

Learning with a large number of labels


Krzysztof Dembczyński


Intelligent Decision Support Systems Laboratory (IDSS)
Poznań University of Technology, Poland





Pre-doc Summer School on Learning Systems
Monday, July 3, 2017
ETH Zürich, Switzerland











Alan Turing, 1912 births, 1954 deaths

20th-century mathematicians, 20th-century philosophers

Academics of the University of Manchester Institute of Science and Technology

Alumni of King's College, Cambridge Artificial intelligence researchers

Atheist philosophers, Bayesian statisticians, British cryptographers, British logicians

British long-distance runners, British male athletes, British people of World War II

Computability theorists, Computer designers, English atheists

English computer scientists, English inventors, English logicians

English long-distance runners, English mathematicians

English people of Scottish descent, English philosophers, Former Protestants

Fellows of the Royal Society, Gay men

Government Communications Headquarters people, History of artificial intelligence

Inventors who committed suicide, LGBT scientists

LGBT scientists from the United Kingdom, Male long-distance runners

Mathematicians who committed suicide, Officers of the Order of the British Empire

People associated with Bletchley Park, People educated at Sherborne School

People from Maida Vale, People from Wilmslow

People prosecuted under anti-homosexuality laws, Philosophers of mind

Philosophers who committed suicide, Princeton University alumni, 1930-39

Programmers who committed suicide, People who have received posthumous pardons

Recipients of British royal pardons, Academics of the University of Manchester

Suicides by cyanide poisoning, Suicides in England, Theoretical computer scientists

Setting

- **Multi-class classification:**

$$\mathbf{x} = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d \xrightarrow{h(\mathbf{x})} y \in \{1, \dots, m\}$$

	x_1	x_2	\dots	x_d	y
\mathbf{x}	4.0	2.5		-1.5	5

Setting

- **Multi-class classification:**

$$\mathbf{x} = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d \xrightarrow{h(\mathbf{x})} y \in \{1, \dots, m\}$$

	x_1	x_2	\dots	x_d	y
\mathbf{x}	4.0	2.5		-1.5	5

- **Multi-label classification:**

$$\mathbf{x} = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d \xrightarrow{h(\mathbf{x})} \mathbf{y} = (y_1, y_2, \dots, y_m) \in \{0, 1\}^m$$

	x_1	x_2	\dots	x_d	y_1	y_2	\dots	y_m
\mathbf{x}	4.0	2.5		-1.5	1	1		0

Extreme classification

Extreme classification \Rightarrow a **large** number of **labels** m ($\geq 10^5$)

Extreme classification

Extreme classification \Rightarrow a **large** number of **labels** m ($\geq 10^5$)

- **Predictive performance:**

Extreme classification

Extreme classification \Rightarrow a **large** number of **labels** m ($\geq 10^5$)

- **Predictive performance:**
 - ▶ Performance measures: Hamming loss, $\text{prec}@k$, $\text{NDCG}@k$, Macro F

Extreme classification

Extreme classification \Rightarrow a **large** number of **labels** $m (\geq 10^5)$

- **Predictive performance:**
 - ▶ Performance measures: Hamming loss, $\text{prec}@k$, $\text{NDCG}@k$, Macro F
- **Computational complexity:**

Extreme classification

Extreme classification \Rightarrow a **large** number of **labels** $m (\geq 10^5)$

- **Predictive performance:**
 - ▶ Performance measures: Hamming loss, $\text{prec}@k$, $\text{NDCG}@k$, Macro F
- **Computational complexity:**
 - ▶ time vs. space

Extreme classification

Extreme classification \Rightarrow a **large** number of **labels** $m (\geq 10^5)$

- **Predictive performance:**
 - ▶ Performance measures: Hamming loss, $\text{prec}@k$, $\text{NDCG}@k$, Macro F
- **Computational complexity:**
 - ▶ time vs. space
 - ▶ #examples vs. #features vs. #labels

Extreme classification

Extreme classification \Rightarrow a **large** number of **labels** $m (\geq 10^5)$


- **Predictive performance:**
 - ▶ Performance measures: Hamming loss, $\text{prec}@k$, $\text{NDCG}@k$, Macro F
- **Computational complexity:**
 - ▶ time vs. space
 - ▶ #examples vs. #features vs. #labels
 - ▶ training vs. validation vs. prediction

Supervised learning


Supervised learning

Test example x


Supervised learning




Supervised learning



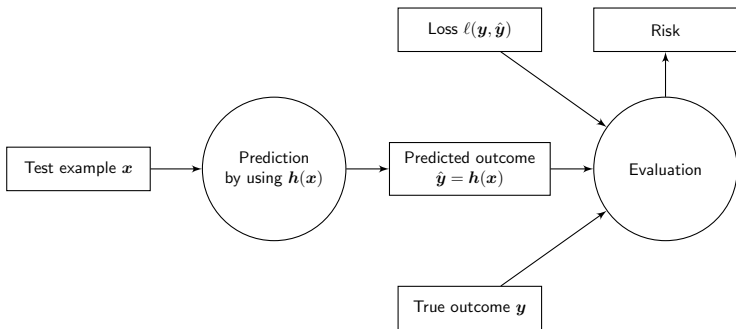
Supervised learning



Supervised learning




Supervised learning



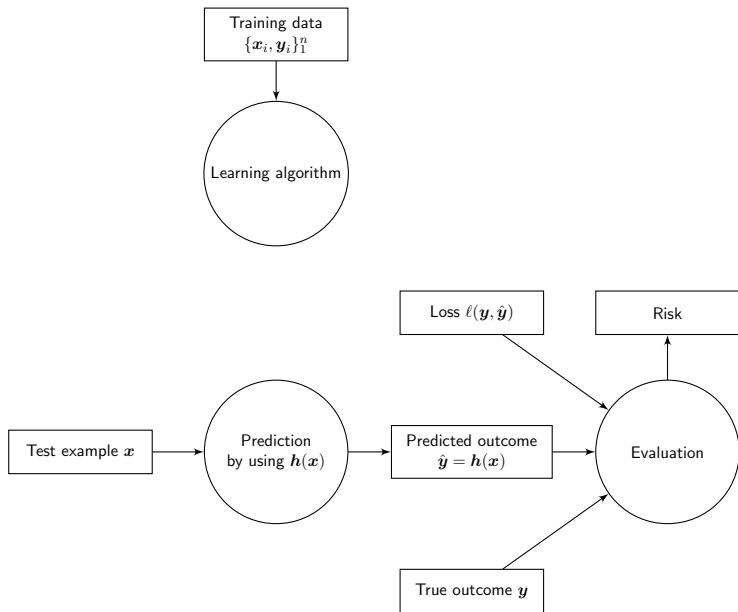
Supervised learning

Training data


$$\{\mathbf{x}_i, \mathbf{y}_i\}_1^n$$



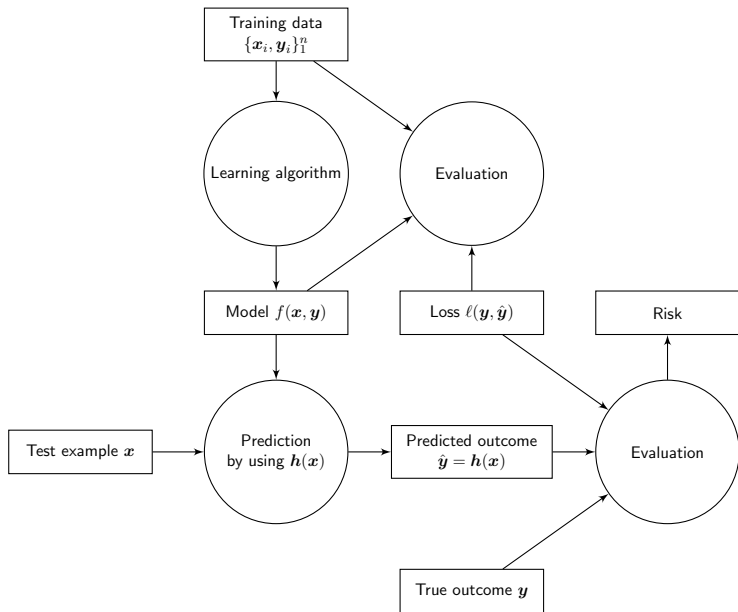
Supervised learning



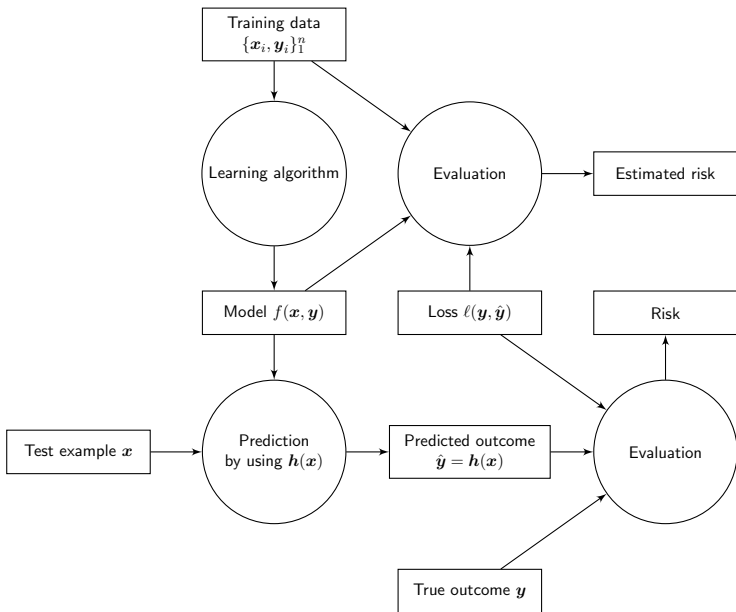
Supervised learning




Supervised learning




Supervised learning



Supervised learning



Supervised learning



Statistical learning framework

Statistical learning framework

- **Input** $x \in \mathcal{X}$ drawn from a distribution $\mathbf{P}(x)$.
 - ▶ usually a feature vector, $\mathcal{X} \subseteq \mathbb{R}^d$.

