

Multi-label learning from batch and streaming data

Jesse Read

Télécom ParisTech
École Polytechnique



Summer School on Mining Big and Complex Data
5 September 2016 — Ohrid, Macedonia

Introduction

$\mathbf{x} =$



Binary classification

$$y \in \{\text{sunset}, \text{non_sunset}\}$$

Introduction

$\mathbf{x} =$



Multi-*class* classification

$y \in \{\text{sunset}, \text{people}, \text{foliage}, \text{beach}, \text{urban}\}$

Introduction

$\mathbf{x} =$



Multi-label classification

$$\begin{aligned} \mathbf{y} &\subseteq \{\text{sunset}, \text{people}, \text{foliage}, \text{beach}, \text{urban}\} \\ &\in \{0, 1\}^5 = [1, 0, 1, 0, 0] \end{aligned}$$

i.e., **multiple** labels per instance instead of a single label.

Introduction

	$K = 2$	$K > 2$
$L = 1$	binary	multi-class
$L > 1$	multi-label	multi-output [†]

[†] also known as multi-target, multi-dimensional.

Figure: For L target variables (labels), each of K values.

- multi-output can be cast to multi-label, just as multi-class can be cast to binary
- set of labels (L) is predefined (contrast to *tagging/keyword assignment*)

Introduction

Table: Academic articles containing the phrase “*multi-label classification*” (Google Scholar, 2016)

year	in text	in title
1996-2000	23	1
2001-2005	188	18
2006-2010	1470	164
2011-2015	5280	629

Single-label vs. Multi-label

Table: Single-label $Y \in \{0, 1\}$

X_1	X_2	X_3	X_4	X_5	Y
1	0.1	3	1	0	0
0	0.9	1	0	1	1
0	0.0	1	1	0	0
1	0.8	2	0	1	1
1	0.0	2	0	1	0
0	0.0	3	1	1	?

Table: Multi-label $Y \subseteq \{\lambda_1, \dots, \lambda_L\}$

X_1	X_2	X_3	X_4	X_5	Y
1	0.1	3	1	0	$\{\lambda_2, \lambda_3\}$
0	0.9	1	0	1	$\{\lambda_1\}$
0	0.0	1	1	0	$\{\lambda_2\}$
1	0.8	2	0	1	$\{\lambda_1, \lambda_4\}$
1	0.0	2	0	1	$\{\lambda_4\}$
0	0.0	3	1	1	?

Single-label vs. Multi-label

Table: Single-label $Y \in \{0, 1\}$

X_1	X_2	X_3	X_4	X_5	Y
1	0.1	3	1	0	0
0	0.9	1	0	1	1
0	0.0	1	1	0	0
1	0.8	2	0	1	1
1	0.0	2	0	1	0
0	0.0	3	1	1	?

Table: Multi-label $[Y_1, \dots, Y_L] \in 2^L$

X_1	X_2	X_3	X_4	X_5	Y_1	Y_2	Y_3	Y_4
1	0.1	3	1	0	0	1	1	0
0	0.9	1	0	1	1	0	0	0
0	0.0	1	1	0	0	1	0	0
1	0.8	2	0	1	1	0	0	1
1	0.0	2	0	1	0	0	0	1
0	0.0	3	1	1	?	?	?	?

Outline

- 1 Introduction
- 2 Applications**
- 3 Methods
- 4 Label Dependence
- 5 Multi-label Classification in Data Streams

Text Categorization and Tag Recommendation

For example, the IMDb dataset: Textual movie **plot summaries** associated with **genres** (labels). Also: Bookmarks, Bibtex, del.icio.us datasets.



The Lord of the Rings: The Fellowship of the Ring (2001) 

PG-13 | 178 min | **Adventure, Fantasy** | 19 December 2001 (USA)

Your rating: ★★★★★★★★ -/10
Ratings: 8.8/10 from 1,110,948 users | Metascore: 92/100
Reviews: 4,988 user | 294 critic | 34 from Metacritic.com

A meek hobbit of the Shire and eight companions set out on a journey to Mount Doom to destroy the One Ring and the dark lord Sauron.

Director: [Peter Jackson](#)
Writers: [J.R.R. Tolkien](#) (novel), [Fran Walsh](#) (screenplay), [2 more credits](#) »
Stars: [Elijah Wood](#), [Ian McKellen](#), [Orlando Bloom](#) | [See full cast and crew](#) »

Text Categorization and Tag Recommendation

For example, the IMDb dataset: Textual movie **plot summaries** associated with **genres** (labels). Also: Bookmarks, Bibtex, del.icio.us datasets.

	<i>abandoned</i>	<i>accident</i>	...	<i>violent</i>	<i>wedding</i>	<i>horror</i>	<i>romance</i>	...	<i>comedy</i>	<i>action</i>
<i>i</i>	X_1	X_2	...	X_{1000}	X_{1001}	Y_1	Y_2	...	Y_{27}	Y_{28}
1	1	0	...	0	1	0	1	...	0	0
2	0	1	...	1	0	1	0	...	0	0
3	0	0	...	0	1	0	1	...	0	0
4	1	1	...	0	1	1	0	...	0	1
5	1	1	...	0	1	0	1	...	0	1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
120919	1	1	...	0	0	0	0	...	0	1

Labelling Images



Images are labelled to indicate

- multiple concepts
- multiple objects
- multiple people

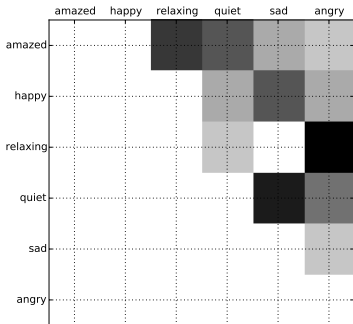
e.g., Associating **Scenes** with **concepts**

\subseteq {beach, sunset, foliage, field, mountain, urban}

Labelling Audio

For example, labelling **music** with **emotions**, **concepts**, etc.

