PD} in $\leq r$ dimensions

Our results: no convex calibrated surrogates for \mathbf{L}^{PD} in $< r(r-1)/2 - 2$ dimensions!

Our Work

When is a surrogate
calibrated for L ?

What is the smallest
dimension d that supports
a convex calibrated
surrogate for L ?


Can we design explicit
low-dimensional
surrogates for L ?

Explicit Convex Calibrated Least Squares Surrogate with $d = \text{rank}(\mathbf{L})$

$$\begin{bmatrix} \mathbf{L} \end{bmatrix}_{n \times k} = \begin{bmatrix} \mathbf{A} \end{bmatrix}_{n \times d} \begin{bmatrix} \mathbf{B} \end{bmatrix}_{d \times k}$$

[Ramaswamy, Agarwal & Tewari, 2013]

Explicit Convex Calibrated Least Squares Surrogate with $d = \text{rank}(\mathbf{L})$


$$\psi^{\text{LS}}(y, \mathbf{u}) = \sum_{j=1}^d (u_j - A_{yj})^2$$
$$\text{pred}^{\text{LS}}(\mathbf{u}) \in \operatorname{argmin}_{\hat{y} \in [k]} \sum_{j=1}^d u_j B_{j\hat{y}}$$

[Ramaswamy, Agarwal & Tewari, 2013]

Explicit Convex Calibrated Output Code Based Surrogate with $d = \text{rank}(\mathbf{L})$

$$\begin{bmatrix} \mathbf{L} \end{bmatrix}_{n \times k} = \begin{bmatrix} \mathbf{A} \end{bmatrix}_{n \times d} \begin{bmatrix} \mathbf{B} \end{bmatrix}_{d \times k}$$

[Ramaswamy, Babu, Agarwal & Williamson, 2014]

Explicit Convex Calibrated Output Code Based Surrogate with $d = \text{rank}(\mathbf{L})$

$$\psi^{\text{OC}}(y, \mathbf{u}) = \sum_{j=1}^d (C_{yj} \phi(1, u_j) + (1 - C_{yj}) \phi(-1, u_j))$$

$$\text{pred}^{\text{OC}}(\mathbf{u}) \in \operatorname{argmin}_{\hat{y} \in [k]} \sum_{j=1}^d \lambda^{-1}(u_j) \beta_{j\hat{y}}$$

[Ramaswamy, Babu, Agarwal & Williamson, 2014]

Explicit Convex Calibrated Output Code Based Surrogate with $d = \text{rank}(\mathbf{L})$

$$\psi^{\text{OC}}(y, \mathbf{u}) = \sum_{j=1}^d (C_{yj} \phi(1, u_j) + (1 - C_{yj}) \phi(-1, u_j))$$

Strictly proper composite
binary surrogate

Explicit Convex Calibrated Output Code Based Surrogate with $d = \text{rank}(\mathbf{L})$

$$\psi^{\text{OC}}(y, \mathbf{u}) = \sum_{j=1}^d (C_{yi} \phi(1, u_j) + (1 - C_{yi}) \phi(-1, u_j))$$

Code matrix
constructed from \mathbf{A}

Explicit Convex Calibrated Output Code Based Surrogate with $d = \text{rank}(\mathbf{L})$

$$\text{pred}^{\text{OC}}(\mathbf{u}) \in \operatorname{argmin}_{\hat{y} \in [k]} \sum_{j=1}^d \lambda^{-1}(u_j) \beta_{j\hat{y}}$$

Inverse link
associated with ϕ

[Ramaswamy, Babu, Agarwal & Williamson, 2014]

Explicit Convex Calibrated Output Code Based Surrogate with $d = \text{rank}(\mathbf{L})$

$$\text{pred}^{\text{OC}}(\mathbf{u}) \in \operatorname{argmin}_{\hat{y} \in [k]} \sum_{j=1}^d \lambda^{-1}(u_j) \beta_{j\hat{y}}$$

Obtained from \mathbf{B}

[Ramaswamy, Babu, Agarwal & Williamson, 2014]

Current/Future Directions

Relaxing universal
consistency requirement

Relaxing exact consistency
requirement

Complex performance
measures

Current/Future Directions

Relaxing universal
consistency requirement

Relaxing exact consistency
requirement


Complex performance
measures

Surrogates for Multiclass 0-1 Classification

Popular Crammer-Singer surrogate:


$$\psi^{\text{CS}}(y, \mathbf{u}) = \max_{\hat{y} \in [k]} (1 - (u_y - u_{\hat{y}}))_+$$

$$\text{pred}^{\text{CS}}(\mathbf{u}) \in \arg \max_{\hat{y} \in [k]} u_{\hat{y}}$$



[Crammer & Singer, 2001]


Surrogates for Multiclass 0-1 Classification



Crammer-Singer algorithm is 0-1
consistent only for
distributions D satisfying
'dominant-label' condition

[Zhang, 2004; Tewari & Bartlett, 2007]

Surrogates for Multiclass 0-1 Classification



Under dominant-label
condition, $\log n$
dimensions suffice!