Statistical learning framework

- **Input** $x \in \mathcal{X}$ drawn from a distribution $\mathbf{P}(x)$.
 - ▶ usually a feature vector, $\mathcal{X} \subseteq \mathbb{R}^d$.
- **Outcome** $y \in \mathcal{Y}$ drawn from a distribution $\mathbf{P}(y | x)$.
 - ▶ a vector of labels $y = (y_1, y_2, \dots, y_m)$.

Statistical learning framework

- **Input** $x \in \mathcal{X}$ drawn from a distribution $\mathbf{P}(x)$.
 - ▶ usually a feature vector, $\mathcal{X} \subseteq \mathbb{R}^d$.
- **Outcome** $y \in \mathcal{Y}$ drawn from a distribution $\mathbf{P}(y | x)$.
 - ▶ a vector of labels $y = (y_1, y_2, \dots, y_m)$.
- **Prediction** $\hat{y} = h(x)$ by means of **prediction function** $h: \mathcal{X} \rightarrow \mathcal{Y}$.
 - ▶ h returns prediction $\hat{y} = h(x)$ for every input x .

Statistical learning framework

- **Input** $x \in \mathcal{X}$ drawn from a distribution $\mathbf{P}(x)$.
 - ▶ usually a feature vector, $\mathcal{X} \subseteq \mathbb{R}^d$.
- **Outcome** $y \in \mathcal{Y}$ drawn from a distribution $\mathbf{P}(y | x)$.
 - ▶ a vector of labels $y = (y_1, y_2, \dots, y_m)$.
- **Prediction** $\hat{y} = h(x)$ by means of **prediction function** $h: \mathcal{X} \rightarrow \mathcal{Y}$.
 - ▶ h returns prediction $\hat{y} = h(x)$ for every input x .
- **Loss** of our prediction: $\ell(y, \hat{y})$.
 - ▶ $\ell: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is a task-specific **loss function**.

Statistical learning framework

- **Input** $x \in \mathcal{X}$ drawn from a distribution $\mathbf{P}(x)$.
 - ▶ usually a feature vector, $\mathcal{X} \subseteq \mathbb{R}^d$.
- **Outcome** $y \in \mathcal{Y}$ drawn from a distribution $\mathbf{P}(y | x)$.
 - ▶ a vector of labels $y = (y_1, y_2, \dots, y_m)$.
- **Prediction** $\hat{y} = h(x)$ by means of **prediction function** $h: \mathcal{X} \rightarrow \mathcal{Y}$.
 - ▶ h returns prediction $\hat{y} = h(x)$ for every input x .
- **Loss** of our prediction: $\ell(y, \hat{y})$.
 - ▶ $\ell: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is a task-specific **loss function**.
- **Goal**: find a prediction function with small loss.

Risk

- **Goal:** minimize the **expected** loss over all examples (**risk**):

$$L_{\ell}(h) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathbf{P}} [\ell(\mathbf{y}, \mathbf{h}(\mathbf{x}))].$$

Risk

- **Goal:** minimize the **expected** loss over all examples (**risk**):

$$L_\ell(h) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathbf{P}} [\ell(\mathbf{y}, \mathbf{h}(\mathbf{x}))].$$

- The **optimal** prediction function over all possible functions expressed conditionally for a given \mathbf{x} :

$$\mathbf{h}^*(\mathbf{x}) = \arg \min_{\mathbf{h}} L_\ell(\mathbf{h}|\mathbf{x}),$$

(so called **Bayes prediction function**).

Hamming loss

- **Hamming loss:**

$$\ell_H(\mathbf{y}, \mathbf{h}(\mathbf{x})) = \frac{1}{m} \sum_{j=1}^m \mathbb{I}[y_j \neq h_j(\mathbf{x})],$$

¹ K. Dembczyński, W. Waegeman, W. Cheng, and E. Hüllermeier. On loss minimization and label dependence in multi-label classification. *Machine Learning*, 88:5–45, 2012

Hamming loss

- **Hamming loss:**

$$\ell_H(\mathbf{y}, \mathbf{h}(\mathbf{x})) = \frac{1}{m} \sum_{j=1}^m \mathbb{1}[y_j \neq h_j(\mathbf{x})],$$

- Sparse labels \Rightarrow Hamming loss of an **all-zero** classifier close to **0**.

¹ K. Dembczyński, W. Waegeman, W. Cheng, and E. Hüllermeier. On loss minimization and label dependence in multi-label classification. *Machine Learning*, 88:5–45, 2012

Hamming loss

- **Hamming loss:**

$$\ell_H(\mathbf{y}, \mathbf{h}(\mathbf{x})) = \frac{1}{m} \sum_{j=1}^m \mathbb{I}[y_j \neq h_j(\mathbf{x})],$$

- Sparse labels \Rightarrow Hamming loss of an **all-zero** classifier close to **0**.
- **The optimal strategy:**¹

$$h_j^*(\mathbf{x}) = \mathbb{I}[\eta_j(\mathbf{x}) > 0.5],$$

where $\eta_j(\mathbf{x}) = \mathbf{P}(y_j = 1 \mid \mathbf{x})$.

¹ K. Dembczyński, W. Waegeman, W. Cheng, and E. Hüllermeier. On loss minimization and label dependence in multi-label classification. *Machine Learning*, 88:5–45, 2012

Hamming loss


- **Hamming loss:**

$$\ell_H(\mathbf{y}, \mathbf{h}(\mathbf{x})) = \frac{1}{m} \sum_{j=1}^m \mathbb{I}[y_j \neq h_j(\mathbf{x})],$$

- Sparse labels \Rightarrow Hamming loss of an **all-zero** classifier close to **0**.
- **The optimal strategy:**¹

$$h_j^*(\mathbf{x}) = \mathbb{I}[\eta_j(\mathbf{x}) > 0.5],$$

where $\eta_j(\mathbf{x}) = \mathbf{P}(y_j = 1 \mid \mathbf{x})$.



¹ K. Dembczyński, W. Waegeman, W. Cheng, and E. Hüllermeier. On loss minimization and label dependence in multi-label classification. *Machine Learning*, 88:5–45, 2012

Precision

- **Precision at position k :**

$$\text{prec@}k(\mathbf{y}, f, \mathbf{x}) = \frac{1}{k} \sum_{r=1}^k \mathbb{I}[y_{\sigma(r)} = 1],$$

where σ is a permutation of labels for \mathbf{x} returned by ranker f .

Precision

- **Precision at position k :**

$$\text{prec@}k(\mathbf{y}, f, \mathbf{x}) = \frac{1}{k} \sum_{r=1}^k \mathbb{I}[y_{\sigma(r)} = 1],$$

where σ is a permutation of labels for \mathbf{x} returned by ranker f .

- **The optimal strategy:** select top k labels according to $\eta_j(\mathbf{x})$.


Precision

- **Precision at position k :**

$$\text{prec@}k(\mathbf{y}, f, \mathbf{x}) = \frac{1}{k} \sum_{r=1}^k \mathbb{1}[y_{\sigma(r)} = 1],$$

where σ is a permutation of labels for \mathbf{x} returned by ranker f .

- **The optimal strategy:** select top k labels according to $\eta_j(\mathbf{x})$.



Normalized Discounted Cumulative Gain

- **Normalized Discounted Cumulative Gain at position k :**

$$\text{NDCG}@k(\mathbf{y}, f, \mathbf{x}) = N_k(\mathbf{y}) \sum_{r=1}^k \frac{y_{\sigma(r)}}{\log(1+r)},$$

where σ is a permutation of labels for \mathbf{x} returned by ranker f , and $N_k(\mathbf{y})$ normalizes $\text{NDCG}@k$ to the interval $[0, 1]$:

$$N_k(\mathbf{y}) = \left(\sum_{r=1}^{\max(k, \sum_{i=1}^m y_i)} \frac{1}{\log(1+r)} \right)^{-1}$$

Normalized Discounted Cumulative Gain


- **The optimal strategy:** rank labels according to the following marginal quantities:

$$\Delta_j(\mathbf{x}) = \sum_{\mathbf{y}:y_j=1} N_k(\mathbf{y})\mathbf{P}(\mathbf{y} | \mathbf{x})$$

Normalized Discounted Cumulative Gain

- **The optimal strategy:** rank labels according to the following marginal quantities:

$$\Delta_j(\mathbf{x}) = \sum_{\mathbf{y}:y_j=1} N_k(\mathbf{y})\mathbf{P}(\mathbf{y} | \mathbf{x})$$



Macro-averaging of the F-measure

- The **macro** F-measure (F-score):

$$F_M(\mathbf{Y}, \hat{\mathbf{Y}}) = \frac{1}{m} \sum_{j=1}^m F(\mathbf{y}_{\cdot j}, \hat{\mathbf{y}}_{\cdot j}) = \frac{1}{m} \sum_{j=1}^m \frac{2 \sum_{i=1}^n y_{ij} \hat{y}_{ij}}{\sum_{i=1}^n y_{ij} + \sum_{i=1}^n \hat{y}_{ij}}.$$