- amazed-surprised
- happy-pleased
- relaxing-calm
- quiet-still
- sad-lonely
- angry-aggressive



Related Tasks

- **multi-output¹ classification**: outputs are nominal

X_1	X_2	X_3	X_4	X_5	rank	gender	group
x_1	x_2	x_3	x_4	x_5	1	M	2
x_1	x_2	x_3	x_4	x_5	4	F	2
x_1	x_2	x_3	x_4	x_5	2	M	1

- **multi-output regression**: outputs are real-valued

X_1	X_2	X_3	X_4	X_5	price	age	percent
x_1	x_2	x_3	x_4	x_5	37.00	25	0.88
x_1	x_2	x_3	x_4	x_5	22.88	22	0.22
x_1	x_2	x_3	x_4	x_5	88.23	11	0.77

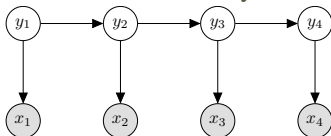
- **label ranking**, i.e., preference learning

$$\lambda_3 \succ \lambda_1 \succ \lambda_4 \succ \dots \succ \lambda_2$$

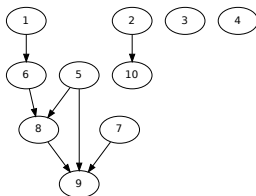
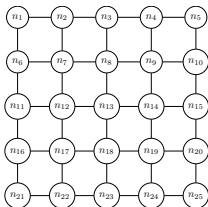
¹aka multi-target, multi-dimensional

Related Areas

- **multi-task learning**: multiple tasks, shared representation
- **sequential learning**: predict across time indices instead of across label indices; each label may have a different input



- **structured output prediction**: assume particular structure among outputs, e.g., pixels, hierarchy



Streaming Multi-label Data

Many advanced applications must deal with **data streams**:

- Data arrives **continuously**, potentially infinitely
- Prediction must be made **immediately**
- Expect **concept drift**

For example,

- Demand prediction
- Intrusion detection
- Pollution detection

Demand Prediction

Outputs (labels) represent the demand at multiple points.

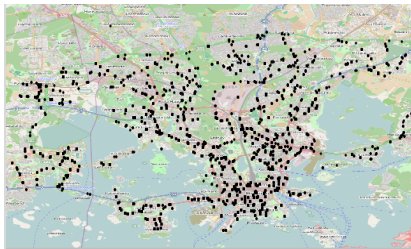


Figure: Stops in the greater Helsinki region. The *Kutsuplus* taxi service could be called to any of these.

We are interested in predicting, for each label $[y_1, \dots, y_L]$,

$p(y_j = 1 | \mathbf{x})$ • probability of demand at j -th node

Localization and Tracking

Outputs represent points in space which may contain an object ($y_j = 1$) or not ($y_j = 0$). Observations are given as \mathbf{x} .

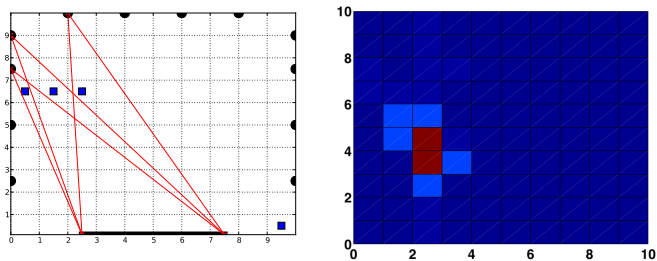


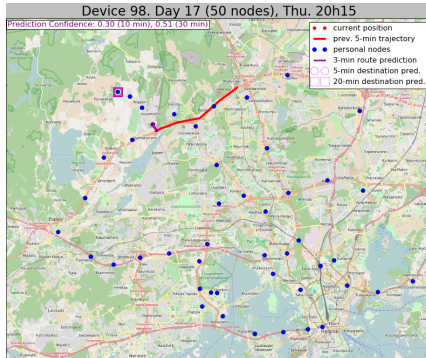
Figure: Modelled on a real-world scenario; a room with a single light source and a number of light sensors.

We are interested in predicting, for each label $[y_1, \dots, y_L]$,

$$y_j = 1 \quad \bullet \text{ if } j\text{-th tile occupied}$$

Route/Destination Forecasting

Personal nodes of a traveller and predicted trajectory



L ● number of geographic points of interest

\mathbf{x} ● observed data (e.g., GPS, sensor activity, time of day)

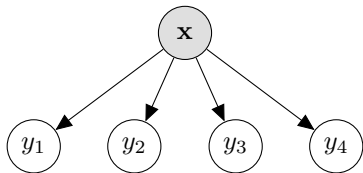
$p(y_j = 1|\mathbf{x})$ ● probability an object is 1 present at the j -th node

$\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$ ● training data

Outline

- 1 Introduction
- 2 Applications
- 3 Methods**
- 4 Label Dependence
- 5 Multi-label Classification in Data Streams

Multi-label Classification



$$\hat{y}_j = h_j(\mathbf{x}) = \underset{y_j \in \{0,1\}}{\operatorname{argmax}} p(y_j|\mathbf{x}) \quad \bullet \text{ for index, } j = 1, \dots, L$$

and then,

$$\begin{aligned} \hat{\mathbf{y}} &= \mathbf{h}(\mathbf{x}) = [\hat{y}_1, \dots, \hat{y}_4] \\ &= \left[\underset{y_1 \in \{0,1\}}{\operatorname{argmax}} p(y_1|\mathbf{x}), \dots, \underset{y_4 \in \{0,1\}}{\operatorname{argmax}} p(y_4|\mathbf{x}) \right] \\ &= [f_1(\mathbf{x}), \dots, f_4(\mathbf{x})] = f(\mathbf{W}^\top \mathbf{x}) \end{aligned}$$

This is the **Binary Relevance** method (BR).

BR Transformation

- ① Transform dataset ...

\mathbf{X}	Y_1	Y_2	Y_3	Y_4
$\mathbf{x}^{(1)}$	0	1	1	0
$\mathbf{x}^{(2)}$	1	0	0	0
$\mathbf{x}^{(3)}$	0	1	0	0
$\mathbf{x}^{(4)}$	1	0	0	1
$\mathbf{x}^{(5)}$	0	0	0	1

...into L separate binary problems (one for each label)

\mathbf{X}	Y_1	\mathbf{X}	Y_2	\mathbf{X}	Y_3	\mathbf{X}	Y_4
$\mathbf{x}^{(1)}$	0	$\mathbf{x}^{(1)}$	1	$\mathbf{x}^{(1)}$	1	$\mathbf{x}^{(1)}$	0
$\mathbf{x}^{(2)}$	1	$\mathbf{x}^{(2)}$	0	$\mathbf{x}^{(2)}$	0	$\mathbf{x}^{(2)}$	0
$\mathbf{x}^{(3)}$	0	$\mathbf{x}^{(3)}$	1	$\mathbf{x}^{(3)}$	0	$\mathbf{x}^{(3)}$	0
$\mathbf{x}^{(4)}$	1	$\mathbf{x}^{(4)}$	0	$\mathbf{x}^{(4)}$	0	$\mathbf{x}^{(4)}$	1
$\mathbf{x}^{(5)}$	0	$\mathbf{x}^{(5)}$	0	$\mathbf{x}^{(5)}$	0	$\mathbf{x}^{(5)}$	1

- ② and **train** with any off-the-shelf binary **base classifier**.

