LTN Examples and Code

Luciano Serafini Michael Spranger

FBK-IRST, Trento, Italy
Sony Computer Science Laboratories Inc., Tokyo, Japan

August 29, 2018

Logic Tensor Networks

- https://github.com/logictensornetworks/
- git clone https://github.com/logictensornetworks/logictensornetworks.git

Classification

Domain A set D of points in the $[0, 1]^2$ square;

Predicate A unary predicate A (class)

Supervisions

1. **Positive examples** a set of points in D, which are known to be instances of A;
2. **Negative examples** a set of points in D, which are known not to be instances of A;

Task For all the other point in D determine if they are instances of A.

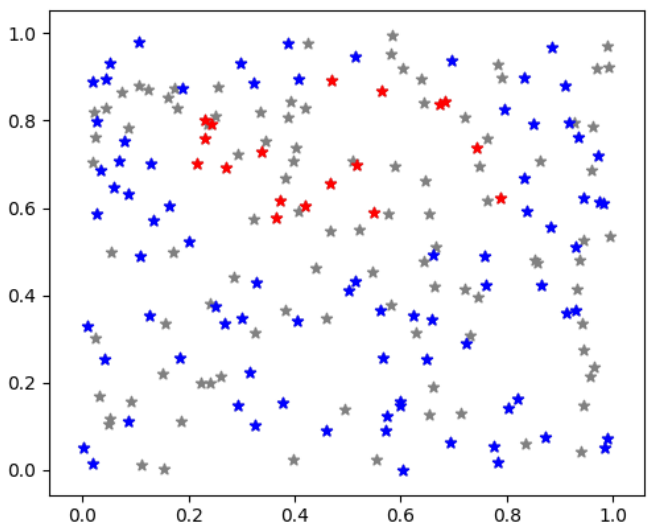

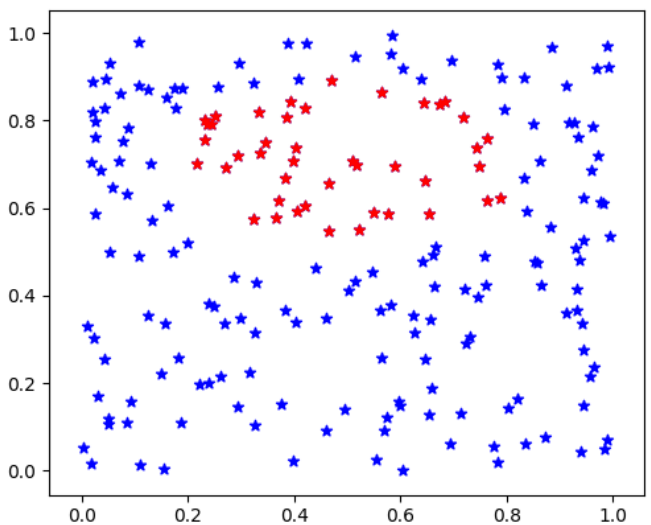

Classification in LTN

$$D = \text{random subset of } [0,1]^2$$

$$P = \{p \in D \mid p \text{ is a positive example of } A\}$$

$$N = \{p \in D \mid p \text{ is a negative example of } A\}$$

$$\forall x \in P : A(x)$$

$$\forall x \in N : \neg A(x)$$

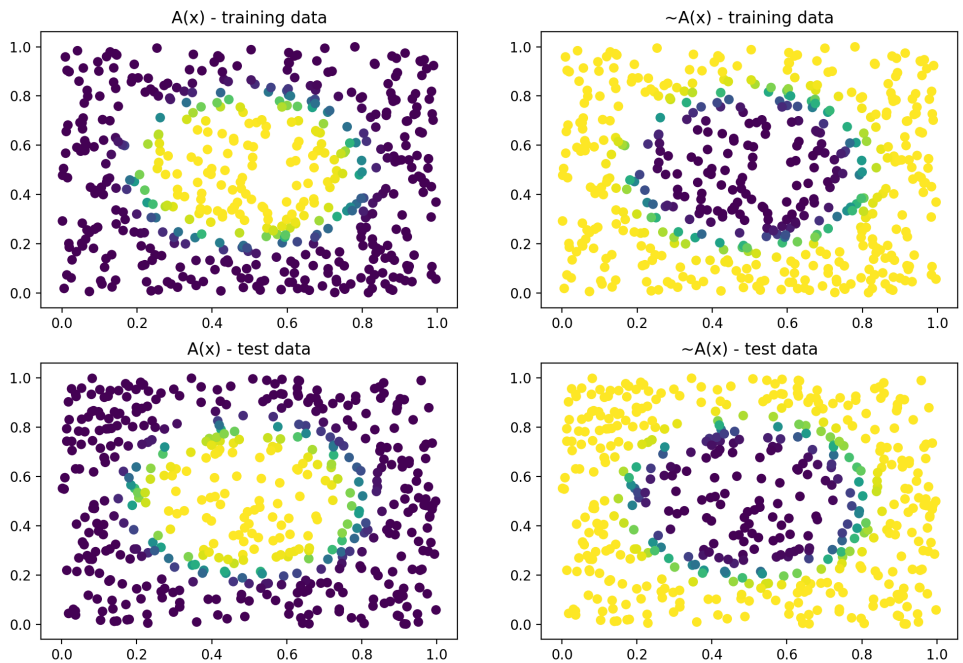

Regression

Domain A set $D = [0, 1]$

Function A unary function $f : D \to D$

Supervisions A set of supervision pairs $S = \{\langle x_1, y_1 \rangle, \ldots, \langle x_n, y_n \rangle\}$ such that $y_i = f(x_i)$

Task For all the other points $x \in D$ predict $f(x)$

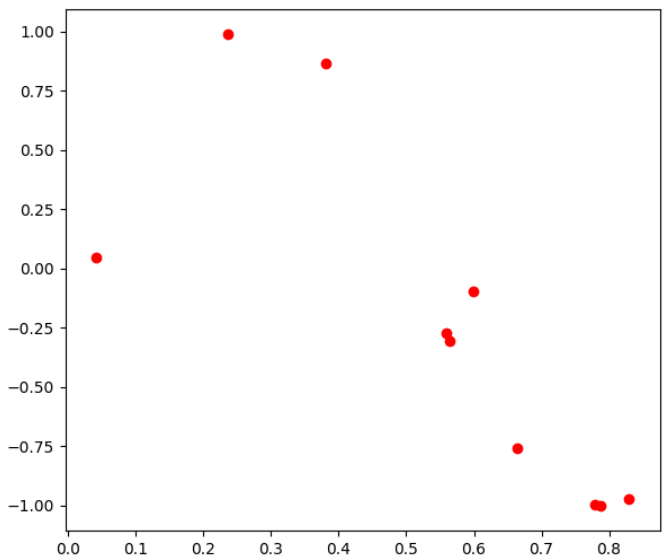

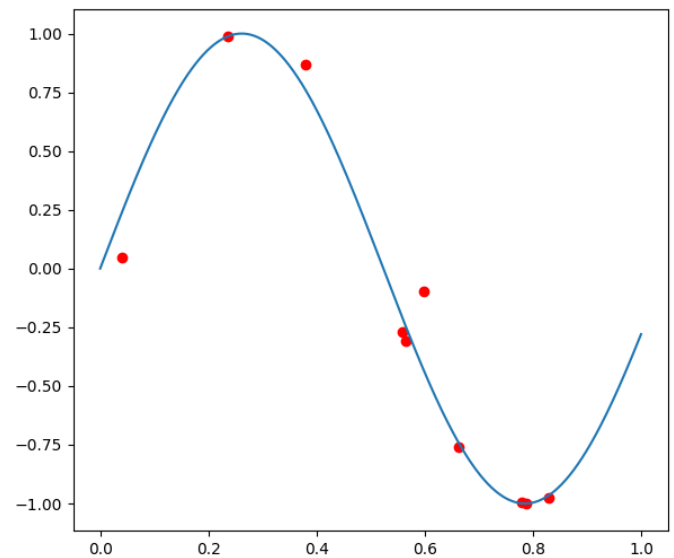

Regression in LTN

$$
\begin{aligned}
D &= \text{a uniform sampling of } [0,1] \\
S &= \{\langle x_1, y_1\rangle, \ldots, \langle x_n, y_n\rangle\}
\end{aligned}
$$