True labels

y_{11}	y_{12}	y_{13}	y_{14}
y_{21}	y_{22}	y_{23}	y_{24}
y_{31}	y_{32}	y_{33}	y_{34}
y_{41}	y_{42}	y_{43}	y_{44}
y_{51}	y_{52}	y_{53}	y_{54}
y_{61}	y_{62}	y_{63}	y_{64}

Predicted labels

\hat{y}_{11}	\hat{y}_{12}	\hat{y}_{13}	\hat{y}_{14}
\hat{y}_{21}	\hat{y}_{22}	\hat{y}_{23}	\hat{y}_{24}
\hat{y}_{31}	\hat{y}_{32}	\hat{y}_{33}	\hat{y}_{34}
\hat{y}_{41}	\hat{y}_{42}	\hat{y}_{43}	\hat{y}_{44}
\hat{y}_{51}	\hat{y}_{52}	\hat{y}_{53}	\hat{y}_{54}
\hat{y}_{61}	\hat{y}_{62}	\hat{y}_{63}	\hat{y}_{64}

Macro-averaging of the F-measure

- The **macro** F-measure (F-score):

$$F_M(\mathbf{Y}, \hat{\mathbf{Y}}) = \frac{1}{m} \sum_{j=1}^m F(\mathbf{y}_{\cdot j}, \hat{\mathbf{y}}_{\cdot j}) = \frac{1}{m} \sum_{j=1}^m \frac{2 \sum_{i=1}^n y_{ij} \hat{y}_{ij}}{\sum_{i=1}^n y_{ij} + \sum_{i=1}^n \hat{y}_{ij}}.$$

True labels

y_{11}	y_{12}	y_{13}	y_{14}
y_{21}	y_{22}	y_{23}	y_{24}
y_{31}	y_{32}	y_{33}	y_{34}
y_{41}	y_{42}	y_{43}	y_{44}
y_{51}	y_{52}	y_{53}	y_{54}
y_{61}	y_{62}	y_{63}	y_{64}

Predicted labels

\hat{y}_{11}	\hat{y}_{12}	\hat{y}_{13}	\hat{y}_{14}
\hat{y}_{21}	\hat{y}_{22}	\hat{y}_{23}	\hat{y}_{24}
\hat{y}_{31}	\hat{y}_{32}	\hat{y}_{33}	\hat{y}_{34}
\hat{y}_{41}	\hat{y}_{42}	\hat{y}_{43}	\hat{y}_{44}
\hat{y}_{51}	\hat{y}_{52}	\hat{y}_{53}	\hat{y}_{54}
\hat{y}_{61}	\hat{y}_{62}	\hat{y}_{63}	\hat{y}_{64}

Macro-averaging of the F-measure

- The **macro** F-measure (F-score):

$$F_M(\mathbf{Y}, \hat{\mathbf{Y}}) = \frac{1}{m} \sum_{j=1}^m F(\mathbf{y}_{\cdot j}, \hat{\mathbf{y}}_{\cdot j}) = \frac{1}{m} \sum_{j=1}^m \frac{2 \sum_{i=1}^n y_{ij} \hat{y}_{ij}}{\sum_{i=1}^n y_{ij} + \sum_{i=1}^n \hat{y}_{ij}}.$$

True labels

y_{11}	y_{12}	y_{13}	y_{14}
y_{21}	y_{22}	y_{23}	y_{24}
y_{31}	y_{32}	y_{33}	y_{34}
y_{41}	y_{42}	y_{43}	y_{44}
y_{51}	y_{52}	y_{53}	y_{54}
y_{61}	y_{62}	y_{63}	y_{64}

Predicted labels

\hat{y}_{11}	\hat{y}_{12}	\hat{y}_{13}	\hat{y}_{14}
\hat{y}_{21}	\hat{y}_{22}	\hat{y}_{23}	\hat{y}_{24}
\hat{y}_{31}	\hat{y}_{32}	\hat{y}_{33}	\hat{y}_{34}
\hat{y}_{41}	\hat{y}_{42}	\hat{y}_{43}	\hat{y}_{44}
\hat{y}_{51}	\hat{y}_{52}	\hat{y}_{53}	\hat{y}_{54}
\hat{y}_{61}	\hat{y}_{62}	\hat{y}_{63}	\hat{y}_{64}

Macro-averaging of the F-measure

- The **macro** F-measure (F-score):

$$F_M(\mathbf{Y}, \hat{\mathbf{Y}}) = \frac{1}{m} \sum_{j=1}^m F(\mathbf{y}_{\cdot j}, \hat{\mathbf{y}}_{\cdot j}) = \frac{1}{m} \sum_{j=1}^m \frac{2 \sum_{i=1}^n y_{ij} \hat{y}_{ij}}{\sum_{i=1}^n y_{ij} + \sum_{i=1}^n \hat{y}_{ij}}.$$

True labels

y_{11}	y_{12}	y_{13}	y_{14}
y_{21}	y_{22}	y_{23}	y_{24}
y_{31}	y_{32}	y_{33}	y_{34}
y_{41}	y_{42}	y_{43}	y_{44}
y_{51}	y_{52}	y_{53}	y_{54}
y_{61}	y_{62}	y_{63}	y_{64}

Predicted labels

\hat{y}_{11}	\hat{y}_{12}	\hat{y}_{13}	\hat{y}_{14}
\hat{y}_{21}	\hat{y}_{22}	\hat{y}_{23}	\hat{y}_{24}
\hat{y}_{31}	\hat{y}_{32}	\hat{y}_{33}	\hat{y}_{34}
\hat{y}_{41}	\hat{y}_{42}	\hat{y}_{43}	\hat{y}_{44}
\hat{y}_{51}	\hat{y}_{52}	\hat{y}_{53}	\hat{y}_{54}
\hat{y}_{61}	\hat{y}_{62}	\hat{y}_{63}	\hat{y}_{64}

Macro-averaging of the F-measure

- The **macro** F-measure (F-score):

$$F_M(\mathbf{Y}, \hat{\mathbf{Y}}) = \frac{1}{m} \sum_{j=1}^m F(\mathbf{y}_{\cdot j}, \hat{\mathbf{y}}_{\cdot j}) = \frac{1}{m} \sum_{j=1}^m \frac{2 \sum_{i=1}^n y_{ij} \hat{y}_{ij}}{\sum_{i=1}^n y_{ij} + \sum_{i=1}^n \hat{y}_{ij}}.$$

True labels

y_{11}	y_{12}	y_{13}	y_{14}
y_{21}	y_{22}	y_{23}	y_{24}
y_{31}	y_{32}	y_{33}	y_{34}
y_{41}	y_{42}	y_{43}	y_{44}
y_{51}	y_{52}	y_{53}	y_{54}
y_{61}	y_{62}	y_{63}	y_{64}

Predicted labels

\hat{y}_{11}	\hat{y}_{12}	\hat{y}_{13}	\hat{y}_{14}
\hat{y}_{21}	\hat{y}_{22}	\hat{y}_{23}	\hat{y}_{24}
\hat{y}_{31}	\hat{y}_{32}	\hat{y}_{33}	\hat{y}_{34}
\hat{y}_{41}	\hat{y}_{42}	\hat{y}_{43}	\hat{y}_{44}
\hat{y}_{51}	\hat{y}_{52}	\hat{y}_{53}	\hat{y}_{54}
\hat{y}_{61}	\hat{y}_{62}	\hat{y}_{63}	\hat{y}_{64}

Macro-averaging of the F-measure

- Can be **solved** by **reduction** to m independent **binary** problems.²

² O. Koyejo, N. Natarajan, P. Ravikumar, and I. Dhillon. Consistent multilabel classification. In *NIPS*, 2015

Macro-averaging of the F-measure

- Can be **solved** by **reduction** to m independent **binary** problems.²
- **Thresholding** the conditional probabilities:

$$F(\tau) = \frac{2 \int_{\mathcal{X}} \eta(\mathbf{x}) \mathbb{I}[\eta(\mathbf{x}) \geq \tau] d\mu(\mathbf{x})}{\int_{\mathcal{X}} \eta(\mathbf{x}) d\mu(\mathbf{x}) + \int_{\mathcal{X}} \mathbb{I}[\eta(\mathbf{x}) \geq \tau] d\mu(\mathbf{x})}.$$

² O. Koyejo, N. Natarajan, P. Ravikumar, and I. Dhillon. Consistent multilabel classification. In *NIPS*, 2015

Macro-averaging of the F-measure

- Can be **solved** by **reduction** to m independent **binary** problems.²
- **Thresholding** the conditional probabilities:

$$F(\tau) = \frac{2 \int_{\mathcal{X}} \eta(\mathbf{x}) \mathbb{I}[\eta(\mathbf{x}) \geq \tau] d\mu(\mathbf{x})}{\int_{\mathcal{X}} \eta(\mathbf{x}) d\mu(\mathbf{x}) + \int_{\mathcal{X}} \mathbb{I}[\eta(\mathbf{x}) \geq \tau] d\mu(\mathbf{x})}.$$

- The **optimal F-measure** is $F(\tau^*)$: no binary classifier can be better.