Why Not Binary Relevance?

BR ignores **label dependence**, i.e.,

$$p(\mathbf{y}|\mathbf{x}) = \prod_{j=1}^L p(y_j|\mathbf{x})$$

which may not always hold!

Example (Film Genre Classification)

$$p(y_{\text{romance}}|\mathbf{x}) \neq p(y_{\text{romance}}|\mathbf{x}, y_{\text{horror}})$$

Why Not Binary Relevance?

BR ignores **label dependence**, i.e.,

$$p(\mathbf{y}|\mathbf{x}) = \prod_{j=1}^L p(y_j|\mathbf{x})$$

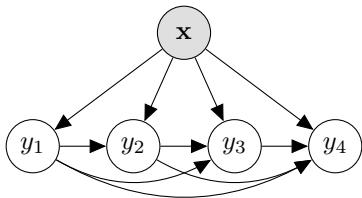
which may not always hold!

Table: Average **predictive performance** (5 fold CV, EXACT MATCH)

	L	BR	MCC
Music	6	0.30	0.37
Scene	6	0.54	0.68
Yeast	14	0.14	0.23
Genbase	27	0.94	0.96
Medical	45	0.58	0.62
Enron	53	0.07	0.09
Reuters	101	0.29	0.37

Classifier Chains

Modelling label dependence,



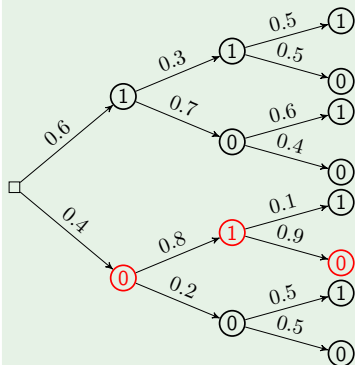
$$p(\mathbf{y}|\mathbf{x}) = p(y_1|\mathbf{x}) \prod_{j=2}^L p(y_j|\mathbf{x}, y_1, \dots, y_{j-1})$$

and,

$$\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y} \in \{0,1\}^L} p(\mathbf{y}|\mathbf{x})$$

Bayes Optimal CC

Example



① $p(\mathbf{y} = [0, 0, 0]) = 0.040$

② $p(\mathbf{y} = [0, 0, 1]) = 0.040$

③ $p(\mathbf{y} = [0, 1, 0]) = 0.288$

④ ...

⑥ $p(\mathbf{y} = [1, 0, 1]) = 0.252$

⑦ ...

⑧ $p(\mathbf{y} = [1, 1, 1]) = 0.090$

return $\operatorname{argmax}_{\mathbf{y}} p(\mathbf{y}|\tilde{\mathbf{x}})$

- Search space of $\{0, 1\}^L$ paths is too much

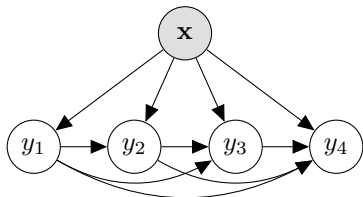
CC Transformation

Similar to BR: make L binary problems, but **include previous predictions as feature attributes**,

\mathbf{X}	Y_1	\mathbf{X}	Y_1	Y_2	\mathbf{X}	Y_1	Y_2	Y_3	\mathbf{X}	Y_1	Y_3	Y_3	Y_4
$\mathbf{x}^{(1)}$	0	$\mathbf{x}^{(1)}$	0	1	$\mathbf{x}^{(1)}$	0	1	1	$\mathbf{x}^{(1)}$	0	1	1	0
$\mathbf{x}^{(2)}$	1	$\mathbf{x}^{(2)}$	1	0	$\mathbf{x}^{(2)}$	1	0	0	$\mathbf{x}^{(2)}$	1	0	0	0
$\mathbf{x}^{(3)}$	0	$\mathbf{x}^{(3)}$	0	1	$\mathbf{x}^{(3)}$	0	1	0	$\mathbf{x}^{(3)}$	0	1	0	0
$\mathbf{x}^{(4)}$	1	$\mathbf{x}^{(4)}$	1	0	$\mathbf{x}^{(4)}$	1	0	0	$\mathbf{x}^{(4)}$	1	0	0	1
$\mathbf{x}^{(5)}$	0	$\mathbf{x}^{(5)}$	0	0	$\mathbf{x}^{(5)}$	0	0	0	$\mathbf{x}^{(5)}$	0	0	0	1

and, again, apply any classifier (not necessarily a probabilistic one)!

Greedy CC



L classifiers for L labels. For test instance $\tilde{\mathbf{x}}$, classify

① $\hat{y}_1 = h_1(\tilde{\mathbf{x}})$

② $\hat{y}_2 = h_2(\tilde{\mathbf{x}}, \hat{y}_1)$

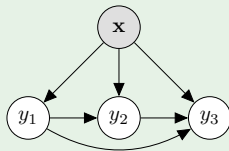
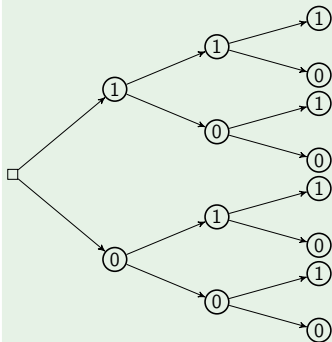
③ $\hat{y}_3 = h_3(\tilde{\mathbf{x}}, \hat{y}_1, \hat{y}_2)$

④ $\hat{y}_4 = h_4(\tilde{\mathbf{x}}, \hat{y}_1, \hat{y}_2, \hat{y}_3)$

and return

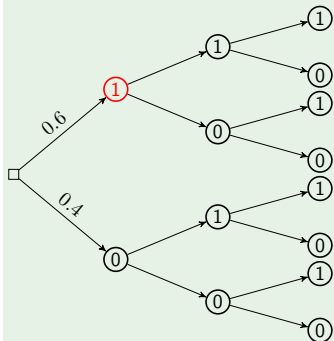
$$\hat{\mathbf{y}} = [\hat{y}_1, \dots, \hat{y}_L]$$