$$\forall \langle x, y\rangle \in S : f(x) = y$$

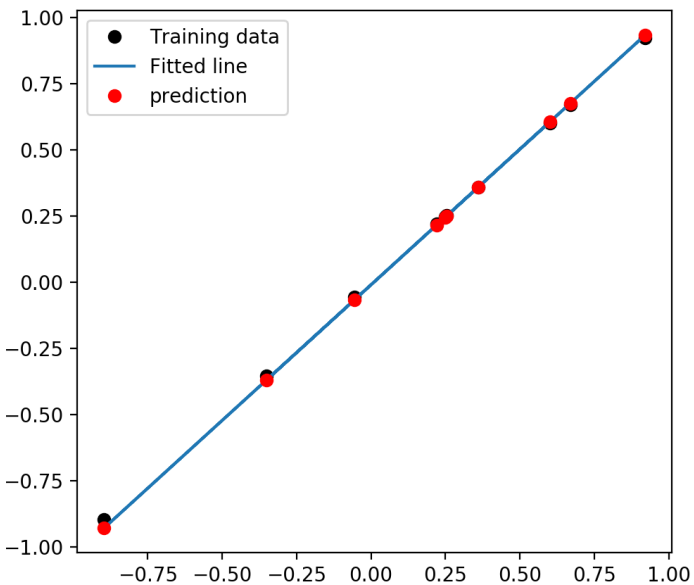

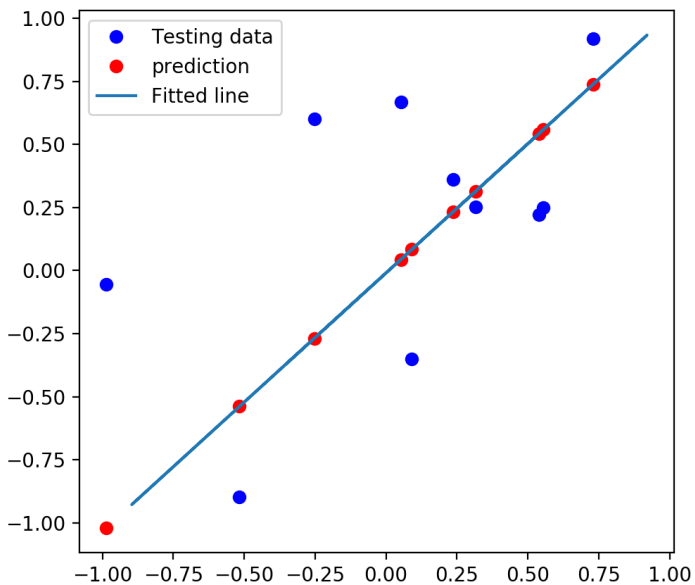

Multi-Label Classification with Label Constraints

Domain A set D of points in the $[0,1]^2$ square;

Predicate A set unary predicate $A_1, \ldots, A_n$ (class/labels)

Supervisions

1. **Positive examples for** A_i a set of points in D, which are known to be instances of A_i;
2. **Negative examples for** A_i a set of points in D, which are known not to be instances of A_i;

Label constraints

- **Subset constraint** if x is labelled with A_i then it is labelled with A_j;
- **Disjoint constraint** x is labelled with A_i iff then x is labelled with $\neg A_j$ and viceversa;

Task For all the points in D determine if they are labelled with A_i or $\neg A_i$ (Notice that it is possible that x is labelled neither with A_i nor $\neg A_i$)

Multi-Label Classification with Label Constraints in LTN

$$\forall x \in A_i^+ : A_i(x)$$
$$\forall x \in A_i^- : \neg A_i(x)$$
$$\forall x \in D : A_i(x) \rightarrow A_j(x)$$
$$\forall x \in D : A_i(x) \leftrightarrow \neg A_j(x) \wedge \neg A_i(x) \leftrightarrow A_j(x)$$

Notice that FOL allows to express more general constraints between labels, as for instance

