Aristote Seminar - Hybrid AI

Combining explicit knowledge with data-driven AI : some examples on image interpretation

Céline Hudelot, Professeur
Laboratoire MICS - CentraleSupélec

18 janvier 2024

Plan

1. Introduction : An overview of Neuro-Symbolic AI
2. Some works on image interpretation
   - Classification with explanations and spatial relations
   - Knowledge integration into the learning algorithm
3. Conclusion

AI : two antagonistic approaches [a]

a. D. Cardon et al - La Revanche des neurones - https://hal.archives-ouvertes.fr/hal-02005537/document

Two different assumptions

- Human reasoning and knowledge are complex : knowledge implicitly in data.
  - Statistic or data-centric AI - Connectionist approaches - Learning from data.
  - Exploitation of the past experience represented by annotated data, building calibrated predictive models from it.
- Human reasoning can be captured, even if partially incomplete : explicit representation of knowledge (using symbols rather that statistics to represent the world).
  - Symbolic AI - Based on the modeling of logical reasoning, on formalisms for knowledge representation and reasoning.

AI : two antagonistic approaches [a]

a. D. Cardon et al - La Revanche des neurones - https://hal.archives-ouvertes.fr/hal-02005537/document

Two different assumptions

- Human reasoning and knowledge are complex : knowledge implicitly in data.
  - Statistic or data-centric AI - Connectionist approaches - Learning from data.
  - Exploitation of the past experience represented by annotated data, building calibrated predictive models from it.
- Human reasoning can be captured, even if partially incomplete : explicit representation of knowledge (using symbols rather that statistics to represent the world).
  - Symbolic AI - Based on the modeling of logical reasoning, on formalisms for knowledge representation and reasoning.

AI : two antagonistic approaches [a]

a. D. Cardon et al - La Revanche des neurones - https://hal.archives-ouvertes.fr/hal-02005537/document

Two different assumptions

- Human reasoning and knowledge are complex : knowledge implicitly in data.
  - Statistic or data-centric AI - Connectionist approaches - Learning from data.
  - Exploitation of the past experience represented by annotated data, building calibrated predictive models from it.
- Human reasoning can be captured, even if partially incomplete : explicit representation of knowledge (using symbols rather that statistics to represent the world).
  - Symbolic AI - Based on the modeling of logical reasoning, on formalisms for knowledge representation and reasoning.

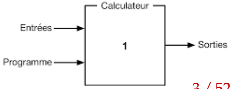

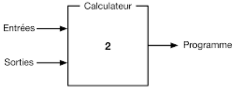

AI : two antagonistic approaches [a]

a. D. Cardon et al - La Revanche des neurones - https://hal.archives-ouvertes.fr/hal-02005537/document

Two different assumptions

- Human reasoning and knowledge are complex : knowledge implicitly in data.
  - Statistic or data-centric AI - Connectionist approaches - Learning from data.
  - Exploitation of the past experience represented by annotated data, building calibrated predictive models from it.
- Human reasoning can be captured, even if partially incomplete : explicit representation of knowledge (using symbols rather that statistics to represent the world).
  - Symbolic AI - Based on the modeling of logical reasoning, on formalisms for knowledge representation and reasoning.

AI : two antagonistic approaches [a]

a. D. Cardon et al - La Revanche des neurones - https://hal.archives-ouvertes.fr/hal-02005537/document

Two different assumptions

- Human reasoning and knowledge are complex : knowledge implicitly in data.
  - Statistic or data-centric AI - Connectionist approaches - Learning from data.
  - Exploitation of the past experience represented by annotated data, building calibrated predictive models from it.
- Human reasoning can be captured, even if partially incomplete : explicit representation of knowledge (using symbols rather that statistics to represent the world).
  - Symbolic AI - Based on the modeling of logical reasoning, on formalisms for knowledge representation and reasoning.

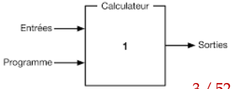

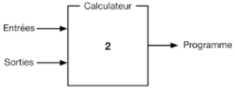

AI : two antagonistic approaches [a]

a. D. Cardon et al - La Revanche des neurones - https://hal.archives-ouvertes.fr/hal-02005537/document

Two different assumptions

- Human reasoning and knowledge are complex : knowledge implicitly in data.
  - Statistic or data-centric AI - Connectionist approaches - Learning from data.
  - Exploitation of the past experience represented by annotated data, building calibrated predictive models from it.
- Human reasoning can be captured, even if partially incomplete : explicit representation of knowledge (using symbols rather that statistics to represent the world).
  - Symbolic AI - Based on the modeling of logical reasoning, on formalisms for knowledge representation and reasoning.