² O. Koyejo, N. Natarajan, P. Ravikumar, and I. Dhillon. Consistent multilabel classification. In *NIPS*, 2015

Macro-averaging of the F-measure

- Can be **solved** by **reduction** to m independent **binary** problems.²
- **Thresholding** the conditional probabilities:

$$F(\tau) = \frac{2 \int_{\mathcal{X}} \eta(\mathbf{x}) \mathbb{I}[\eta(\mathbf{x}) \geq \tau] d\mu(\mathbf{x})}{\int_{\mathcal{X}} \eta(\mathbf{x}) d\mu(\mathbf{x}) + \int_{\mathcal{X}} \mathbb{I}[\eta(\mathbf{x}) \geq \tau] d\mu(\mathbf{x})}.$$

- The **optimal F-measure** is $F(\tau^*)$: no binary classifier can be better.
- The optimal solution **satisfies** the following **condition**: $F(\tau^*) = 2\tau^*$.


² O. Koyejo, N. Natarajan, P. Ravikumar, and I. Dhillon. Consistent multilabel classification. In *NIPS*, 2015

Macro-averaging of the F-measure

- Can be **solved** by **reduction** to m independent **binary** problems.²
- **Thresholding** the conditional probabilities:

$$F(\tau) = \frac{2 \int_{\mathcal{X}} \eta(\mathbf{x}) \mathbb{I}[\eta(\mathbf{x}) \geq \tau] d\mu(\mathbf{x})}{\int_{\mathcal{X}} \eta(\mathbf{x}) d\mu(\mathbf{x}) + \int_{\mathcal{X}} \mathbb{I}[\eta(\mathbf{x}) \geq \tau] d\mu(\mathbf{x})}.$$

- The **optimal F-measure** is $F(\tau^*)$: no binary classifier can be better.
- The optimal solution **satisfies** the following **condition**: $F(\tau^*) = 2\tau^*$.



² O. Koyejo, N. Natarajan, P. Ravikumar, and I. Dhillon. Consistent multilabel classification. In *NIPS*, 2015

Predictive model

- From the above analysis we can conclude:

Predictive model

- From the above analysis we can conclude:

We need to train models that accurately estimate marginal probabilities or other related marginal quantities

Predictive model

- From the above analysis we can conclude:

We need to train models that accurately estimate marginal probabilities or other related marginal quantities

- We could use to this end the well-known **1-vs-All** approach.

Computational challenges: naive solution

- Size of the problem:

Computational challenges: naive solution

- Size of the problem:
 - ▶ # examples: $n > 10^6$

Computational challenges: naive solution

- Size of the problem:
 - ▶ # examples: $n > 10^6$
 - ▶ # features: $d > 10^6$

Computational challenges: naive solution

- Size of the problem:
 - ▶ # examples: $n > 10^6$
 - ▶ # features: $d > 10^6$
 - ▶ # labels: $m > 10^5$

Computational challenges: naive solution

- Size of the problem:
 - ▶ # examples: $n > 10^6$
 - ▶ # features: $d > 10^6$
 - ▶ # labels: $m > 10^5$
- Naive solution: A dense linear model for each label (1-vs-All):

$$\hat{\mathbf{y}} = \mathbf{W}^\top \mathbf{x}$$

Computational challenges: naive solution

- Size of the problem:
 - ▶ # examples: $n > 10^6$
 - ▶ # features: $d > 10^6$
 - ▶ # labels: $m > 10^5$
- Naive solution: A dense linear model for each label (1-vs-All):

$$\hat{\mathbf{y}} = \mathbf{W}^\top \mathbf{x}$$

- ▶ Train time complexity:

Computational challenges: naive solution

- Size of the problem:
 - ▶ # examples: $n > 10^6$
 - ▶ # features: $d > 10^6$
 - ▶ # labels: $m > 10^5$
- Naive solution: A dense linear model for each label (1-vs-All):

$$\hat{\mathbf{y}} = \mathbf{W}^\top \mathbf{x}$$

- ▶ Train time complexity: $> 10^{17}$

Computational challenges: naive solution

- Size of the problem:
 - ▶ # examples: $n > 10^6$
 - ▶ # features: $d > 10^6$
 - ▶ # labels: $m > 10^5$
- Naive solution: A dense linear model for each label (1-vs-All):

$$\hat{\mathbf{y}} = \mathbf{W}^\top \mathbf{x}$$

- ▶ Train time complexity: $> 10^{17}$
- ▶ Space complexity:

Computational challenges: naive solution

- Size of the problem:
 - ▶ # examples: $n > 10^6$
 - ▶ # features: $d > 10^6$
 - ▶ # labels: $m > 10^5$
- Naive solution: A dense linear model for each label (1-vs-All):

$$\hat{\mathbf{y}} = \mathbf{W}^\top \mathbf{x}$$

- ▶ Train time complexity: $> 10^{17}$
- ▶ Space complexity: $> 10^{11}$

Computational challenges: naive solution

- Size of the problem:
 - ▶ # examples: $n > 10^6$
 - ▶ # features: $d > 10^6$
 - ▶ # labels: $m > 10^5$
- Naive solution: A dense linear model for each label (1-vs-All):

$$\hat{\mathbf{y}} = \mathbf{W}^\top \mathbf{x}$$

- ▶ Train time complexity: $> 10^{17}$
- ▶ Space complexity: $> 10^{11}$
- ▶ Test time complexity:

Computational challenges: naive solution

- Size of the problem:
 - ▶ # examples: $n > 10^6$
 - ▶ # features: $d > 10^6$
 - ▶ # labels: $m > 10^5$
- Naive solution: A dense linear model for each label (1-vs-All):

$$\hat{\mathbf{y}} = \mathbf{W}^\top \mathbf{x}$$

- ▶ Train time complexity: $> 10^{17}$
- ▶ Space complexity: $> 10^{11}$
- ▶ Test time complexity: $> 10^{11}$

Computational challenges: naive solution

- **It does not have to be so hard:**

Computational challenges: naive solution

- **It does not have to be so hard:**
 - ▶ High performance computing resources available.

Computational challenges: naive solution

- **It does not have to be so hard:**
 - ▶ High performance computing resources available.
 - ▶ Large data \rightarrow sparse data (sparse features and labels).

Computational challenges: naive solution

- **It does not have to be so hard:**
 - ▶ High performance computing resources available.
 - ▶ Large data \rightarrow sparse data (sparse features and labels).
 - ▶ Fast learning algorithms for standard learning problems exist.

```

0.912227 0.905463 22 22.0 1.0000 -0.1043 87
0.861865 0.811503 44 44.0 -1.0000 -0.0604 65
0.823944 0.785142 87 87.0 1.0000 -0.2309 60
0.766675 0.709405 174 174.0 1.0000 0.0754 25
0.642809 0.518943 348 348.0 1.0000 0.3440 47
0.540082 0.437356 696 696.0 1.0000 0.9767 24
0.450636 0.361190 1392 1392.0 1.0000 0.6204 181
0.376935 0.303234 2784 2784.0 1.0000 0.4380 50
0.320936 0.264938 5568 5568.0 -1.0000 -0.9257 89
0.281048 0.241153 11135 11135.0 1.0000 1.0000 62
0.249233 0.217415 22269 22269.0 1.0000 1.0000 140
0.221765 0.194296 44537 44537.0 1.0000 1.0000 41
0.201490 0.181213 89073 89073.0 -1.0000 -1.0000 27
0.187823 0.174157 178146 178146.0 1.0000 1.0000 49
0.176267 0.164711 356291 356291.0 -1.0000 -1.0000 100
0.165728 0.155188 712582 712582.0 -1.0000 -1.0000 69

finished run
number of examples = 781265
weighted example sum = 7.813e+05
weighted label sum = -4.018e+04
average loss = 0.1645
best constant = -0.05143
total feature number = 59936409
vw -c rcv1.train.txt 1.46s user 0.21s system 189% cpu 0.883 total
S:29PM 1-of-3-8: | ~/rcv1/norm [jl/ttypts/18]

```

Figure: Vowpal Wabbit³ at a lecture of John Langford⁴

³ Vowpal Wabbit, <http://hunch.net/~vw>

⁴ <http://cilvr.cs.nyu.edu/doku.php?id=courses:bigdata:slides:start>

Fast binary classification

- Data set: **RCV1**
- Predicted category: CCAT
- # training examples: 781 265
- # features: 60M
- Size: 1.1 GB
- Command line: `time vw -sgd rcv1.train.txt -c`
- Learning time: 1-3 secs on a laptop.

Reducing computational costs of the naive solution

- Linear models

Reducing computational costs of the naive solution

- Linear models
- Decision trees

Reducing computational costs of the naive solution

- Linear models
- Decision trees
- Label trees

Linear models

Linear models

- Fast training by least squares:⁵

⁵ T. Hastie, R. Tibshirani, and J.H. Friedman. *Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, second edition, 2009

Linear models

- Fast training by least squares:⁵

$$\mathbf{W} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$$

⁵ T. Hastie, R. Tibshirani, and J.H. Friedman. *Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, second edition, 2009

Linear models

- Fast training by least squares:⁵

$$\mathbf{W} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$$

- Works well in low dimensional feature spaces.
- Does not really improve space and test time complexity.