$$\forall x \in D : (A(x) \wedge B(x)) \rightarrow (C(x) \vee D(x))$$

Multi-Label Classification with Label Constraints in LTN

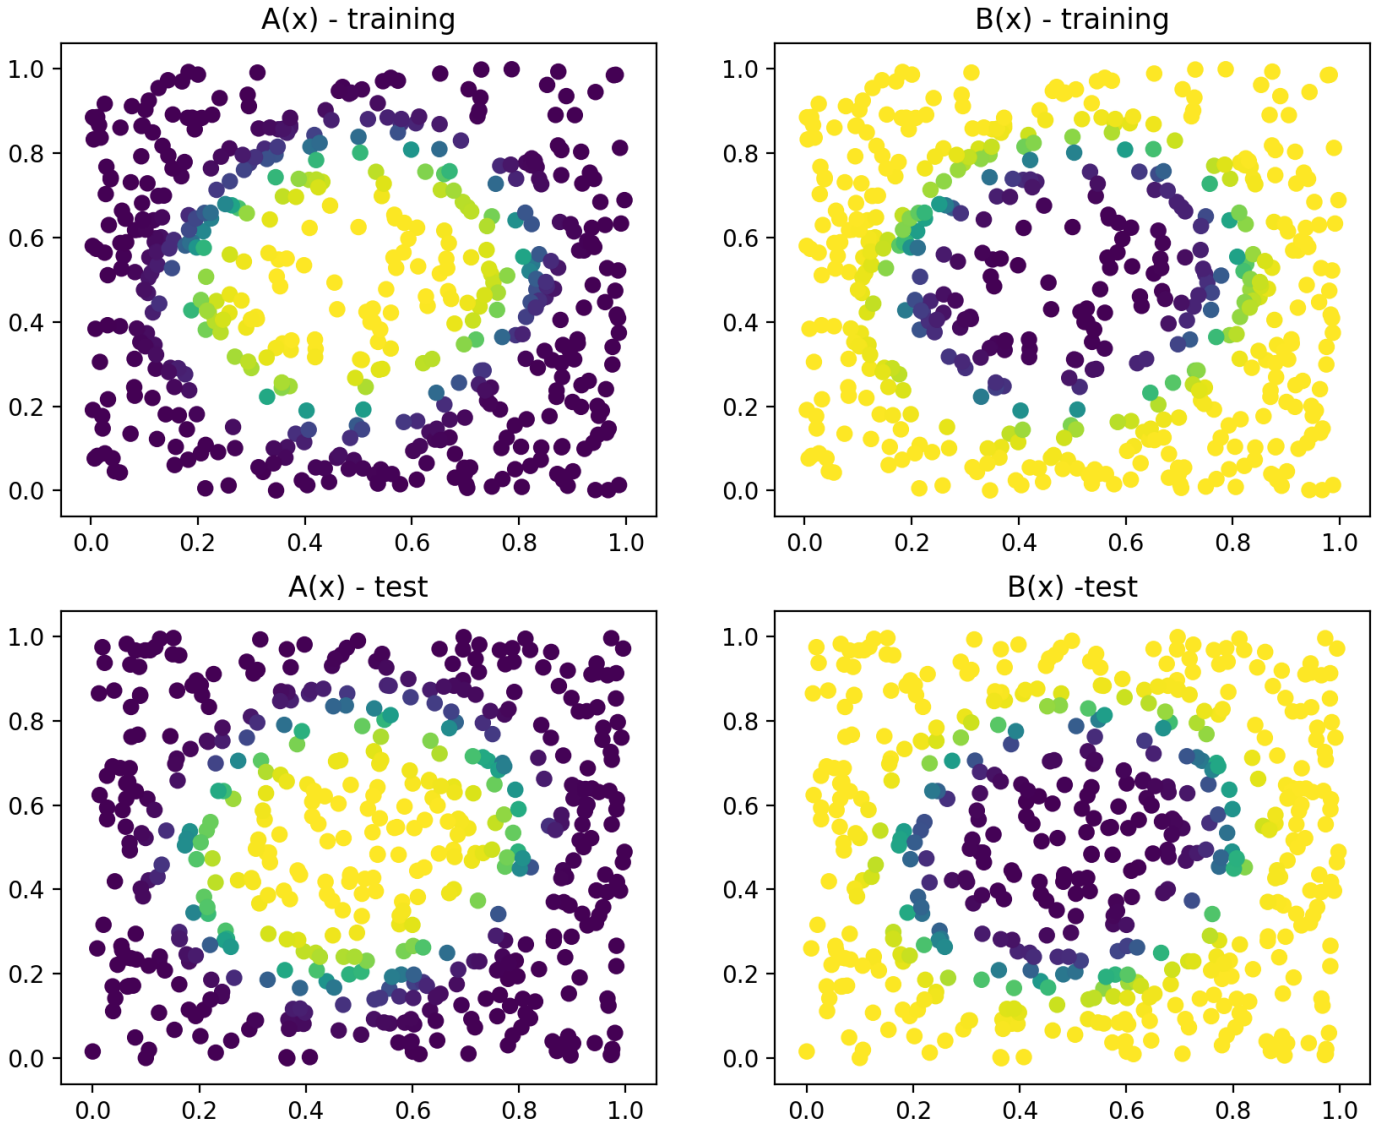

Unsupervised learning - Clustering

Domain A set $D \subset [0, 1]$

Cluster labels A set of cluster labels $C_1, \ldots, C_n$

Task Put each point in some cluster C_i minimizing intra-cluster distance and maximizing inter-cluster distance

Clustering Constraints

- Every point belongs to a claster

$$\forall x : C_1(x) \vee \cdots \vee C_n(x)$$

- Every point cannot belong to two clusters

$$\forall x : \neg(C_i(x) \wedge C_j(x) \qquad \text{for } i \neq j \in \{1, \ldots, n\}$$

- Clusters are not empty

$$\exists x : C_i(x) \qquad \text{for } i \in \{1, \ldots, n\}$$

- Close points should belong to the same cluster.

$$\forall x, y : Close(x, y) \rightarrow (C_i(x) \leftrightarrow C_i(y)) \text{for } i \in \{1, \ldots, n\}$$

Exemplary results

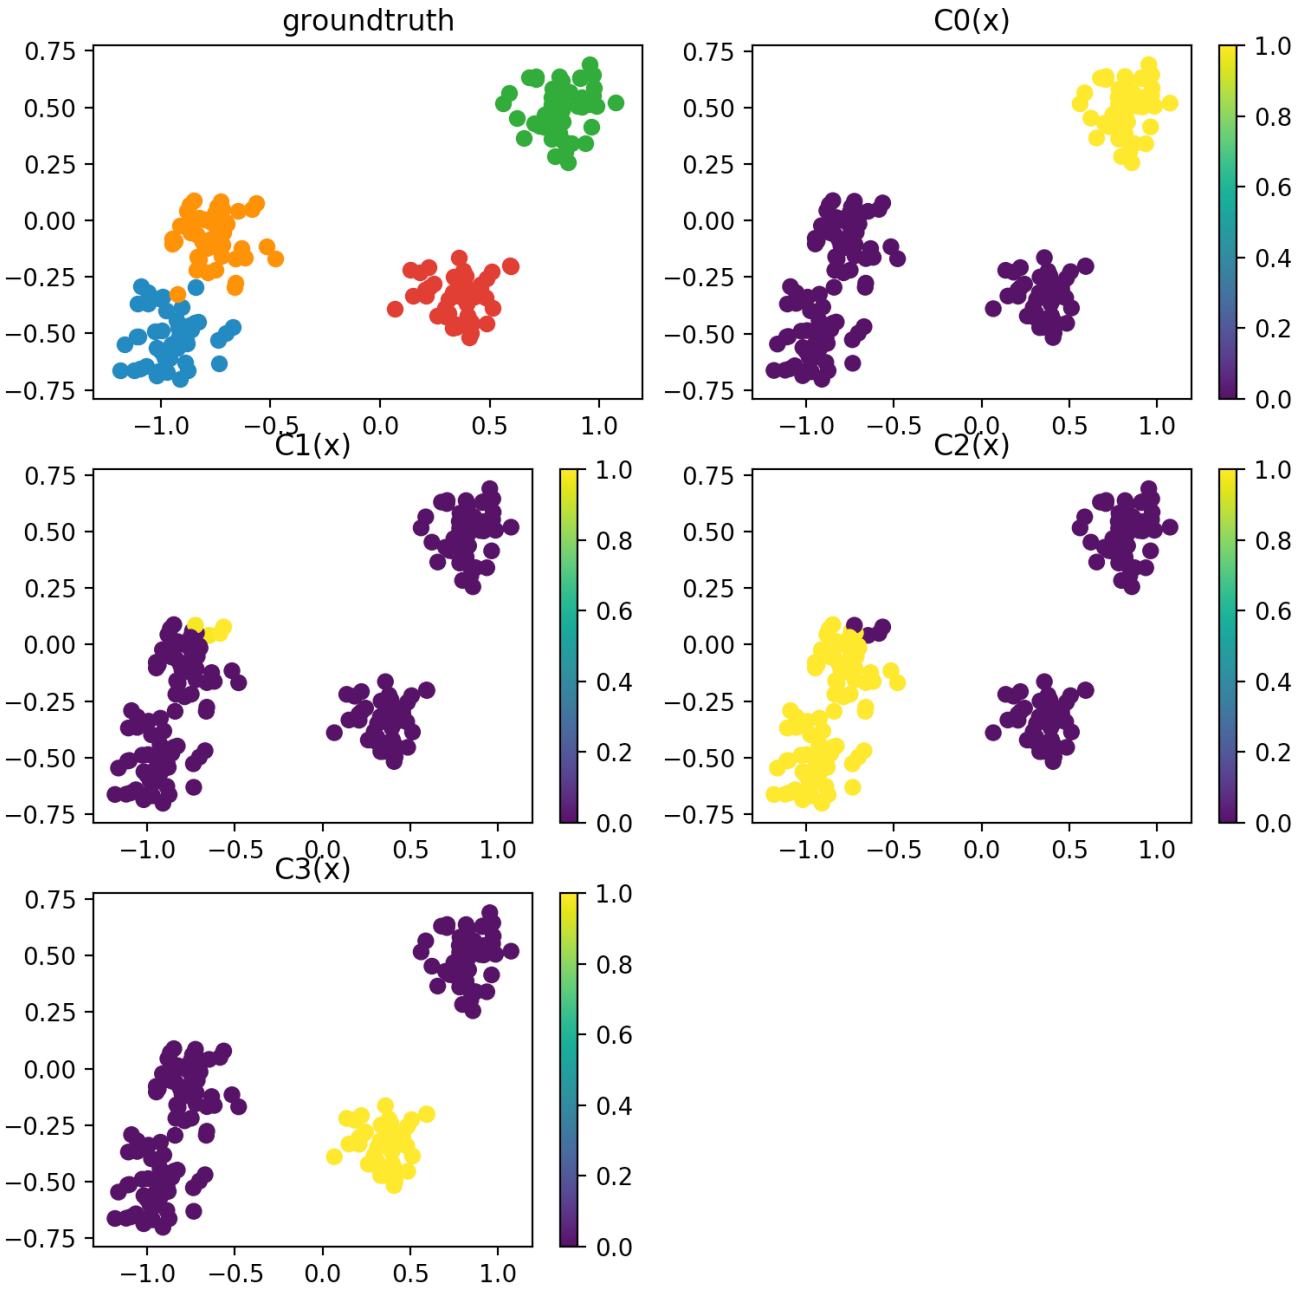

Relations

Domain A set D of points in the $[0, 1]^2$ square;

Predicate A set 2-ary relational predicates $R_1, \ldots, R_n$

Supervision

1. **Positive examples for** $R_i(x, y)$ a set of points in $D \times D$, which are known to be instances of R_i;
2. **Negative examples for** $R_i(x, y)$ a set of points in $D \times D$, which are known not to be instances of R_i;

Constraints

- **Symmetry** if $\forall x, y \in D \times D : R(x, y) \rightarrow R(y, x)$
- **Subset** if $\forall (x, y) \in D \times D : R(x, y) \rightarrow R'(x, y)$
- ...