Example

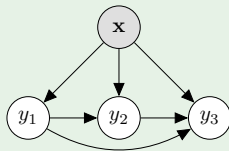


$$\hat{\mathbf{y}} = \mathbf{h}(\tilde{\mathbf{x}}) = [?, ?, ?]$$

Example

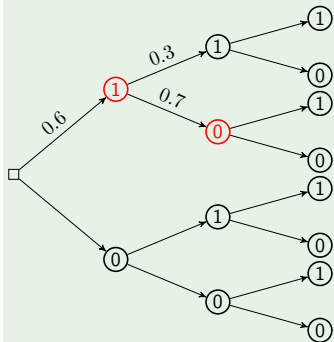


$$\hat{\mathbf{y}} = \mathbf{h}(\tilde{\mathbf{x}}) = [1, ?, ?]$$

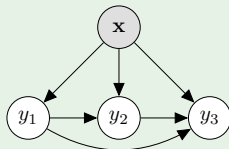


$$\textcircled{1} \hat{y}_1 = h_1(\tilde{\mathbf{x}}) = \operatorname{argmax}_{y_1} p(y_1|\tilde{\mathbf{x}}) = 1$$

Example

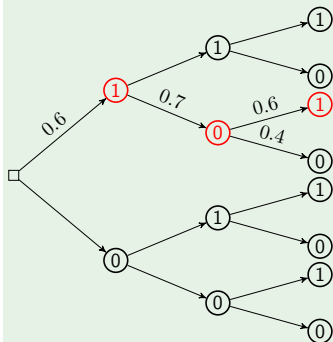


$$\hat{\mathbf{y}} = \mathbf{h}(\tilde{\mathbf{x}}) = [1, \mathbf{0}, ?]$$

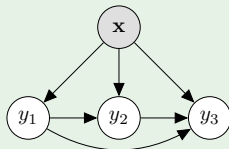


- 1 $\hat{y}_1 = h_1(\tilde{\mathbf{x}}) = \operatorname{argmax}_{y_1} p(y_1|\tilde{\mathbf{x}}) = 1$
- 2 $\hat{y}_2 = h_2(\tilde{\mathbf{x}}, \hat{y}_1) = \dots = 0$

Example

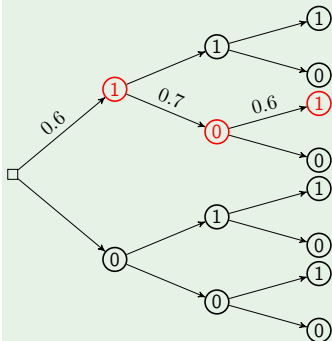


$$\hat{\mathbf{y}} = \mathbf{h}(\tilde{\mathbf{x}}) = [1, 0, 1]$$

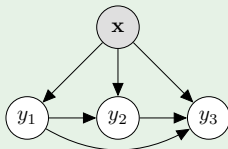


- 1 $\hat{y}_1 = h_1(\tilde{\mathbf{x}}) = \operatorname{argmax}_{y_1} p(y_1|\tilde{\mathbf{x}}) = 1$
- 2 $\hat{y}_2 = h_2(\tilde{\mathbf{x}}, \hat{y}_1) = \dots = 0$
- 3 $\hat{y}_3 = h_3(\tilde{\mathbf{x}}, \hat{y}_1, \hat{y}_2) = \dots = 1$

Example



$$\hat{\mathbf{y}} = \mathbf{h}(\tilde{\mathbf{x}}) = [1, 0, 1]$$

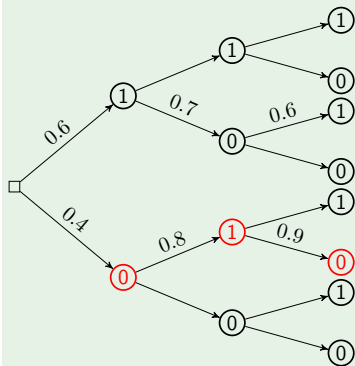


- 1 $\hat{y}_1 = h_1(\tilde{\mathbf{x}}) = \operatorname{argmax}_{y_1} p(y_1|\tilde{\mathbf{x}}) = 1$
- 2 $\hat{y}_2 = h_2(\tilde{\mathbf{x}}, \hat{y}_1) = \dots = 0$
- 3 $\hat{y}_3 = h_3(\tilde{\mathbf{x}}, \hat{y}_1, \hat{y}_2) = \dots = 1$

- Improves over BR; similar build time (if $L < D$);
- able to use any off-the-shelf classifier for h_j ; parallelizable
- But, **errors may be propagated down the chain**

Monte-Carlo search for CC

Example



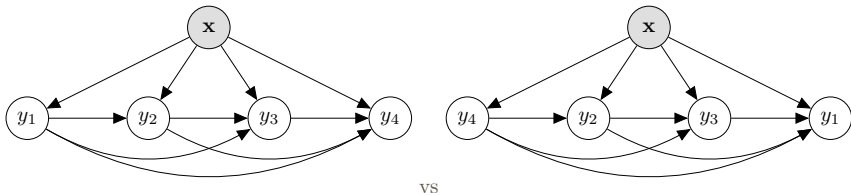
Sample T times ...

- $p([1, 0, 1]) = 0.6 \cdot 0.7 \cdot 0.6 = 0.252$
- $p([0, 1, 0]) = 0.4 \cdot 0.8 \cdot 0.9 = 0.288$

return $\operatorname{argmax}_{\mathbf{y}_t} p(\mathbf{y}_t | \mathbf{x})$

- **Tractable**, with **similar accuracy** to (Bayes Optimal) PCC
- Can use other search algorithms, e.g., beam search

Does Label-*order* Matter?



In theory, models are equivalent, since

$$p(\mathbf{y}|\mathbf{x}) = p(y_1|\mathbf{x})p(y_2|y_1, \mathbf{x}) = p(y_2|\mathbf{x})p(y_1|y_2, \mathbf{x})$$

- **but** we are estimating p from **finite** and **noisy** data; thus

$$\hat{p}(y_1|\mathbf{x})\hat{p}(y_2|y_1, \mathbf{x}) \neq \hat{p}(y_2|\mathbf{x})\hat{p}(y_1|y_2, \mathbf{x})$$

- and in the greedy case,

$$\hat{p}(y_2|y_1, \mathbf{x}) \approx \hat{p}(y_2|\hat{y}_1, \mathbf{x}) = \hat{p}(y_2|y_1 = \underset{y_1}{\operatorname{argmax}} \hat{p}(y_1|\mathbf{x})|\mathbf{x})$$

Does Label-*order* Matter?