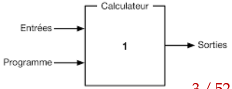

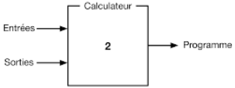

AI : two antagonistic approaches [a]

a. D. Cardon et al - La Revanche des neurones - https://hal.archives-ouvertes.fr/hal-02005537/document

Two different assumptions

- Human reasoning and knowledge are complex : knowledge implicitly in data.
  - Statistic or data-centric AI - Connectionist approaches - Learning from data.
  - Exploitation of the past experience represented by annotated data, building calibrated predictive models from it.
- Human reasoning can be captured, even if partially incomplete : explicit representation of knowledge (using symbols rather that statistics to represent the world).
  - Symbolic AI - Based on the modeling of logical reasoning, on formalisms for knowledge representation and reasoning.

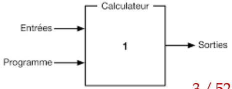

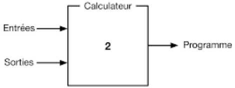

Hybrid AI

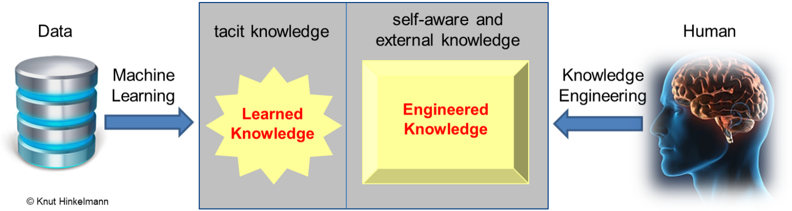

FIGURE – Source : `https://www.aaai-make.info/`

Bringing together, for added value, data-driven AI with symbolic and knowledge-oriented AI to answer their respective weaknesses.

Hybrid AI : why ?

Because :

- Lack of high level reasoning in deep-learning [Bottou,2011] [a].
- Deep neural models are *black box* models that can be easily fooled.
- More to predict that what is visible or readable (the knowledge is not totally inside the data).
- For some decision-based AI systems, the rules are to be told (ethics, policies, laws...) : ⇒ need of Knowledge Representation and Reasoning.
- Hybrid systems such as neuro-symbolic models are very promising (see `https://www.youtube.com/watch?v=4PuuziOgSU4`)

a. Bottou, Leon. (2011). From Machine Learning to Machine Reasoning. Computing Research Repository - CORR. 94. 10.1007/s10994-013-5335-x.

Hybrid AI

An important topic

With different names or sub-fields

- **Neuro-Symbolic Artificial Intelligence** : bringing together the neural and the symbolic traditions in AI (https://people.cs.ksu.edu/~hitzler/nesy/)
  - Neural : use of artificial neural networks, or connectionist systems.
  - Symbolic : AI approaches that are based on explicit symbol manipulation.
- **Informed Machine Learning** : integrating Prior Knowledge into Learning Systems.
  - (von Rueden et al, 21) [a] Learning from a hybrid information source that consists of data and prior knowledge. The prior knowledge comes from an independent source, is given by formal representations and is explicitly integrated into the ML pipeline
- **Knowledge Reasoning meets Machine Learning**
  - KR2ML workshops (https://kr2ml.github.io/)

Hybrid AI

Kautz's Neurosymbolic taxonomy

- **Symbolic Neuro Symbolic**
- **Symbolic[Neuro]**
- **Neuro ;Symbolic**
- **Neuro :Symbolic $\rightarrow$ Neuro**
- **Neuro$_{\text{Symbolic}}$**
- **Neuro[Symbolic]**