[Ramaswamy, Tewari & Agarwal, 2015]

Surrogates for Document Ranking with Pairwise Disagreement Loss

Least Squares surrogate:

$$\psi^{\text{LS}}(y, \mathbf{u}) = \sum_{i=1}^r \sum_{j \neq i} (u_{ij} - y_{ij})^2$$

$\Theta(r^2)$ -dimensional surrogate

$$\text{pred}^{\text{LS}}(\mathbf{u}) \in \operatorname{argmin}_{\sigma \in S_r} \sum_{i=1}^r \sum_{j \neq i} u_{ij} \mathbf{1}(\sigma(i) > \sigma(j))$$

[Ramaswamy, Agarwal & Tewari, 2013]


Surrogates for Document Ranking with Pairwise Disagreement Loss

Resulting algorithm is
universally consistent, BUT...

[Ramaswamy, Agarwal & Tewari, 2013]

Surrogates for Document Ranking with Pairwise Disagreement Loss

Least Squares surrogate:

$$\psi^{\text{LS}}(y, \mathbf{u}) = \sum_{i=1}^r \sum_{j \neq i} (u_{ij} - y_{ij})^2$$
$$\text{pred}^{\text{LS}}(\mathbf{u}) \in \operatorname{argmin}_{\sigma \in S_r} \sum_{i=1}^r \sum_{j \neq i} u_{ij} \mathbf{1}(\sigma(i) > \sigma(j))$$


[Ramaswamy, Agarwal & Tewari, 2013]

Surrogates for Document Ranking with Pairwise Disagreement Loss



Efficient implementation
under 'DAG' condition

[Ramaswamy, Agarwal & Tewari, 2013]

Current/Future Directions




Relaxing universal
consistency requirement

Relaxing exact consistency
requirement


Complex performance
measures

Approximate Consistency


Data Space



Model Space



Approximate Consistency via Approximately Calibrated Surrogates



[Ramaswamy & Agarwal, in progress]

Current/Future Directions

Relaxing universal
consistency requirement

Relaxing exact consistency
requirement

Complex performance
measures

Example: F-Measure

$$Y = \hat{Y} = \{\pm 1\}$$

$$F_D[h] = \frac{2 \cdot \text{Prec}_D[h] \cdot \text{Rec}_D[h]}{\text{Prec}_D[h] + \text{Rec}_D[h]}$$

[Narasimhan, Vaish & Agarwal, 2014;
Narasimhan, Ramaswamy & Agarwal, 2015]

Summary

Given a learning problem defined by
a loss matrix \mathbf{L} , need \mathbf{L} -consistent
learning algorithms

Summary

Given a learning problem defined by
a loss matrix \mathbf{L} , need \mathbf{L} -consistent
learning algorithms

Many popular
learning algorithms
minimize a convex
surrogate loss


Summary

Given a learning problem defined by a loss matrix \mathbf{L} , need \mathbf{L} -consistent learning algorithms

Many popular learning algorithms minimize a convex surrogate loss

Our work:
General methodology for designing convex calibrated surrogates for any loss \mathbf{L}

Open Questions



General characterization of
convex calibration
dimension of a loss \mathbf{L} ?

Open Questions

General characterization of
convex calibration
dimension of a loss \mathbf{L} ?

Convex calibrated surrogates
with $d = \text{CCdim}(\mathbf{L})$?

Open Questions

General characterization of
convex calibration
dimension of a loss \mathbf{L} ?

Convex calibrated surrogates
with $d = \text{CCdim}(\mathbf{L})$?

Other structure in
loss matrices?

Acknowledgments




Harish G. Ramaswamy



Harikrishna Narasimhan



Rohit Vaish



Balaji S. Babu



Ambuj Tewari



Robert C. Williamson



सत्यमेव जयते

Department of
Science & Technology
Government of India



Indo-US Science &
Technology Forum