⁵ T. Hastie, R. Tibshirani, and J.H. Friedman. *Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, second edition, 2009

Linear models

- Training time complexity:

-
- ⁶ L. Bottou. Large-scale machine learning with stochastic gradient descent. In Yves Lechevallier and Gilbert Saporta, editors, *COMPSTAT*, pages 177–187, Paris, France, August 2010. Springer
- ⁷ R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008
- ⁸ John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *JMLR*, 10:2899–2934, 2009
- ⁹ Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, pages 160–167, 2008
- ¹⁰ Rohit Babbar and Bernhard Schölkopf. Dismec - distributed sparse machines for extreme multi-label classification. *CoRR*, 2016
- ¹¹ K.Q. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML*, pages 1113–1120. ACM, 2009

Linear models

- Training time complexity:
 - ▶ Stochastic gradient descent⁶ or coordinate gradient descent⁷

-
- ⁶ L. Bottou. Large-scale machine learning with stochastic gradient descent. In Yves Lechevallier and Gilbert Saporta, editors, *COMPSTAT*, pages 177–187, Paris, France, August 2010. Springer
- ⁷ R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008
- ⁸ John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *JMLR*, 10:2899–2934, 2009
- ⁹ Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, pages 160–167, 2008
- ¹⁰ Rohit Babbar and Bernhard Schölkopf. Dismec - distributed sparse machines for extreme multi-label classification. *CoRR*, 2016
- ¹¹ K.Q. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML*, pages 1113–1120. ACM, 2009

Linear models

- Training time complexity:
 - ▶ Stochastic gradient descent⁶ or coordinate gradient descent⁷
 - ▶ Sparse feature vectors (e.g., sparse updates in SGD⁸)

-
- ⁶ L. Bottou. Large-scale machine learning with stochastic gradient descent. In Yves Lechevallier and Gilbert Saporta, editors, *COMPSTAT*, pages 177–187, Paris, France, August 2010. Springer
- ⁷ R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008
- ⁸ John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *JMLR*, 10:2899–2934, 2009
- ⁹ Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, pages 160–167, 2008
- ¹⁰ Rohit Babbar and Bernhard Schölkopf. Dismec - distributed sparse machines for extreme multi-label classification. *CoRR*, 2016
- ¹¹ K.Q. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML*, pages 1113–1120. ACM, 2009

Linear models

- Training time complexity:
 - ▶ Stochastic gradient descent⁶ or coordinate gradient descent⁷
 - ▶ Sparse feature vectors (e.g., sparse updates in SGD⁸)
 - ▶ Negative sampling.⁹

-
- ⁶ L. Bottou. Large-scale machine learning with stochastic gradient descent. In Yves Lechevallier and Gilbert Saporta, editors, *COMPSTAT*, pages 177–187, Paris, France, August 2010. Springer
- ⁷ R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008
- ⁸ John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *JMLR*, 10:2899–2934, 2009
- ⁹ Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, pages 160–167, 2008
- ¹⁰ Rohit Babbar and Bernhard Schölkopf. Dismec - distributed sparse machines for extreme multi-label classification. *CoRR*, 2016
- ¹¹ K.Q. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML*, pages 1113–1120. ACM, 2009

Linear models

- Training time complexity:
 - ▶ Stochastic gradient descent⁶ or coordinate gradient descent⁷
 - ▶ Sparse feature vectors (e.g., sparse updates in SGD⁸)
 - ▶ Negative sampling.⁹
- Space complexity:

⁶ L. Bottou. Large-scale machine learning with stochastic gradient descent. In Yves Lechevallier and Gilbert Saporta, editors, *COMPSTAT*, pages 177–187, Paris, France, August 2010. Springer

⁷ R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008

⁸ John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *JMLR*, 10:2899–2934, 2009

⁹ Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, pages 160–167, 2008

¹⁰ Rohit Babbar and Bernhard Schölkopf. Dismec - distributed sparse machines for extreme multi-label classification. *CoRR*, 2016

¹¹ K.Q. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML*, pages 1113–1120. ACM, 2009

Linear models

- Training time complexity:
 - ▶ Stochastic gradient descent⁶ or coordinate gradient descent⁷
 - ▶ Sparse feature vectors (e.g., sparse updates in SGD⁸)
 - ▶ Negative sampling.⁹
- Space complexity:
 - ▶ Proper regularization: L_1 vs L_2 .

⁶ L. Bottou. Large-scale machine learning with stochastic gradient descent. In Yves Lechevallier and Gilbert Saporta, editors, *COMPSTAT*, pages 177–187, Paris, France, August 2010. Springer

⁷ R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008

⁸ John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *JMLR*, 10:2899–2934, 2009

⁹ Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, pages 160–167, 2008

¹⁰ Rohit Babbar and Bernhard Schölkopf. Dismec - distributed sparse machines for extreme multi-label classification. *CoRR*, 2016

¹¹ K.Q. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML*, pages 1113–1120. ACM, 2009

Linear models

- Training time complexity:
 - ▶ Stochastic gradient descent⁶ or coordinate gradient descent⁷
 - ▶ Sparse feature vectors (e.g., sparse updates in SGD⁸)
 - ▶ Negative sampling.⁹
- Space complexity:
 - ▶ Proper regularization: L_1 vs L_2 .
 - ▶ Removing small weights.¹⁰

⁶ L. Bottou. Large-scale machine learning with stochastic gradient descent. In Yves Lechevallier and Gilbert Saporta, editors, *COMPSTAT*, pages 177–187, Paris, France, August 2010. Springer

⁷ R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008

⁸ John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *JMLR*, 10:2899–2934, 2009

⁹ Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, pages 160–167, 2008

¹⁰ Rohit Babbar and Bernhard Schölkopf. Dismec - distributed sparse machines for extreme multi-label classification. *CoRR*, 2016

¹¹ K.Q. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML*, pages 1113–1120. ACM, 2009

Linear models

- Training time complexity:
 - ▶ Stochastic gradient descent⁶ or coordinate gradient descent⁷
 - ▶ Sparse feature vectors (e.g., sparse updates in SGD⁸)
 - ▶ Negative sampling.⁹
- Space complexity:
 - ▶ Proper regularization: L_1 vs L_2 .
 - ▶ Removing small weights.¹⁰
 - ▶ Feature hashing.¹¹

⁶ L. Bottou. Large-scale machine learning with stochastic gradient descent. In Yves Lechevallier and Gilbert Saporta, editors, *COMPSTAT*, pages 177–187, Paris, France, August 2010. Springer

⁷ R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008

⁸ John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *JMLR*, 10:2899–2934, 2009

⁹ Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, pages 160–167, 2008

¹⁰ Rohit Babbar and Bernhard Schölkopf. Dismec - distributed sparse machines for extreme multi-label classification. *CoRR*, 2016

¹¹ K.Q. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *ICML*, pages 1113–1120. ACM, 2009

Feature hashing in extreme classification

- Standard approach
 - ▶ A single slot for each weight and model.
 - ▶ Requires a lot of space.

x

0.1		0.5	1		0.3		-1
-----	--	-----	---	--	-----	--	----

w_1


w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
-------	-------	-------	-------	-------	-------	-------	-------

w_2

w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
-------	-------	-------	-------	-------	-------	-------	-------


Feature hashing in extreme classification

- Standard approach
 - ▶ A single slot for each weight and model.
 - ▶ Requires a lot of space.




Feature hashing in extreme classification

- Standard approach
 - ▶ A single slot for each weight and model.
 - ▶ Requires a lot of space.




Feature hashing in extreme classification

- Standard approach
 - ▶ A single slot for each weight and model.
 - ▶ Requires a lot of space.



Feature hashing in extreme classification

- Standard approach
 - ▶ A single slot for each weight and model.
 - ▶ Requires a lot of space.



Feature hashing in extreme classification

- Hashing to a common space
 - ▶ Hash the label and feature index using $h(j, v)$.
 - ▶ Hash a sign $\xi(j, v)$ to reduce the impact of conflicts.

x


0.1		0.5	1		0.3		-1
-----	--	-----	---	--	-----	--	----

w

w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8	w_9	w_{10}	w_{11}	w_{12}
-------	-------	-------	-------	-------	-------	-------	-------	-------	----------	----------	----------


Feature hashing in extreme classification

- Hashing to a common space
 - ▶ Hash the label and feature index using $h(j, v)$.
 - ▶ Hash a sign $\xi(j, v)$ to reduce the impact of conflicts.




Feature hashing in extreme classification

- Hashing to a common space
 - ▶ Hash the label and feature index using $h(j, v)$.
 - ▶ Hash a sign $\xi(j, v)$ to reduce the impact of conflicts.




Feature hashing in extreme classification

- Hashing to a common space
 - ▶ Hash the label and feature index using $h(j, v)$.
 - ▶ Hash a sign $\xi(j, v)$ to reduce the impact of conflicts.



Feature hashing in extreme classification

- Hashing to a common space
 - ▶ Hash the label and feature index using $h(j, v)$.
 - ▶ Hash a sign $\xi(j, v)$ to reduce the impact of conflicts.



Linear models

- Low-dimensional representation of x , \mathbf{W} , y :

$$y = \mathbf{U}^\dagger \mathbf{V}x$$

- ▶ feature space: PCA on \mathbf{X} .
- ▶ label space: PCA on \mathbf{Y} ,¹² compressed sensing,¹³ etc.
- ▶ both spaces: CCA on both \mathbf{X} and \mathbf{Y} ,¹⁴ etc.
- ▶ matrix factorization of \mathbf{W} .¹⁵
- ▶ A kind of **lossy compression/embedding** methods.