Relations

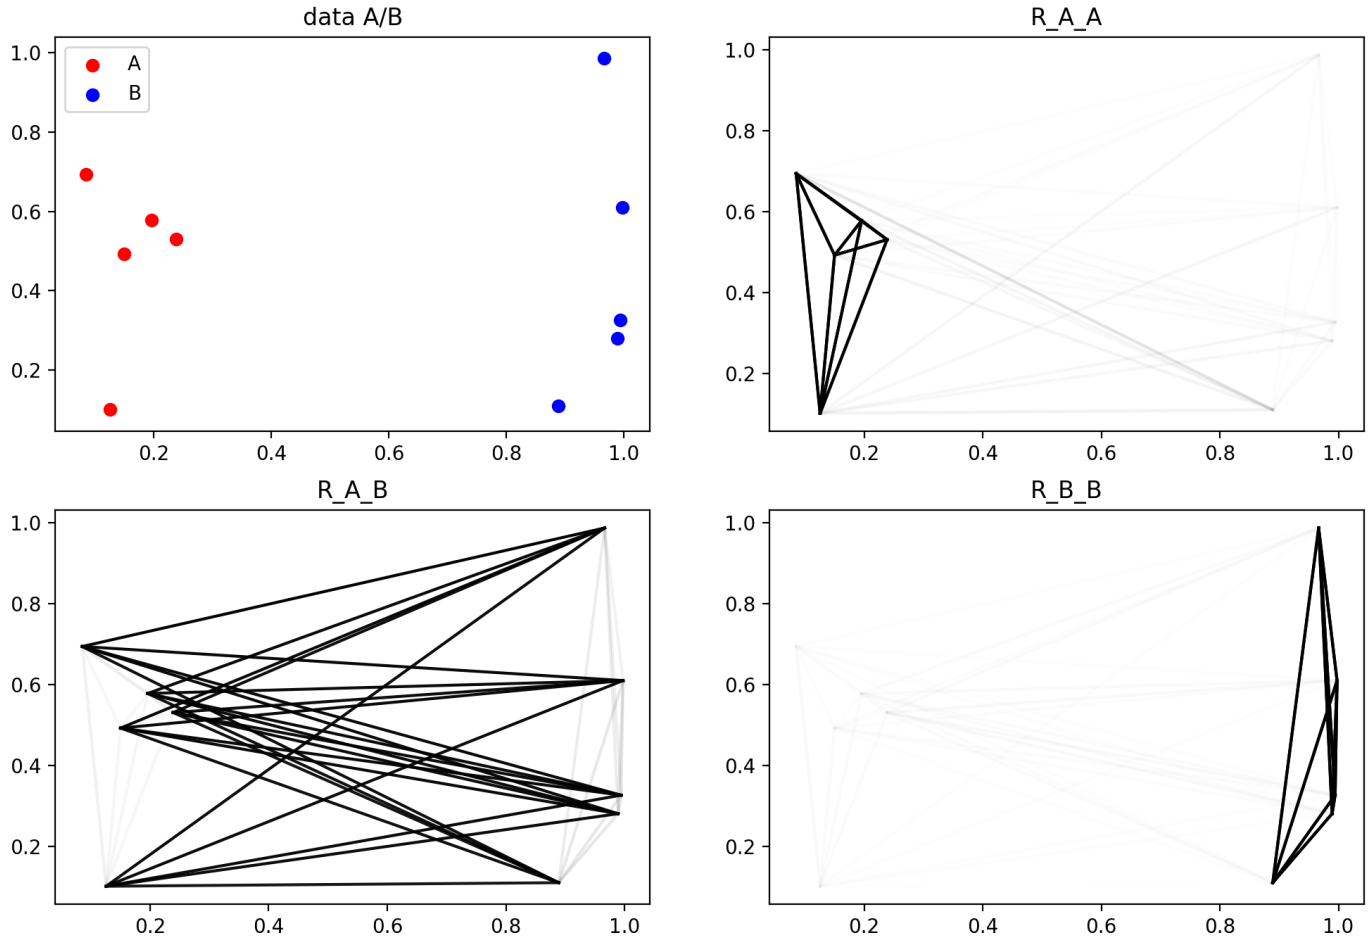

Learning spatial relations

Domain

Our domain is composed of 2d rectangles on a plane; We are interested in the following 6 spatial relations.

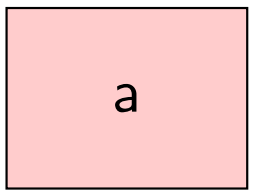

Learning spatial relations

Domain

Our domain is composed of 2d rectangles on a plane; We are interested in the following 6 spatial relations.

- *a* is on the left of *b*

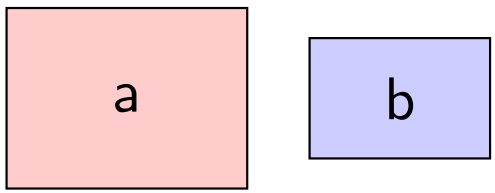

Learning spatial relations

Domain

Our domain is composed of 2d rectangles on a plane; We are interested in the following 6 spatial relations.

- *a* is on the left of *b*
- *a* is on the right of *b*

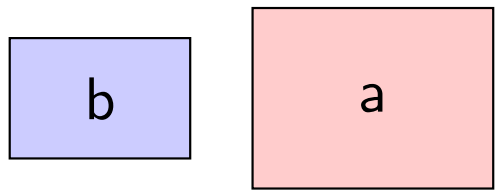

Learning spatial relations

Domain

Our domain is composed of 2d rectangles on a plane; We are interested in the following 6 spatial relations.

- a is on the left of b
- a is on the right of b
- a is above of b

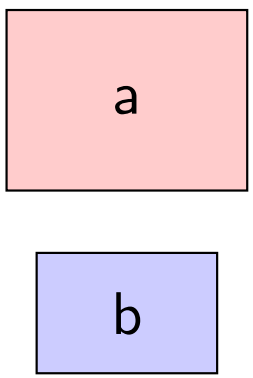

Learning spatial relations

Domain

Our domain is composed of 2d rectangles on a plane; We are interested in the following 6 spatial relations.

- a is on the left of b
- a is on the right of b
- a is above of b
- a is below of b

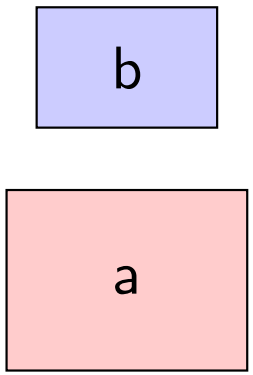

Learning spatial relations

Domain

Our domain is composed of 2d rectangles on a plane; We are interested in the following 6 spatial relations.

- a is on the left of b
- a is on the right of b
- a is above of b
- a is below of b
- a is contained in b

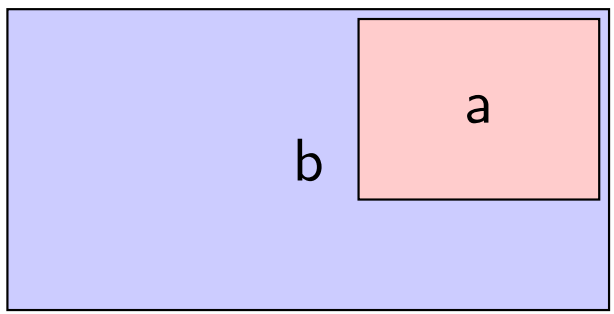

Learning spatial relations

Domain

Our domain is composed of 2d rectangles on a plane; We are interested in the following 6 spatial relations.