In theory, models are equivalent, since

$$p(\mathbf{y}|\mathbf{x}) = p(y_1|\mathbf{x})p(y_2|y_1, \mathbf{x}) = p(y_2|\mathbf{x})p(y_1|y_2, \mathbf{x})$$

- **but** we are estimating p from **finite** and **noisy** data; thus

$$\hat{p}(y_1|\mathbf{x})\hat{p}(y_2|y_1, \mathbf{x}) \neq \hat{p}(y_2|\mathbf{x})\hat{p}(y_1|y_2, \mathbf{x})$$

- and in the greedy case,

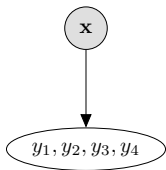
$$\hat{p}(y_2|y_1, \mathbf{x}) \approx \hat{p}(y_2|\hat{y}_1, \mathbf{x}) = \hat{p}(y_2|y_1 = \underset{y_1}{\operatorname{argmax}} \hat{p}(y_1|\mathbf{x})|\mathbf{x})$$

The approximations cause **high variance** on account of **error propagation**. We can

- ① can reduce variance with an *ensemble of classifier chains*
- ② we can **search space of chain orders** (huge space, but a little search makes a difference)

Label Powerset Method (LP)

One multi-class problem (taking many values),



$$\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} p(\mathbf{y}|\mathbf{x}) \quad \bullet \text{ where } |\mathcal{Y}| \leq \{0, 1\}^L$$

- Each value is a label vector, 2^L in total, but
- typically, only the **occurrences of the training set**.
- (in practice, $|\mathcal{Y}| \leq \text{size of training set}$, and $|\mathcal{Y}| \ll 2^L$)

Label Powerset Method (LP)

- ① Transform dataset ...

\mathbf{X}	Y_1	Y_2	Y_3	Y_4
$\mathbf{x}^{(1)}$	0	1	1	0
$\mathbf{x}^{(2)}$	1	0	0	0
$\mathbf{x}^{(3)}$	0	1	1	0
$\mathbf{x}^{(4)}$	1	0	0	1
$\mathbf{x}^{(5)}$	0	0	0	1

... into a multi-*class* problem, taking 2^L possible values:

\mathbf{X}	$Y \in 2^L$
$\mathbf{x}^{(1)}$	0110
$\mathbf{x}^{(2)}$	1000
$\mathbf{x}^{(3)}$	0110
$\mathbf{x}^{(4)}$	1001
$\mathbf{x}^{(5)}$	0001

- ② ... and train any off-the-shelf multi-*class* classifier.

Issues with LP

- **complexity** (up to 2^L combinations)
- **imbalance**: few examples per class label
- **overfitting**: how to predict new value?

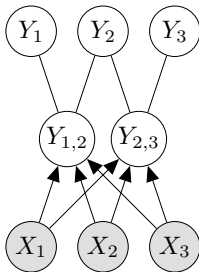
Example

In the Enron dataset, 44% of labelsets are unique (to a single training example or test instance). In del.icio.us dataset, 98% are unique.

Meta Labels

Improving the label-powerset approach:

- **decomposition** of label set into M subsets of size k ($k < L$)
- **pruning**, such that, e.g., $Y_{1,2} \in \{[0, 0], [0, 1], [1, 1]\}$
- combine together with **random subspace** method with a **voting** scheme



Meta Labels

Improving the label-powerset approach:

- **decomposition** of label set into M subsets of size k ($k < L$)
- **pruning**, such that, e.g., $Y_{1,2} \in \{[0, 0], [0, 1], [1, 1]\}$
- combine together with **random subspace** method with a **voting** scheme

Method	Inference Complexity
Label Powerset	$O(2^L \cdot D)$
Pruned Sets	$O(P \cdot D)$
Decomposition / RAKEL	$O(M \cdot 2^k \cdot D)$
Meta Labels	$O(M \cdot P \cdot D')$

where $P < 2^L$ and $P < 2^k, D' < D$.

Summary of Methods

Two views of a multi-label problem of L labels:

- ① L **binary problems**
- ② a **multi-class problem** with up to 2^L classes

Problem Transformation:

- ① Transform data into subproblems (binary or multi-class)
- ② Apply some off-the-shelf base classifier

or, **Algorithm Adaptation:**

- ① Take a suitable single-label classifier (k NN, neural networks, decision trees ...)
- ② Adapt it (if necessary) for multi-label classification

Outline

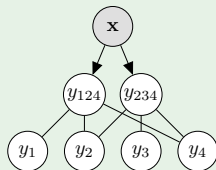
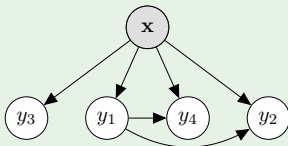
- 1 Introduction
- 2 Applications
- 3 Methods
- 4 Label Dependence**
- 5 Multi-label Classification in Data Streams

Label Dependence in MLC

Common approach: Present methods to

- ① measure **label dependence**
 - ② find a **structure** that best represents this
- and then apply classifiers, compare results to BR.

Example



- Link particular labels (nodes) together (CC-based methods)
- Form particular label subsets (LP-based methods)

Label Dependence in MLC

Common approach: Present methods to

- ① measure **label dependence**
- ② find a **structure** that best represents this

and then apply classifiers, compare results to BR.



Problem

Measuring label dependence is expensive, models built on it often **do not improve over models built on random dependence!**



Problem

For some metrics (such as Hamming-loss / label accuracy), knowledge of **label dependence is theoretically unnecessary!**

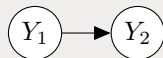
Marginal label dependence

Marginal dependence

When the joint is **not** the product of the marginals, i.e.,

$$p(y_2) \neq p(y_2|y_1)$$

$$p(y_1)p(y_2) \neq p(y_1, y_2)$$



- Estimate from co-occurrence frequencies in training data

Marginal label dependence

Example

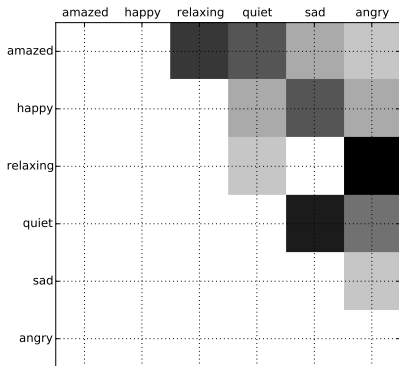


Figure: Music dataset - Mutual Information

Marginal label dependence

Example

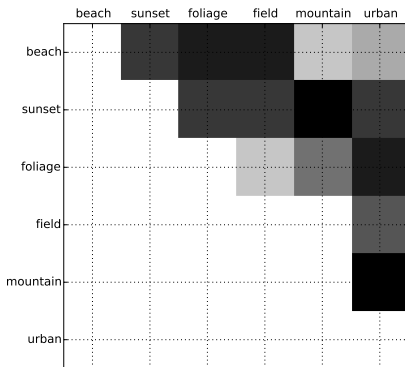


Figure: Scene dataset - Mutual Information

Marginal label dependence

Marginal dependence

When the joint is **not** the product of the marginals, i.e.,

$$p(y_2) \neq p(y_2|y_1)$$

$$p(y_1)p(y_2) \neq p(y_1, y_2)$$



- Estimate from co-occurrence frequencies in training data

Used for **regularization/constraints**:

- 1 $\hat{\mathbf{y}} = h(\mathbf{x})$ makes a prediction
- 2 $\hat{\mathbf{y}}' = g(\hat{\mathbf{y}})$ regularizes the prediction

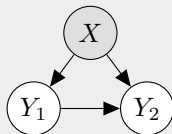
Conditional label dependence

But at classification time, we condition on the input!

Conditional dependence

...conditioned on input observation \mathbf{x} .

$$p(y_2|y_1, \mathbf{x}) \neq p(y_2|\mathbf{x})$$



- Have to build and measure models

Indication of conditional dependence if

- the performance of LP/CC exceeds that of BR
- errors among the binary models are correlated

But what does this mean?

Conditional label dependence

But at classification time, we condition on the input!

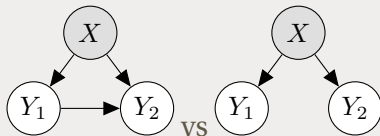
Conditional *independence*

...conditioned on input observation \mathbf{x} .

For example,

$$p(y_2) \neq p(y_2|y_1)$$

$$\text{but } p(y_2|\mathbf{x}) = p(y_2|, y_1, \mathbf{x})$$



- Have to build and measure models

Indication of conditional dependence if

- the performance of LP/CC exceeds that of BR
- errors among the binary models are correlated

But what does this mean?

The LOGICAL Problem

Example (The LOGICAL Toy Problem)

X_1	X_2	OR Y_1	AND Y_2	XOR Y_3
0	0	0	0	0
1	0	1	0	1
0	1	1	0	1
1	1	1	1	0

- Each label is a logical operation (independent of the others!)

The LOGICAL Problem

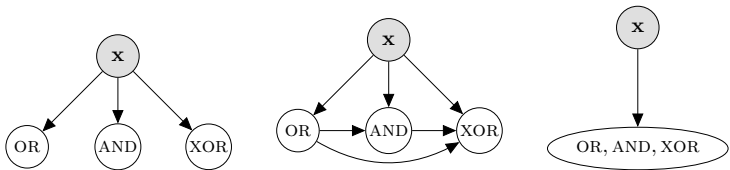


Figure: BR (left), CC (middle), LP (right)

Table: The LOGICAL problem, base classifier logistic regression.

Metric	BR	CC	LP
HAMMING SCORE	0.83	1.00	1.00
EXACT MATCH	0.50	1.00	1.00

- Dependence is introduced by an inadequate model!
- **Dependence depends on the model.**

The LOGICAL Problem

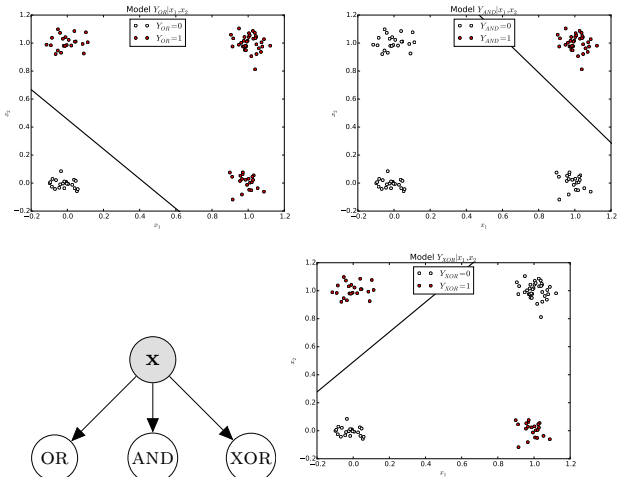


Figure: **Binary Relevance** (BR): linear decision boundary (solid line, estimated with logistic regression) not viable for Y_{XOR} label

Solution via Structure

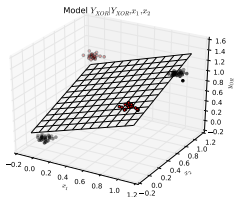
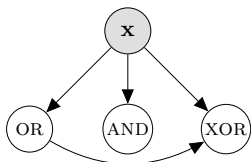
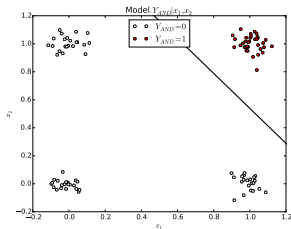
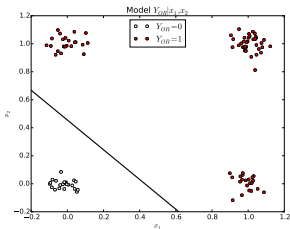


Figure: **Classifier chains** (CC): linear model now applicable to Y_{XOR}

Solution via Multi-class Decomposition

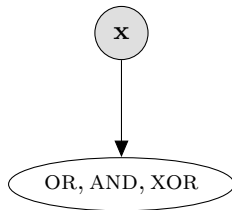
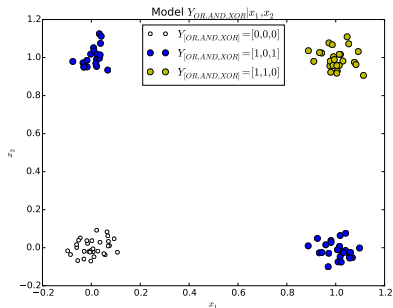


Figure: **Label Powerset** (LP): solvable with one-vs-one multi-class decomposition for any (e.g., linear) base classifier.