¹² F. Tai and H.-T. Lin. Multi-label classification with principal label space transformation. In *Neural Computat.*, volume 9, pages 2508–2542, 2012

¹³ D. Hsu, S. Kakade, J. Langford, and T. Zhang. Multi-label prediction via compressed sensing. In *NIPS*, 2009

¹⁴ Yao-Nan Chen and Hsuan-Tien Lin. Feature-aware label space dimension reduction for multi-label classification. In *NIPS*, pages 1529–1537. Curran Associates, Inc., 2012

¹⁵ Hsiang-Fu Yu, Prateek Jain, Purushottam Kar, and Inderjit S. Dhillon. Large-scale Multi-label Learning with Missing Labels. In *ICML*, 2014

Computational challenges

- Prediction time is still **linear** in the number of labels!

Computational challenges

- Prediction time is still **linear** in the number of labels!
- **Reduce** the test time complexity by:
 - ▶ Maximum inner product search over linear models,
 - ▶ Decision trees,
 - ▶ Label trees.

Test time complexity for linear models

- Classification of a test example in case of linear models can be formulated as:

$$j^* = \arg \max_{j \in \{1, \dots, m\}} \mathbf{w}_j^\top \mathbf{x},$$

i.e., the problem of **maximum inner product search (MIPS)**.

MIPS vs. nearest neighbors

- MIPS is similar, but not the same, to the nearest neighbor search under the square or cosine distance:

$$j^* = \arg \min_{j \in \{1, \dots, m\}} \|\mathbf{w}_j - \mathbf{x}\|_2^2 = \arg \max_{j \in \{1, \dots, m\}} \mathbf{w}_j^\top \mathbf{x} - \frac{\|\mathbf{w}_j\|_2^2}{2}$$

$$j^* = \arg \max_{j \in \{1, \dots, m\}} \frac{\mathbf{w}_j^\top \mathbf{x}}{\|\mathbf{w}_j\| \|\mathbf{x}\|} = \arg \max_{j \in \{1, \dots, m\}} \frac{\mathbf{w}_j^\top \mathbf{x}}{\|\mathbf{w}_j\|}$$

¹⁶ A. Shrivastava and P. Li. Improved asymmetric locality sensitive hashing (ALSH) for maximum inner product search (mips). In *UAI*, 2015

¹⁷ J. H. Friedman, J. L. Bentley, and R. A. Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software*, 3(3):209–226, 1977

¹⁸ Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *ACM Symposium on Theory of Computing*, STOC '98, pages 604–613, New York, NY, USA, 1998. ACM

MIPS vs. nearest neighbors

- MIPS is similar, but not the same, to the nearest neighbor search under the square or cosine distance:

$$j^* = \arg \min_{j \in \{1, \dots, m\}} \|\mathbf{w}_j - \mathbf{x}\|_2^2 = \arg \max_{j \in \{1, \dots, m\}} \mathbf{w}_j^\top \mathbf{x} - \frac{\|\mathbf{w}_j\|_2^2}{2}$$

$$j^* = \arg \max_{j \in \{1, \dots, m\}} \frac{\mathbf{w}_j^\top \mathbf{x}}{\|\mathbf{w}_j\| \|\mathbf{x}\|} = \arg \max_{j \in \{1, \dots, m\}} \frac{\mathbf{w}_j^\top \mathbf{x}}{\|\mathbf{w}_j\|}$$

- Some tricks are used to treat MIPS as nearest neighbor search.¹⁶

¹⁶ A. Shrivastava and P. Li. Improved asymmetric locality sensitive hashing (ALSH) for maximum inner product search (mips). In *UAI*, 2015

¹⁷ J. H. Friedman, J. L. Bentley, and R. A. Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software*, 3(3):209–226, 1977

¹⁸ Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *ACM Symposium on Theory of Computing*, STOC '98, pages 604–613, New York, NY, USA, 1998. ACM

MIPS vs. nearest neighbors

- MIPS is similar, but not the same, to the nearest neighbor search under the square or cosine distance:

$$j^* = \arg \min_{j \in \{1, \dots, m\}} \|\mathbf{w}_j - \mathbf{x}\|_2^2 = \arg \max_{j \in \{1, \dots, m\}} \mathbf{w}_j^\top \mathbf{x} - \frac{\|\mathbf{w}_j\|_2^2}{2}$$

$$j^* = \arg \max_{j \in \{1, \dots, m\}} \frac{\mathbf{w}_j^\top \mathbf{x}}{\|\mathbf{w}_j\| \|\mathbf{x}\|} = \arg \max_{j \in \{1, \dots, m\}} \frac{\mathbf{w}_j^\top \mathbf{x}}{\|\mathbf{w}_j\|}$$

- Some tricks are used to treat MIPS as nearest neighbor search.¹⁶
 - ▶ For low-dimensional problems, efficient tree-based structures exist.¹⁷

¹⁶ A. Shrivastava and P. Li. Improved asymmetric locality sensitive hashing (ALSH) for maximum inner product search (mips). In *UAI*, 2015

¹⁷ J. H. Friedman, J. L. Bentley, and R. A. Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software*, 3(3):209–226, 1977

¹⁸ Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *ACM Symposium on Theory of Computing*, STOC '98, pages 604–613, New York, NY, USA, 1998. ACM

MIPS vs. nearest neighbors

- MIPS is similar, but not the same, to the nearest neighbor search under the square or cosine distance:

$$j^* = \arg \min_{j \in \{1, \dots, m\}} \|\mathbf{w}_j - \mathbf{x}\|_2^2 = \arg \max_{j \in \{1, \dots, m\}} \mathbf{w}_j^\top \mathbf{x} - \frac{\|\mathbf{w}_j\|_2^2}{2}$$

$$j^* = \arg \max_{j \in \{1, \dots, m\}} \frac{\mathbf{w}_j^\top \mathbf{x}}{\|\mathbf{w}_j\| \|\mathbf{x}\|} = \arg \max_{j \in \{1, \dots, m\}} \frac{\mathbf{w}_j^\top \mathbf{x}}{\|\mathbf{w}_j\|}$$

- Some tricks are used to treat MIPS as nearest neighbor search.¹⁶
 - ▶ For low-dimensional problems, efficient tree-based structures exist.¹⁷
 - ▶ Approximate nearest neighbor search via locality-sensitive hashing.¹⁸

¹⁶ A. Shrivastava and P. Li. Improved asymmetric locality sensitive hashing (ALSH) for maximum inner product search (mips). In *UAI*, 2015

¹⁷ J. H. Friedman, J. L. Bentley, and R. A. Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software*, 3(3):209–226, 1977

¹⁸ Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *ACM Symposium on Theory of Computing*, STOC '98, pages 604–613, New York, NY, USA, 1998. ACM

Decision trees

Decision trees

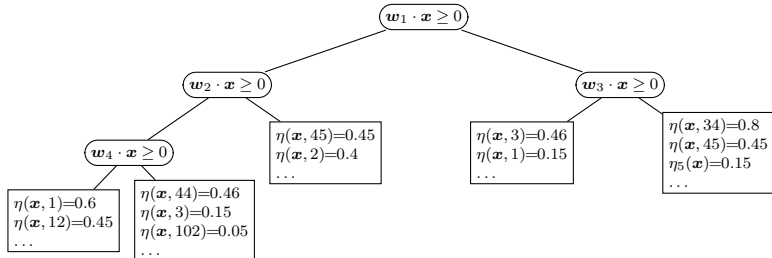
- Fast prediction: logarithmic in n
- Training can be expensive: computation of split criterion
- Two new algorithms: LomTree¹⁹ and **FastXML**²⁰

¹⁹ Anna Choromanska and John Langford. Logarithmic time online multiclass prediction. In *NIPS* 29, 2015

²⁰ Yashoteja Prabhu and Manik Varma. Fastxml: A fast, accurate and stable tree-classifier for extreme multi-label learning. In *KDD*, pages 263–272. ACM, 2014

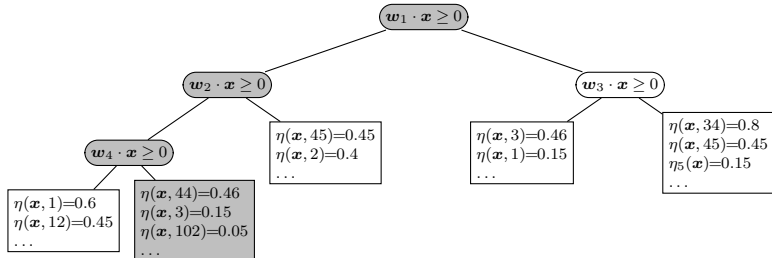
FastXML

- Uses an **ensemble** of standard decision trees.
- **Sparse linear** classifiers trained in internal nodes.
- Very **efficient** training procedure.
- **Empirical distributions** in leaves.
- A test example passes **one path** from the root to a leaf.



FastXML

- Uses an **ensemble** of standard decision trees.
- **Sparse linear** classifiers trained in internal nodes.
- Very **efficient** training procedure.
- **Empirical distributions** in leaves.
- A test example passes **one path** from the root to a leaf.