- a is on the left of b
- a is on the right of b
- a is above of b
- a is below of b
- a is contained in b
- a contains b

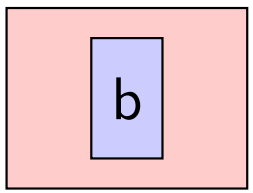

Learning spatial relations

Problem

Given some some examples of pairs of rectangles for each specific relation, and some background knowledge about them as for instance:

- left is the inverse of right
- an object cannot be at the same time on the left and on the right of another object
- if an object a is contained in an object b and b is on the left of c, then a is on the left of c

we want to be able to predict if two randomly generated rectangles, are in one of the 6 spatial relation.

Learning spatial relations

Domain representation

Every rectangle is represented with 4 real numbers,

$$\langle x, y, w, h \rangle$$

encoding the coordinates of the bottom-left corner, and the width and the height

The language

The language is constituted by 6 binary relations

$$\mathit{left}(x, y),\ \mathit{right}(x, y),\ \mathit{above}(x, y),$$
$$\mathit{below}(x, y),\ \mathit{contains}(x, y),\ \mathit{in}(x, y)$$

Learning spatial relations

The constraints

- a set of positive examples for each spatial relation: $left(a, b)$, $right(c, d)$, $above(e, f)$, $below(h, i)$, $contains(j, k)$, $in(l, m)$, ...
- axioms about spatial relations:
  - $\forall x, y : left(x, y) \rightarrow \neg left(y, x)$;
  - $\forall x, y : left(x, y) \rightarrow right(y, x)$;
  - $\forall x, y, z : in(x, y) \wedge left(y, z) \rightarrow left(x, z)$
  - ...

Learning spatial relations

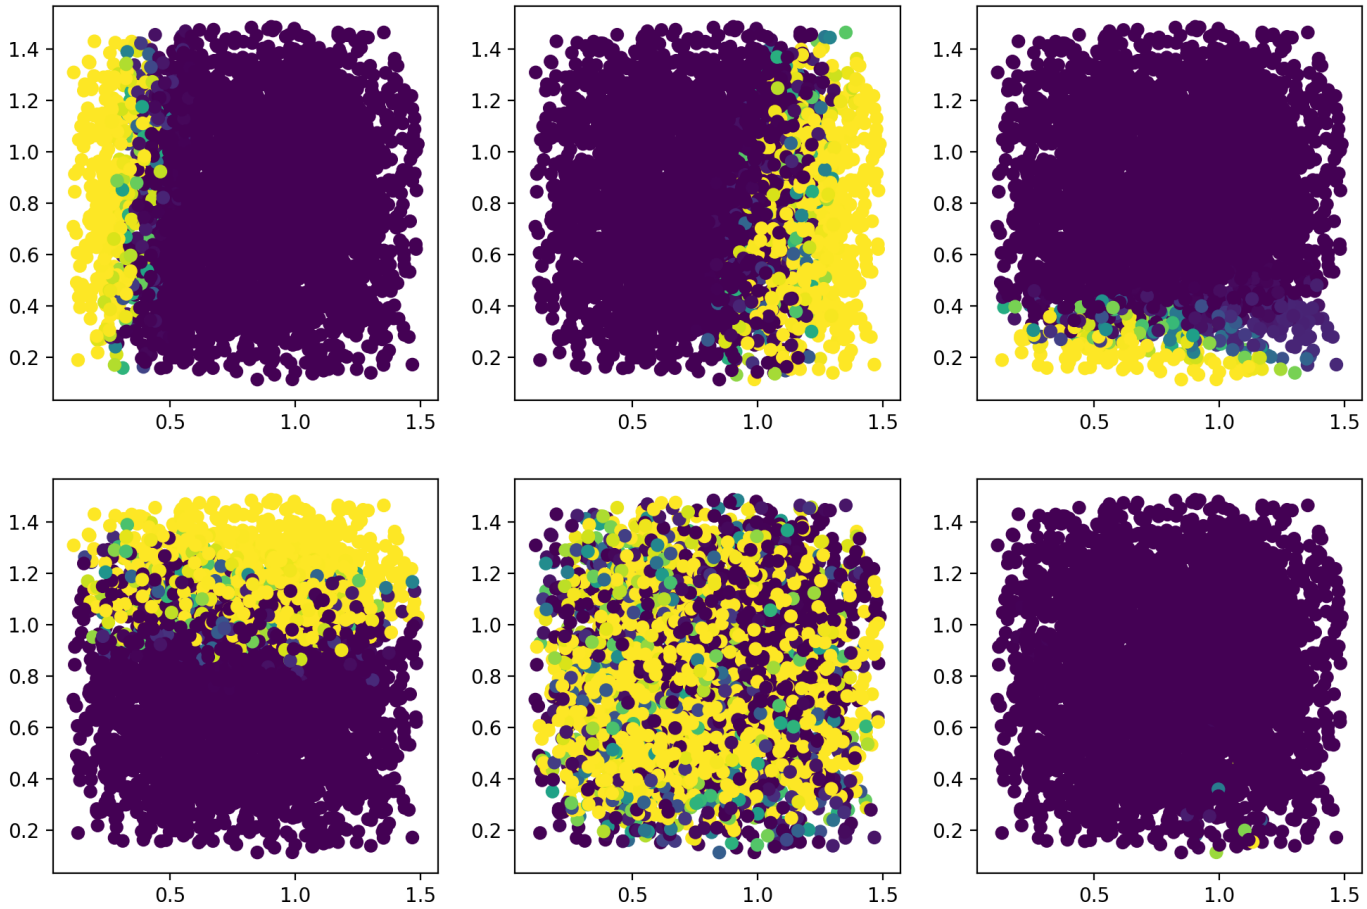

Statistical relational learning

Domain: Smoking-Friends-Cancer, [?]

- Of two groups of people $\{a, b, ..., h\}$ and $\{i, j, ..., n\}$; we know if each of them smokes and the friendship relation within each group;
- for the first group we also know who has a cancer;
- we know that cancer depends on smoking
- and that smoking habits depend on the friendship relation