Solution via Con. Independence

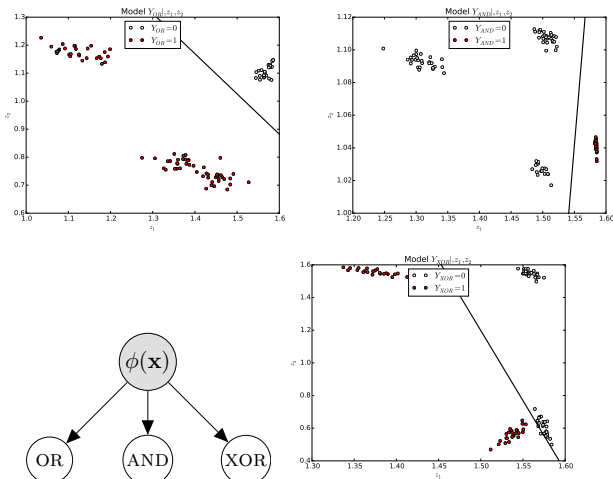


Figure: Solution via **latent structure** (e.g., random RBF) to new input space \mathbf{z} ; creating independence: $p(y_{XOR} | \mathbf{z}, y_{OR}, y_{AND}) \approx p(y_{XOR} | \mathbf{z})$.

Solution via Suitable Base-classifier

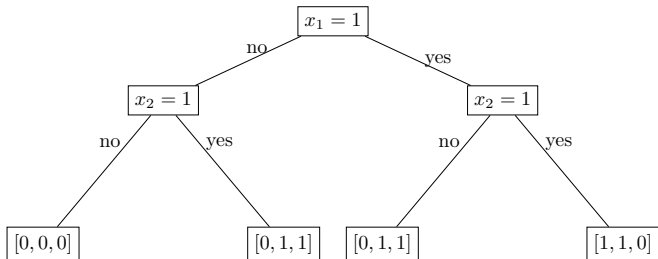


Figure: Solution via non-linear classifier (e.g., [Decision Tree](#)). Leaves hold examples, where $\mathbf{y} = [y_{\text{AND}}, y_{\text{OR}}, y_{\text{XOR}}]$

Detecting Dependence

Conditional label dependence and the choice of base model are inseparable.

$$y_j = h_j(\mathbf{x}) + \epsilon_j$$

$$y_k = h_k(\mathbf{x}) + \epsilon_k$$

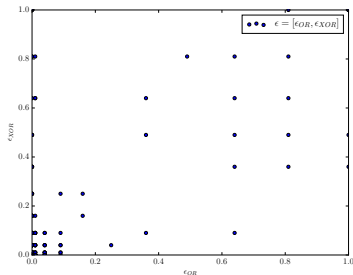
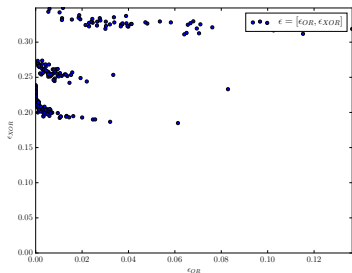


Figure: Errors from logistic regression (left) and decision tree (right).

A fresh look at Problem Transformation

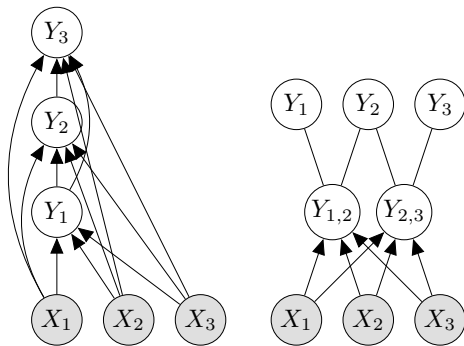


Figure: Standard methods can be viewed as ('deep'/cascaded) basis functions on the label space.

Label Dependence: Summary

- Marginal dependence for regularization
- Conditional dependence
 - ... depends on the model
 - ... may be introduced
- Should consider together:
 - base classifier
 - label structure
 - inner-layer structure
- An open problem
- Much existing research is relevant (latent-variable models, neural networks, deep learning, ...)

Outline

- 1 Introduction
- 2 Applications
- 3 Methods
- 4 Label Dependence
- 5 Multi-label Classification in Data Streams**

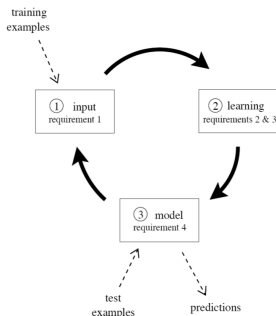
Classification in Data Streams

Setting:

- sequence is **potentially infinite**
- **high speed** of arrival
- stream is **one-way**

Implications

- work in limited memory
- adapt to **concept drift**

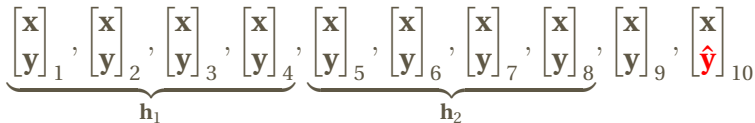


Multi-label Streams Methods

- ① Batch-incremental Ensemble
- ② Problem transformation with an incremental base learner
- ③ Multi-label k NN
- ④ Multi-label incremental decision trees
- ⑤ Neural networks

Batch-Incremental Ensemble

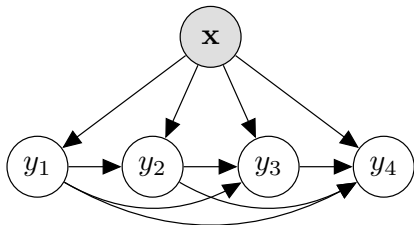
Build regular multi-label models on batches/windows of instances (typically in a [weighted] ensemble).



- A common approach in the literature, and can be surprisingly effective
- Free choice of base classifier (e.g., C4.5, SVM)
- What batch size to use?
 - Too small = models insufficient
 - Too large = slow to adapt
 - Too many batches = too slow

Problem Transformation with Incremental Base Learner

Use an incremental learner (Naive Bayes, SGD, Hoeffding trees) with any problem transformation method (BR, LP, CC, ...)