Optimization in FastXML

- In each internal node FastXML solves:

$$\begin{aligned} \min \quad & \|\mathbf{w}\|_1 + \sum_{i=1}^n C_\delta(\delta_i) \log(1 + \exp(-\delta_i \mathbf{w}^\top \mathbf{x})) \\ & - C_r \sum_{i=1}^n \frac{1}{2} (1 + \delta_i) \text{NDCG}@m(\mathbf{r}^+, \mathbf{y}_i) \\ & - C_r \sum_{i=1}^n \frac{1}{2} (1 - \delta_i) \text{NDCG}@m(\mathbf{r}^-, \mathbf{y}_i) \\ \text{w.r.t.} \quad & \mathbf{w} \in \mathbb{R}^d, \boldsymbol{\delta} \in \{-1, 1\}^m, \mathbf{r}^+, \mathbf{r}^- \in \Pi(1, m) \end{aligned}$$

Optimization in FastXML

- In each internal node FastXML solves:

$$\begin{aligned} \min \quad & \|\mathbf{w}\|_1 + \sum_{i=1}^n C_\delta(\delta_i) \log(1 + \exp(-\delta_i \mathbf{w}^\top \mathbf{x})) \\ & - C_r \sum_{i=1}^n \frac{1}{2} (1 + \delta_i) \text{NDCG}@m(\mathbf{r}^+, \mathbf{y}_i) \\ & - C_r \sum_{i=1}^n \frac{1}{2} (1 - \delta_i) \text{NDCG}@m(\mathbf{r}^-, \mathbf{y}_i) \end{aligned}$$

w.r.t. $\mathbf{w} \in \mathbb{R}^d, \boldsymbol{\delta} \in \{-1, 1\}^m, \mathbf{r}^+, \mathbf{r}^- \in \Pi(1, m)$

linear split

partitioning of training examples

label ranking in positive and negative partition

Optimization in FastXML

- In each internal node FastXML solves:

$$\min \quad \|\mathbf{w}\|_1 + \sum_{i=1}^n C_\delta(\delta_i) \log(1 + \exp(-\delta_i \mathbf{w}^\top \mathbf{x}_i))$$

1. Bernoulli sampling of δ
2. Optimize \mathbf{r}^\pm
3. Optimize δ
4. Optimize \mathbf{w}
5. Repeat 2-4

$$-C_r \sum_{i=1}^n \frac{1}{2} (1 + \delta_i) \text{NDCG}@m(\mathbf{r}^+, \mathbf{y}_i)$$

$$-C_r \sum_{i=1}^n \frac{1}{2} (1 - \delta_i) \text{NDCG}@m(\mathbf{r}^-, \mathbf{y}_i)$$

w.r.t. $\mathbf{w} \in \mathbb{R}^d, \delta \in \{-1, 1\}^m, \mathbf{r}^+, \mathbf{r}^- \in \Pi(1, m)$

linear split


partitioning of training examples

label ranking in positive and negative partition

Label trees

Label trees

- Organize classifiers in a tree structure (one leaf \Leftrightarrow one label).²¹




- Structure of the tree can be given or trained.
- Different training and test procedures for multi-class and multi-label classification.

²¹ S. Bengio, J. Weston, and D. Grangier. Label embedding trees for large multi-class tasks. In *NIPS*, pages 163–171. Curran Associates, Inc., 2010

Probabilistic label trees (PLT)²²

- PLT are based on b -ary **label trees**.




- **Probabilistic classifiers** in **all** nodes of the tree.
- **Internal** node classifier decides whether to **go down the tree**.
- A test example may follow **many paths** from the root to leaves.
- **Batch** and **online** learning possible.

²² K. Jasinska, K. Dembczynski, R. Busa-Fekete, K. Pfannschmidt, T. Klerx, and E. Hüllermeier. Extreme F-measure maximization using sparse probability estimates. In *ICML*, 2016

Probabilistic label trees

- Class probability estimators in nodes for estimating $\mathbf{P}(y_j = 1 \mid \mathbf{x})$.



- Using the **chain rule** of probability

$$\mathbf{P}(y_j = 1 \mid \mathbf{x}) = \eta_j(\mathbf{x}) = \prod_{t \in \text{Path}(j)} \eta(\mathbf{x}, t),$$

$$\text{where } \eta(\mathbf{x}, t) = \begin{cases} \mathbf{P}(z_t = 1 \mid \mathbf{x}) & \text{if } t \text{ is root,} \\ \mathbf{P}(z_t = 1 \mid z_{\text{pa}(t)} = 1, \mathbf{x}) & \text{otherwise.} \end{cases}$$

Probabilistic label trees

- Chain rule of probability:

$$\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1 \mid \mathbf{x}) \times \mathbf{P}(y_1 \vee y_2 = 1 \mid y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x}) =$$

Probabilistic label trees

- Chain rule of probability:

$$\begin{aligned} & \mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1 \mid \mathbf{x}) \times \mathbf{P}(y_1 \vee y_2 = 1 \mid y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x}) = \\ & \frac{\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} \times \frac{\mathbf{P}(y_1 \vee y_2 = 1, y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})}{\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})} = \end{aligned}$$

Probabilistic label trees

- Chain rule of probability:

$$\begin{aligned} & \mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1 \mid \mathbf{x}) \times \mathbf{P}(y_1 \vee y_2 = 1 \mid y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x}) = \\ & \frac{\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} \times \frac{\mathbf{P}(y_1 \vee y_2 = 1, y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})}{\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})} = \\ & \frac{\mathbf{P}(y_1 \vee y_2 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} = \mathbf{P}(y_1 \vee y_2 = 1 \mid \mathbf{x}) \end{aligned}$$

Probabilistic label trees

- Chain rule of probability:

$$\begin{aligned} & \mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1 \mid \mathbf{x}) \times \mathbf{P}(y_1 \vee y_2 = 1 \mid y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x}) = \\ & \frac{\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} \times \frac{\mathbf{P}(y_1 \vee y_2 = 1, y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})}{\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})} = \\ & \frac{\mathbf{P}(y_1 \vee y_2 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} = \mathbf{P}(y_1 \vee y_2 = 1 \mid \mathbf{x}) \end{aligned}$$

$$\mathbf{P}(y_1 \vee y_2 = 1 \mid \mathbf{x}) \times \mathbf{P}(y_1 = 1 \mid y_1 \vee y_2 = 1, \mathbf{x}) =$$

Probabilistic label trees

- Chain rule of probability:

$$\begin{aligned} & \mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1 \mid \mathbf{x}) \times \mathbf{P}(y_1 \vee y_2 = 1 \mid y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x}) = \\ & \frac{\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} \times \frac{\mathbf{P}(y_1 \vee y_2 = 1, y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})}{\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})} = \\ & \frac{\mathbf{P}(y_1 \vee y_2 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} = \mathbf{P}(y_1 \vee y_2 = 1 \mid \mathbf{x}) \end{aligned}$$

$$\begin{aligned} & \mathbf{P}(y_1 \vee y_2 = 1 \mid \mathbf{x}) \times \mathbf{P}(y_1 = 1 \mid y_1 \vee y_2 = 1, \mathbf{x}) = \\ & \frac{\mathbf{P}(y_1 \vee y_2 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} \times \frac{\mathbf{P}(y_1 = 1, y_1 \vee y_2 = 1, \mathbf{x})}{\mathbf{P}(y_1 \vee y_2 = 1, \mathbf{x})} = \end{aligned}$$

Probabilistic label trees


- Chain rule of probability:

$$\begin{aligned} & \mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1 \mid \mathbf{x}) \times \mathbf{P}(y_1 \vee y_2 = 1 \mid y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x}) = \\ & \frac{\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} \times \frac{\mathbf{P}(y_1 \vee y_2 = 1, y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})}{\mathbf{P}(y_1 \vee y_2 \vee y_3 \vee y_4 = 1, \mathbf{x})} = \\ & \frac{\mathbf{P}(y_1 \vee y_2 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} = \mathbf{P}(y_1 \vee y_2 = 1 \mid \mathbf{x}) \end{aligned}$$

$$\begin{aligned} & \mathbf{P}(y_1 \vee y_2 = 1 \mid \mathbf{x}) \times \mathbf{P}(y_1 = 1 \mid y_1 \vee y_2 = 1, \mathbf{x}) = \\ & \frac{\mathbf{P}(y_1 \vee y_2 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} \times \frac{\mathbf{P}(y_1 = 1, y_1 \vee y_2 = 1, \mathbf{x})}{\mathbf{P}(y_1 \vee y_2 = 1, \mathbf{x})} = \\ & \frac{\mathbf{P}(y_1 = 1, \mathbf{x})}{\mathbf{P}(\mathbf{x})} = \mathbf{P}(y_1 = 1 \mid \mathbf{x}) \end{aligned}$$


Probabilistic label trees

- Class probability estimators in nodes for estimating $\mathbf{P}(y_i = 1 \mid \mathbf{x})$.



Probabilistic label trees


- Class probability estimators in nodes for estimating $\mathbf{P}(y_i = 1 \mid \mathbf{x})$.



- Training: reduced complexity by the **conditions** used in the **nodes**.

Probabilistic label trees

- Class probability estimators in nodes for estimating $\mathbf{P}(y_i = 1 | \mathbf{x})$.



- Training: reduced complexity by the **conditions** used in the **nodes**.
- Prediction: **priority queue search** or **branch and bound**.