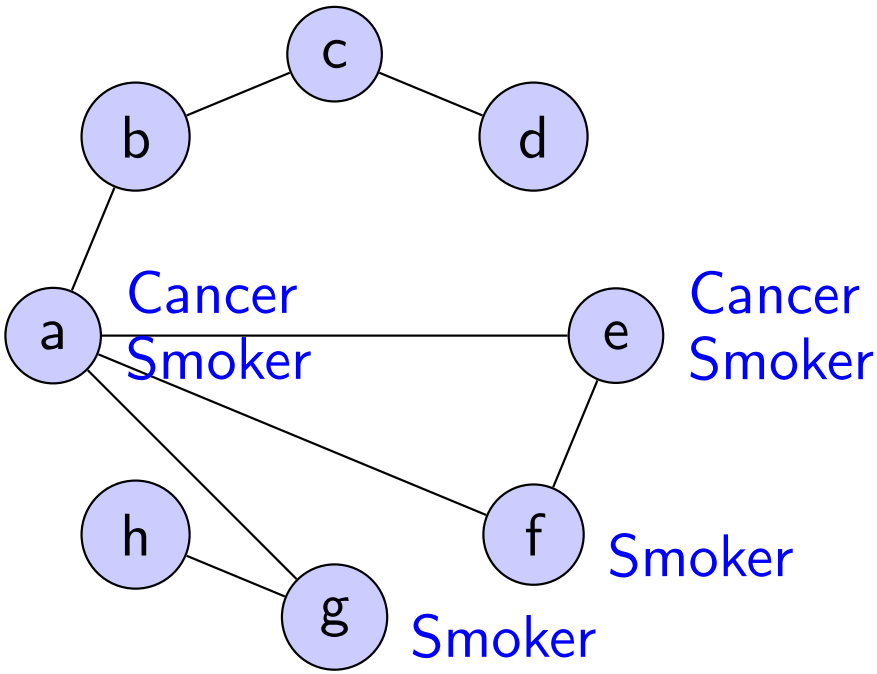

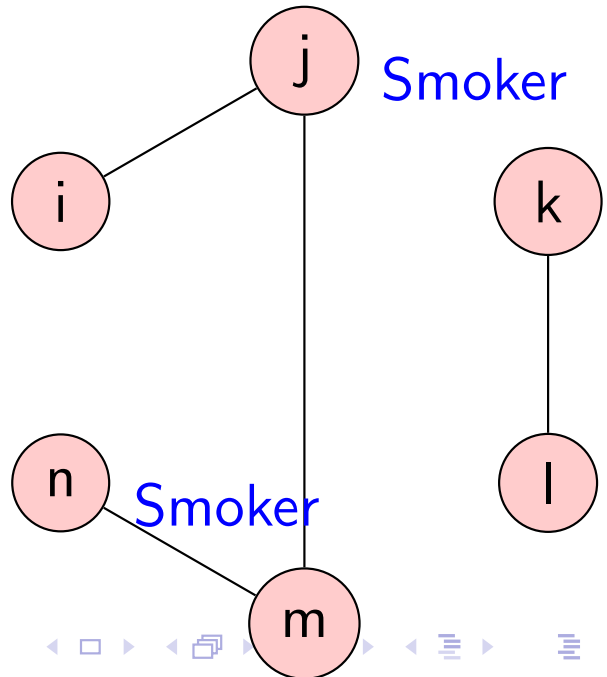

Statistical relational learning

Task 1

For each of the person of the second group we have to predict if he/she has a cancer or not

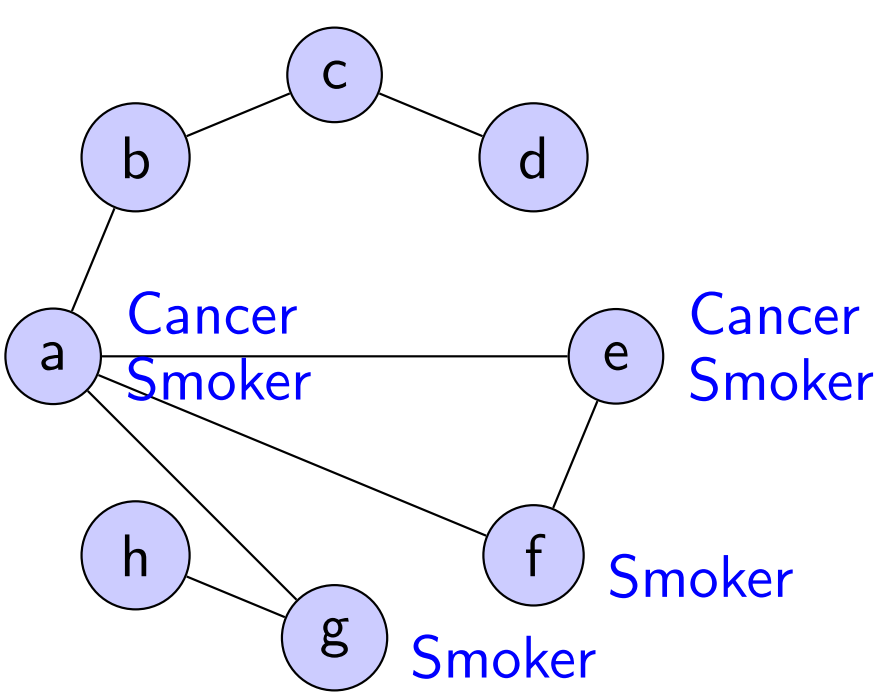

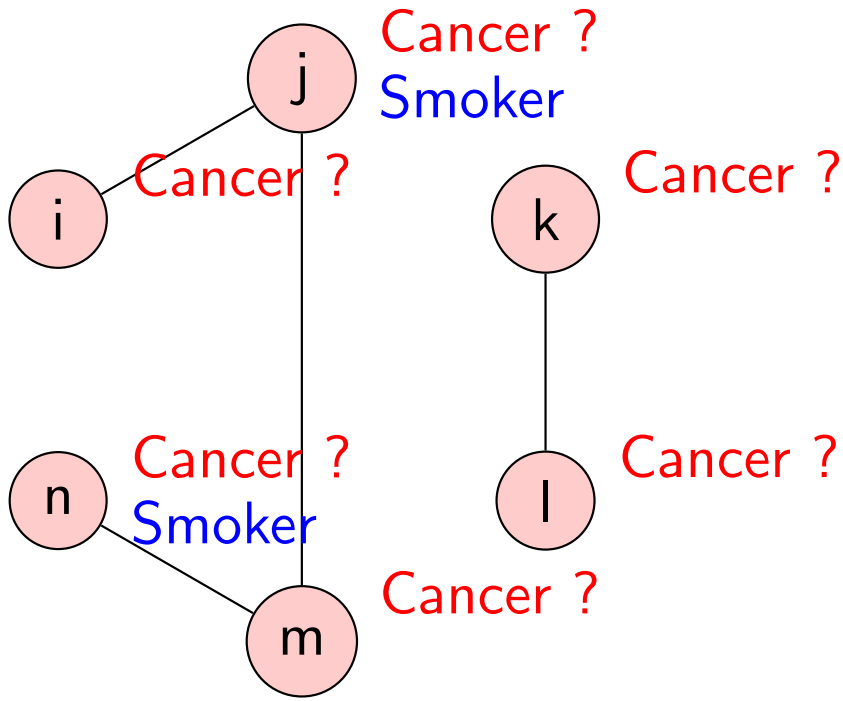

Statistical relational learning

Task 2

For each person we want to find a semantic embedding in $\mathbb{R}^k$ consistend with the semantics and the structure. For instance:

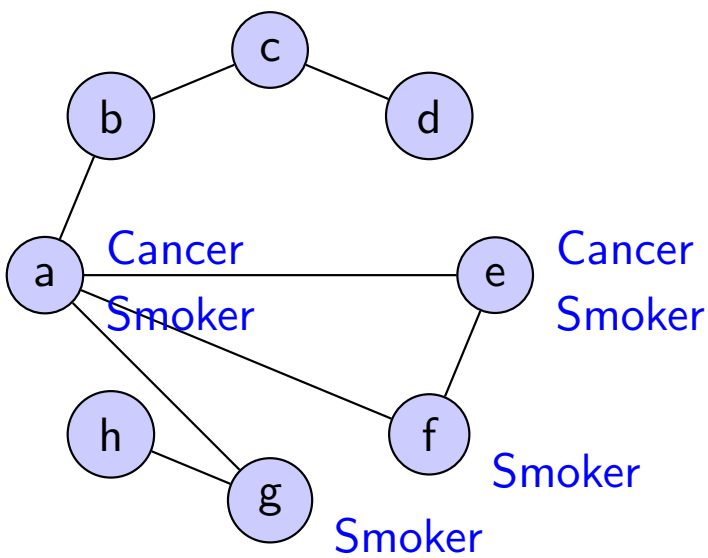

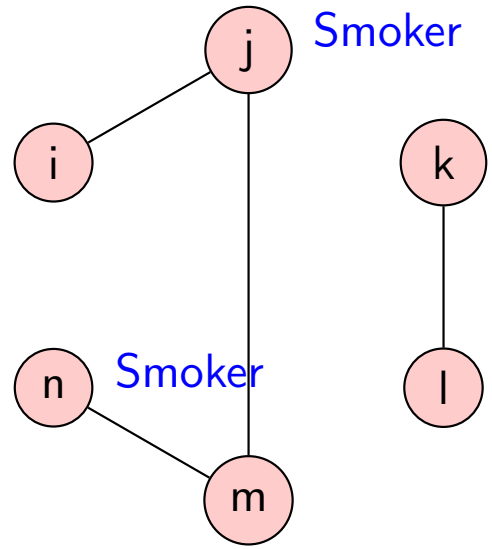

Statistical relational learning

Task 2

For each person we want to find a semantic embedding in $\mathbb{R}^k$ consistend with the semantics and the structure. For instance:

- $a^{\mathcal{G}} \approx e^{\mathcal{G}}$ since both a and e smoke and have cancer, and they have two friends that smoke, one of which has a cancer

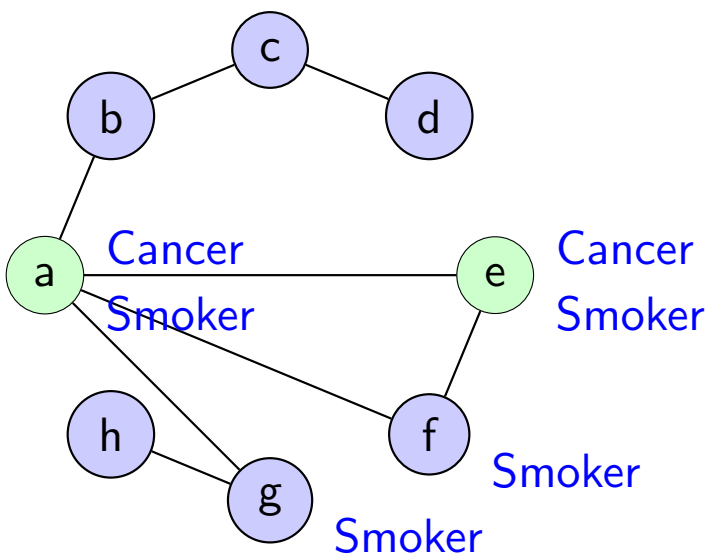

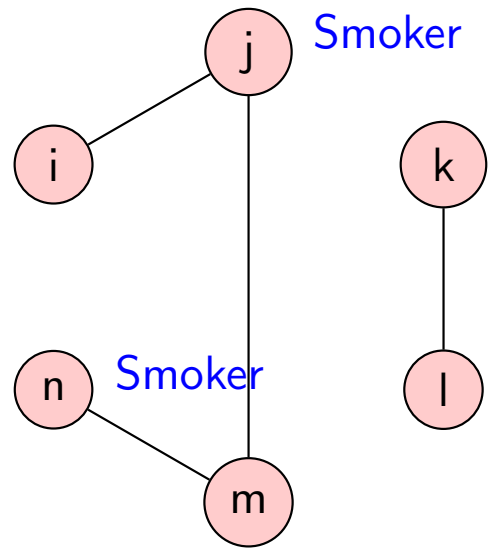

Statistical relational learning

Task 2

For each person we want to find a semantic embedding in $\mathbb{R}^k$ consistend with the semantics and the structure. For instance:

- $a^{\mathcal{G}} \approx e^{\mathcal{G}}$ since both a and e smoke and have cancer, and they have two friends that smoke, one of which has a cancer
- $d^{\mathcal{G}} \approx k^{\mathcal{G}} \approx i^{\mathcal{G}}$ because they don't smoke and don't have cancer, and they have only one friend, who does not smoke and does not have a cancer

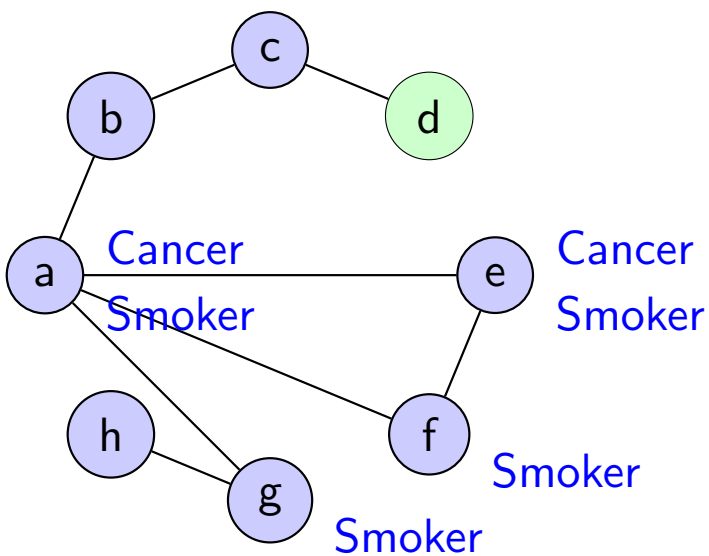

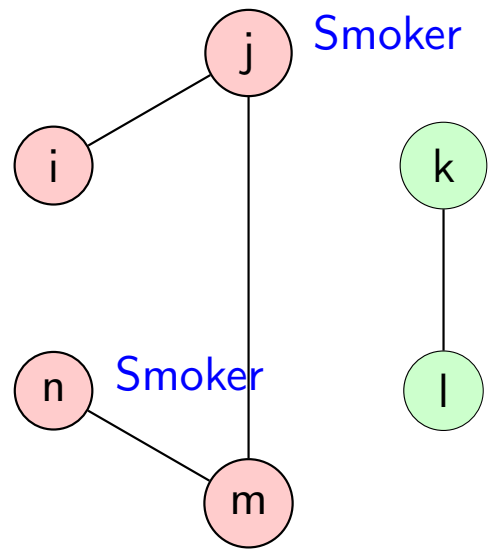

Statistical relational learning

Task 3

We want to know the truth value of certain formulas,

- e.g., the correlation between friendship and smoking habits;
- the correlation between smoking habits and cancer