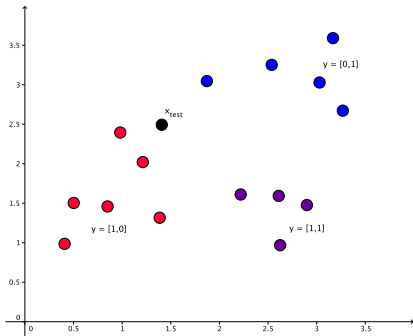


- Simple implementation,
- Risk of overfitting (e.g., with classifier chains)
- Concept drift may invalidate structure
- Limited choice of base learner (must be incremental)

Multi-label k NN

Maintain a dynamic buffer of instances, compare each test instance $\tilde{\mathbf{x}}$ to the k neighbouring instances,

$$\hat{y}_j = \begin{cases} 1 & \left(\frac{1}{k} \sum_{i|\mathbf{x}^{(i)} \in \text{Ne}(\tilde{\mathbf{x}})} y_j^{(i)} > 0.5 \right) \\ 0 & \text{otherwise} \end{cases}$$



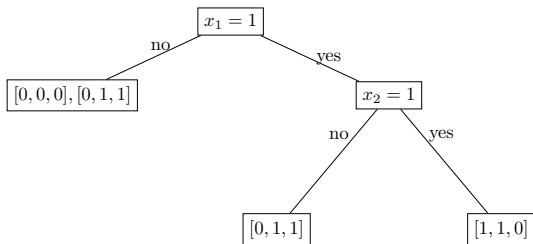
- efficient wrt L
- ... but not wrt D
- **limited buffer size,**
- not suitable for all problems

ML Incremental Decision Trees

- A small sample can suffice to choose a splitting attribute (Hoeffding bound gives guarantees)
- As in regular tree, with modified splitting criteria, e.g.,

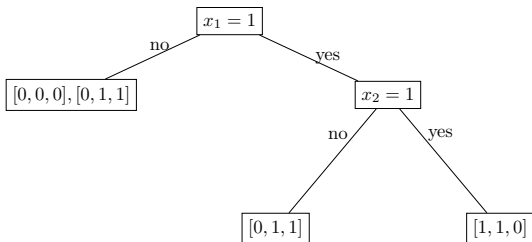
$$H_{\text{ML}}(S) = - \sum_{j=1}^L \sum_{c \in \{0,1\}} P(y_j = c) \log_2 P(y_j = c)$$

- Examples with *multiple labels* collect at the leaves.



ML Incremental Decision Trees

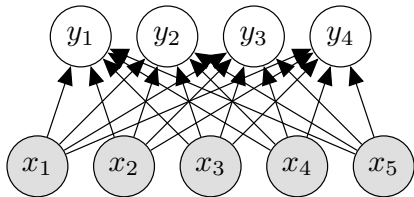
- *Fast*, and usually *competitive*,
- But tree may grow very **conservatively**,
- ... and need to replace it (or part of it) when concept changes.



- Place multi-label classifiers at the leaves of the tree
- and wrap it in an ensemble.

Neural networks

Each label is a node. Trained with SGD



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

$$\mathbf{g} = \nabla E(\mathbf{W})$$

$$W_{j,k}^{(t+1)} = W_{j,k}^{(t)} + \lambda g_{j,k}$$

- Can be applied natively
- One layer = BR, should use **hidden layers** to model label dependence / improve performance
- Hyper-parameter tuning can be tedious
- Relatively *poor performance* in empirical comparisons on standard data streams (improving now with recent advances in SGD, more common use of basis expansion)

Multi-label Data Streams: Issues

- Overfitting
- Class imbalance
- **Multi-dimensional concept drift**
- **Labelled examples difficult to obtain** (semi-supervised)
- Dynamic label set
- Time dependence

Multi-label Concept Drift

Consider the **relative frequencies** of the j -th and k -th labels:

$$\begin{bmatrix} p_j & p_{j,k} \\ & p_k \end{bmatrix}$$

(if **marginal independence** then $p_{j,k} = p_j p_k$).

Possible **drift**:

- p_j increases (j -th label relatively **more frequent**)
- p_j and p_k both decrease (**label cardinality** decreasing)
- $p_{j,k}$ changes relative to $p_j p_k$ (change in marginal **dependence** relation between the labels)

Multi-label Concept Drift

And when **conditioned on input \mathbf{x}** , we consider the **relative frequencies/values** of the j -th and k -th **errors**:

$$\begin{bmatrix} \epsilon_j & \epsilon_{j,k} \\ & \epsilon_k \end{bmatrix}$$

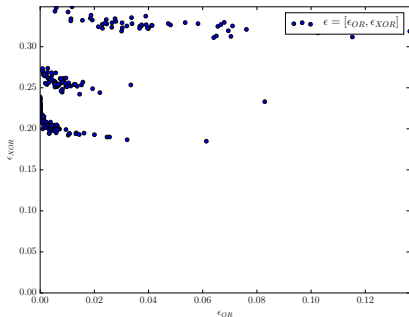
(if **conditional independence**, then $\epsilon_{j,k} \approx \epsilon_j \cdot \epsilon_k$).

Possible drift:

- ϵ_j increases (**more errors** on j -th label)
- ϵ_j and ϵ_k both increase (**more errors**)
- $\epsilon_{j,k}$ changes relative to ϵ_j, ϵ_k (change in **conditional dependence** relation)

Example

Recall the distribution of errors



This shape may change over time – and structures may need to be adjusted to cope

Dealing with Concept Drift

Possible approaches

- Just **ignore it** – batch models must be replaced anyway, k NN and SGD adapt; in other cases can use weighted ensembles/fading factor
- **Monitor a predictive performance statistic** with a **change detector** (e.g., window based-detection, ADWIN) and reset models
- **Monitor the distribution** with a **change detector** (e.g., window based, KL divergence) and reset/recalibrate models

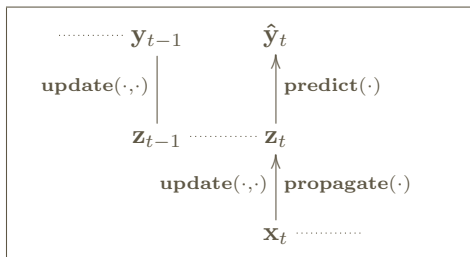
(similar to single-labelled data, except more complex measurement)

Dealing with Unlabelled Instances

- Ignore instances with no label
- Use **active learning** to get good labels
- Use predicted labels (**self-training**)
- Use an **unsupervised process** for example **clustering**, **latent-variable representations**.

Dealing with Unlabelled Instances

- Use an **unsupervised process** for example **clustering**, **latent-variable representations**.
 - 1 $\mathbf{z}_t = g(\mathbf{x}_t)$
 - 2 $\hat{\mathbf{y}}_t = h(\mathbf{z}_t)$
 - 3 update g with $(\mathbf{x}_t, \mathbf{z}_t)$
 - 4 update h with $(\mathbf{z}_{t-1}, \mathbf{y}_{t-1})$ (if \mathbf{y}_{t-1} is available)



Can also be as one single model

Summary

- Multi-label classification is an active area of research, relevant to many real-world problems
- Methods that deal appropriately with label dependence can achieve significant gains over a naive approach
- Many multi-label problems come in the form of a data stream, incurring particular challenges

Multi-label learning from batch and streaming data

Jesse Read

Télécom ParisTech
École Polytechnique



Summer School on Mining Big and Complex Data
5 September 2016 — Ohrid, Macedonia