Probabilistic label trees

- The same idea under different names:
 - ▶ **Conditional probability trees**²³
 - ▶ **Probabilistic classifier chains**²⁴
 - ▶ **Hierarchical softmax**²⁵
 - ▶ **Homer**²⁶
 - ▶ **Nested dichotomies**²⁷
 - ▶ **Multi-stage classification**²⁸

²³ A. Beygelzimer, J. Langford, Y. Lifshits, G. B. Sorkin, and A. L. Strehl. Conditional probability tree estimation analysis and algorithms. In *UAI*, pages 51–58, 2009

²⁴ K. Dembczyński, W. Cheng, and E. Hüllermeier. Bayes optimal multilabel classification via probabilistic classifier chains. In *ICML*, pages 279–286. Omnipress, 2010

²⁵ Frederic Morin and Yoshua Bengio. Hierarchical probabilistic neural network language model. In *AISTATS*, pages 246–252, 2005

²⁶ G. Tsoumakas, I. Katakis, and I. Vlahavas. Effective and efficient multilabel classification in domains with large number of labels. In *Proc. ECML/PKDD 2008 Workshop on Mining Multidimensional Data*, 2008

²⁷ J. Fox. *Applied regression analysis, linear models, and related methods*. Sage, 1997

²⁸ Marek Kurzynski. On the multistage bayes classifier. *Pattern Recognition*, 21(4):355–365, 1988

FastXML vs. PLT

	FastXML	PLT
tree structure	✓	✓
structure learning	✓	×
number of trees	≥ 1	1
number of leaves	$O(n)$	m
internal nodes models	linear	linear
leaves models	empirical distribution	linear
visited paths during prediction	1 per tree	several
sparse probability estimation	✓	✓

Experimental results

	#labels	#features	#test	#train	inst./lab.	lab./inst.
RCV1	2456	47236	155962	623847	1218.56	4.79
AmazonCat	13330	203882	306782	1186239	448.57	5.04
Wiki10	30938	101938	6616	14146	8.52	18.64
Delicious	205443	782585	100095	196606	72.29	75.54
WikiLSHTC	325056	1617899	587084	1778351	17.46	3.19
Amazon	670091	135909	153025	490449	3.99	5.45

Table: Datasets from the Extreme Classification repository.²⁹

²⁹ <http://manikvarma.org/downloads/XC/XMLRepository.html>

Experimental results

	PLT			FastXML		
	P@1	P@3	P@5	P@1	P@3	P@5
RCV1	90.46	72.4	51.86	91.13	73.35	52.67
AmazonCat	91.47	75.84	61.02	92.95	77.5	62.51
Wiki10	84.34	72.34	62.72	81.71	66.67	56.70
Delicious	45.37	38.94	35.88	42.81	38.76	36.34
WikiLSHTC	45.67	29.13	21.95	49.35	32.69	24.03
Amazon	36.65	32.12	28.85	34.24	29.3	26.12

Experimental results

	PLT					FastXML			
	train [min]	test [ms]	b	depth	#calls	train [min]	test [ms]	depth	#calls
RCV1	64	0.22	32	2,25	43,46	78	0.96	14.95	747
AmazonCat	96	0.17	16	3,43	54,39	561	1.14	17.44	871
Wiki10	290	2.66	32	2,98	121,98	16	3.00	10.83	541
Delicious	1327	32.97	2	17,69	11779,65	458	4.01	14.79	739
WikiLSHTC	653	3.00	32	3,66	622,27	724	1.17	18.01	900
Amazon	54	0.99	8	6,45	374,30	422	1.39	15.92	796

Challenges

Challenges

- Learning theory for an extremely large number of labels:

Challenges

- Learning theory for an extremely large number of labels:
 - ▶ **Statistical guarantees** for the **error** rate that **do not depend**, or depend very weakly (sublinearly), on the **total number of labels**.

Challenges

- Learning theory for an extremely large number of labels:
 - ▶ **Statistical guarantees** for the **error** rate that **do not depend**, or depend very weakly (sublinearly), on the **total number of labels**.
 - ▶ The **bound** on the error rate could be expressed in terms of the average number of **positive labels** (which is certainly much less than the total number of labels).

Challenges

- Learning theory for an extremely large number of labels:
 - ▶ **Statistical guarantees** for the **error** rate that **do not depend**, or depend very weakly (sublinearly), on the **total number of labels**.
 - ▶ The **bound** on the error rate could be expressed in terms of the average number of **positive labels** (which is certainly much less than the total number of labels).
 - ▶ Particular performance guarantees depend on the considered **loss function**.

Challenges

- Training and prediction under limited time and space budget:

Challenges

- Training and prediction under limited time and space budget:
 - ▶ **Restricted** computational **resources** (time and space) for both **training** and **prediction**.

Challenges

- Training and prediction under limited time and space budget:
 - ▶ **Restricted** computational **resources** (time and space) for both **training** and **prediction**.
 - ▶ A **trade-off** between computational (time and space) **complexity** and the **predictive performance**.

Challenges

- Training and prediction under limited time and space budget:
 - ▶ **Restricted** computational **resources** (time and space) for both **training** and **prediction**.
 - ▶ A **trade-off** between computational (time and space) **complexity** and the **predictive performance**.
 - ▶ By imposing hard constraints on time and space budget, the challenge is then to **optimize** the **predictive performance** of an algorithm under these **constraints**.

Challenges

- Unreliable learning information:

Challenges

- Unreliable learning information:
 - ▶ We **cannot** expect that all labels will be properly **checked** and **assigned** to training examples.

Challenges

- Unreliable learning information:
 - ▶ We **cannot** expect that all labels will be properly **checked** and **assigned** to training examples.
 - ▶ Therefore we often deal with a problem of learning with **missing labels** or learning from **positive and unlabeled examples**.

Challenges

- Performance measures:

Challenges

- Performance measures:
 - ▶ Typical performance measures such as **0/1** or **Hamming** loss do **not fit** well to the extreme setting.

Challenges

- Performance measures:
 - ▶ Typical performance measures such as **0/1** or **Hamming** loss do **not fit** well to the extreme setting.
 - ▶ Other measures are often used such as **precision@k** or the **F-measure**.

Challenges

- Performance measures:
 - ▶ Typical performance measures such as **0/1** or **Hamming** loss do **not fit** well to the extreme setting.
 - ▶ Other measures are often used such as **precision@k** or the **F-measure**.
 - ▶ However, it remains an **open question** how to **design loss functions** suitable for extreme classification.

Do we search in the right place?



Figure: ³⁰ A similar comics has been earlier used by Asela Gunawardana.³¹

³⁰ Source: Florence Morning News, Mutt and Jeff Comic Strip, Page 7, Florence, South Carolina, 1942

³¹ Asela Gunawardana, *Evaluating Machine Learned User Experiences*. Extreme Classification Workshop. NIPS 2015

Challenges

- Long-tail label distributions and zero-shot learning:

Challenges

- Long-tail label distributions and zero-shot learning:
 - ▶ A close relation to the problem of **estimating distributions over large alphabets**.

Challenges


- Long-tail label distributions and zero-shot learning:
 - ▶ A close relation to the problem of **estimating distributions over large alphabets**.
 - ▶ The distribution of label frequencies is often characterized by a **long-tail** for which proper **smoothing** (like add-constant or Good-Turing estimates) or **calibration** techniques (like isotonic regression or domain adaptation) have to be used.

Challenges

- Long-tail label distributions and zero-shot learning:
 - ▶ A close relation to the problem of **estimating distributions over large alphabets**.
 - ▶ The distribution of label frequencies is often characterized by a **long-tail** for which proper **smoothing** (like add-constant or Good-Turing estimates) or **calibration** techniques (like isotonic regression or domain adaptation) have to be used.
 - ▶ In practical applications, learning algorithms run in **rapidly changing environments**: **new labels** may appear during testing/prediction phase (\Rightarrow **zero-shot learning**)

Challenges

- Long-tail label distributions and zero-shot learning:
 - ▶ Frequency of labels in the WikiLSHTC dataset:³²




- ▶ Many labels with only few examples (\Rightarrow one-shot learning).

³² <http://manikvarma.org/downloads/XC/XMLRepository.html>

Challenges

- Long-tail label distributions and zero-shot learning:
 - ▶ Frequency of frequencies for the WikiLSHTC dataset:



- ▶ The missing mass obtained by the Good-Turing estimate: 0.014.

Take-away message

Take-away message

- **Extreme classification:** #examples, #features, #labels

Take-away message

- **Extreme classification:** #examples, #features, #labels
- **Complexity:** time vs. space, training vs. validation vs. **prediction**

Take-away message

- **Extreme classification:** #examples, #features, #labels
- **Complexity:** time vs. space, training vs. validation vs. **prediction**
- **Computational challenges:** fast learning of linear models, compression, MIPS, decision trees, label trees.

Take-away message

- **Extreme classification:** #examples, #features, #labels
- **Complexity:** time vs. space, training vs. validation vs. **prediction**
- **Computational challenges:** fast learning of linear models, compression, MIPS, decision trees, label trees.
- **Statistical challenges:** Is learning possible in the extreme setting?

Take-away message

- **Extreme classification:** #examples, #features, #labels
- **Complexity:** time vs. space, training vs. validation vs. **prediction**
- **Computational challenges:** fast learning of linear models, compression, MIPS, decision trees, label trees.
- **Statistical challenges:** Is learning possible in the extreme setting?
- For more check:

Take-away message

- **Extreme classification:** #examples, #features, #labels
- **Complexity:** time vs. space, training vs. validation vs. **prediction**
- **Computational challenges:** fast learning of linear models, compression, MIPS, decision trees, label trees.
- **Statistical challenges:** Is learning possible in the extreme setting?
- For more check:
 - ▶ <http://www.cs.put.poznan.pl/kdembczynski>

Take-away message

- **Extreme classification:** #examples, #features, #labels
- **Complexity:** time vs. space, training vs. validation vs. **prediction**
- **Computational challenges:** fast learning of linear models, compression, MIPS, decision trees, label trees.
- **Statistical challenges:** Is learning possible in the extreme setting?
- For more check:
 - ▶ <http://www.cs.put.poznan.pl/kdembczynski>
 - ▶ **Code:** <https://github.com/busarobi/XMLC>