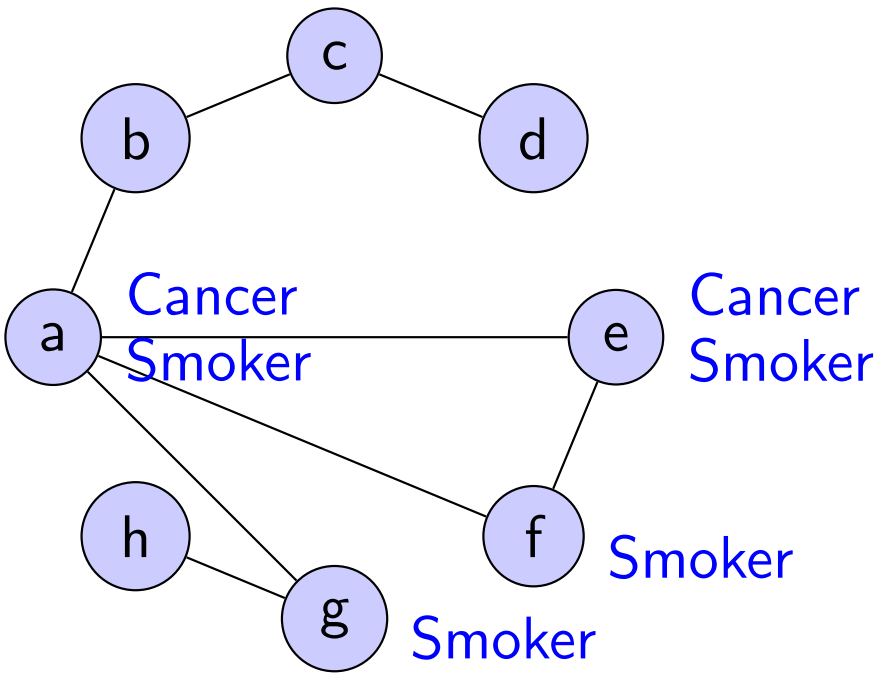

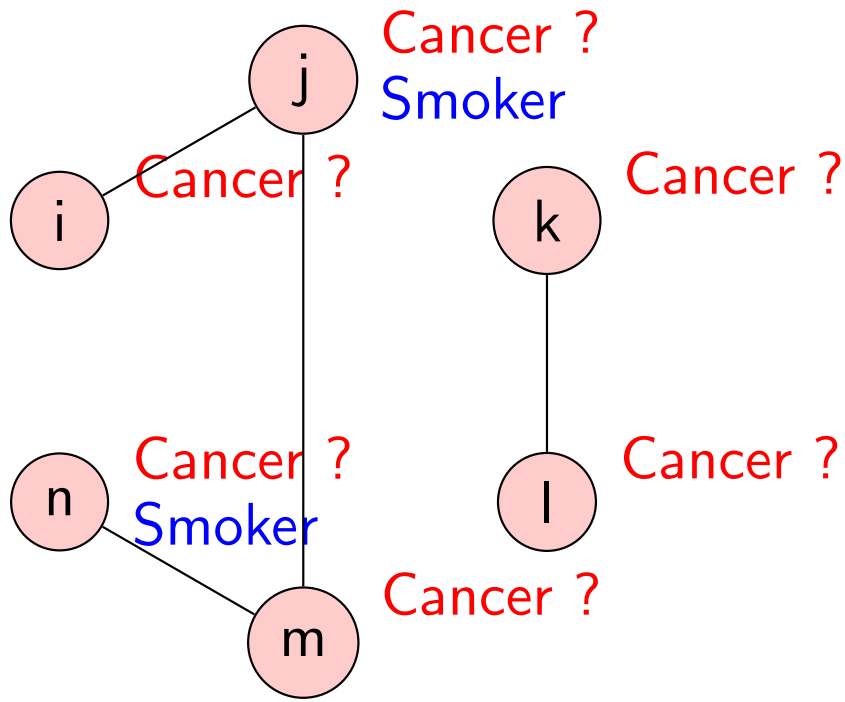

Statistical relational learning

Representation of the domain in $\mathbb{R}^k$

For each individual of the domain $\{a, \ldots, n\}$ we don't provide an explicit mapping to $\mathbb{R}^k$, which instead is generated, as the result of the constraint optimization. We only provide the dimension of the domain (i.e., k)

The language

- unary predicates $S(x)$ and $C(x)$ for "x smokes" and "x has a cancer" and a binary predicate $F(x, y)$ for "y is a fiend of x"

Constraints

- $S(a)$, $\neg S(b)$, $\neg S(c)$, $\neg S(d)$, $S(e)$, $\ldots \neg S(i)$, $\ldots$, $S(n)$;
- $C(a)$, $\neg C(b)$, $\neg C(c)$, $\ldots$, $\neg C(h)$;
- $\forall x : \neg F(x, x)$
- $\forall xy : F(x, y) \rightarrow F(y, x)$

Mnist with constraints

Dataset

Contains pictures resulting from overlaying two MNIST digit pictures, where the smaller digit is in black-on-white and the smallest in white-on-black. I.e., d_x and d_y are the pixel matrices of the digits x and y, then the pixel matrix $d_{xy} = d_x \oplus d_y$ is defined as

$$d_x \oplus d_y = \begin{cases} d_x - w \cdot d_y & \text{if } x \leq y \\ d_y - w \cdot d_x & \text{Otherwise} \end{cases}$$

where w is randomly generated number in $[0, 1]$

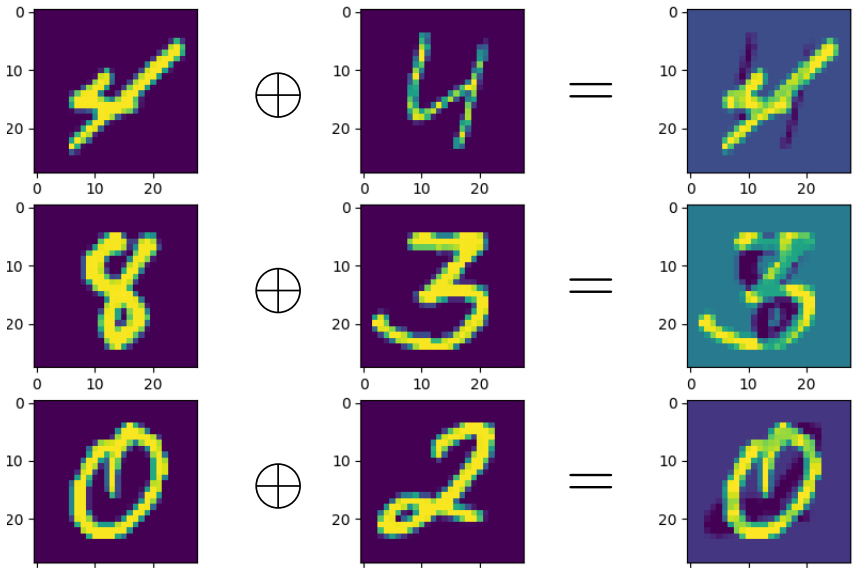

Mnist with constraints

The Task

Given an image $d_{xy} = d_x \oplus d_y$ we have to predict x and y.

Input

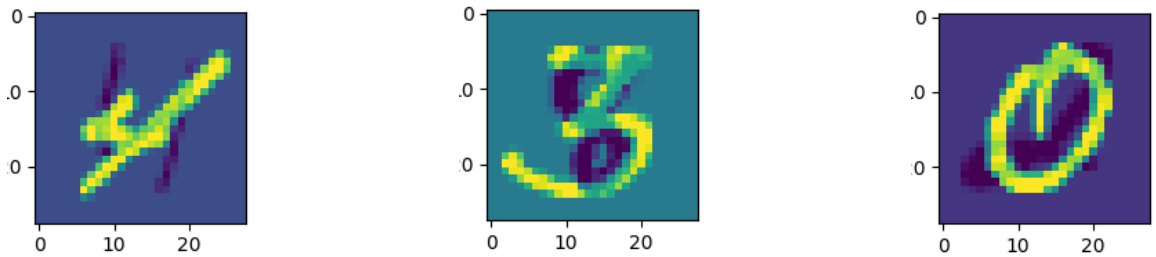

Mnist with constraints

The Task

Given an image $d_{xy} = d_x \oplus d_y$ we have to predict x and y.

Input

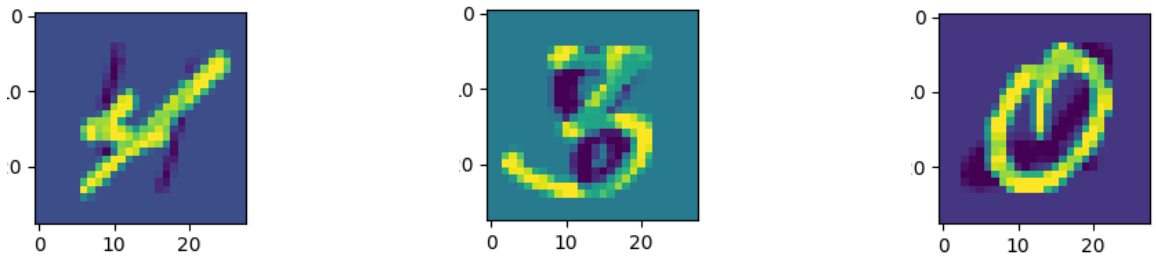

Output

4, 4 3, 8 0, 2

Mnist with constraints

The language

20 unary predicates, two for every digit;

$$\begin{array}{cccc} zero_1(x) & one_1(x) & \ldots & nine_1(x) \\ zero_2(x) & one_2(x) & \ldots & nine_2(x) \end{array}$$

$zero_1(x)$ (resp. $zero_2(x)$) means: "the smaller (resp the larger) digit of x is a 0

Constraints

- $four_1\left(\text{ 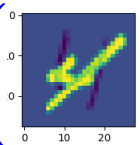 }\right)$, $three_1\left(\text{ 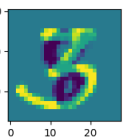 }\right)$, $zero_1\left(\text{ 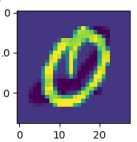 }\right)$,
  $four_2\left(\text{ 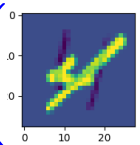 }\right)$, $eight_2\left(\text{ 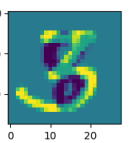 }\right)$, $two_2\left(\text{ 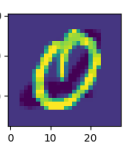 }\right)$,...
- $\forall x : zero_1(x) \rightarrow \neg one_1(x)$, ...
- $\forall x : \neg(one_1(x) \wedge zero_2(x))$, ...