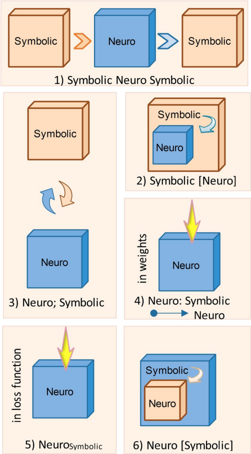

Source : Hassan et al, 22

Henri Kautz, The Third AI Summer : AAAI Robert S. Engelmore Memorial Lecture [1]

1. https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/19122

Hybrid AI : Kautz's Neurosymbolic taxonomy

Symbolic Neuro Symbolic

Input and output are presented in symbolic form, all the processing is neural

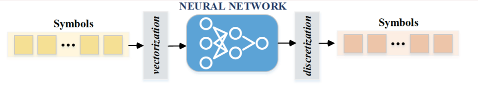

Source : Wang et al, PAMI 2023

Example : Deep learning procedure for NLP

Input symbols (words) are converted to vector embeddings (Glove ; Word2VEc), processed by the neural models whose output embeddings are converted to symbols.

Hybrid AI : Kautz's Neurosymbolic taxonomy

Symbolic[Neuro]

Symbolic systems, where neural modules are internally used as subroutine within a symbolic problem solver

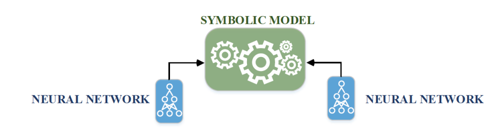

Example : Alpha Go

Monte Carlo Tree Search with a Neural network for the heuristic function evaluation

Hybrid AI : Kautz's Neurosymbolic taxonomy

Neuro ;Symbolic

Neural and symbolic parts focus on different but complementary tasks in a big pipeline., e.g a neural network focusing on one task (e.g. object detection) interacts via its input and output with a symbolic system specialising in a complementary task(e.g. question answering)

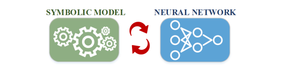

Example : Neuro-symbolic concept learner

A neural perception module learns visual concepts and a symbolic reasoning module executes symbolic programs on the concept representations for question answering (Mao et al, 2019) [a]

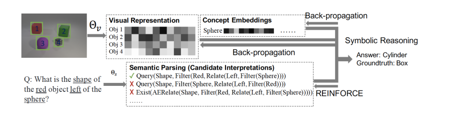

a. https://arxiv.org/abs/1904.12584

Hybrid AI : Kautz's Neurosymbolic taxonomy

Neuro :Symbolic → Neuro

Symbolic knowledge is compiled into the structure of Neural models

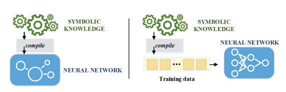

Example : GCN-based embedder

Vector based representations learning of symbolic knowledge to incorporate symbolic domain knowledge into connectionist architectures (Xie et al, 2019) [a]

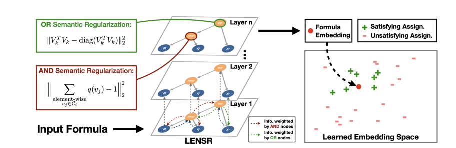

Figure 1: LENSR overview. Our GCN-based embedder projects logic graphs representing formulae or assignments onto a manifold where entailment is related to distance; satisfying assignments are closer to the associated formula. Such a space enables fast approximate entailment checks — we use this embedding space to form logic losses that regularize deep neural networks for a target task.

Hybrid AI : Kautz's Neurosymbolic taxonomy

Neuro$_{\text{Symbolic}}$

Symbolic knowledge is tensorized and neural methods are used to perform reasoning over this tensorized representation, e.g. additional soft-constraints in the loss function used to train DNNs.

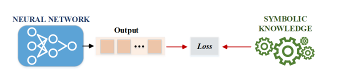

Example

Logic Tensor Networks (LTNs) ((Badreddine et al, 2022)[a]. First-order logic formulae are translated as fuzzy relations on real numbers for neural computing to allow gradient based sub-symbolic learning.

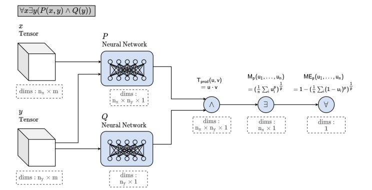

Hybrid AI : Kautz's Neurosymbolic taxonomy

Neuro[Symbolic]

Fully-integrated system, i.e. true symbolic reasoning inside a neural engine.

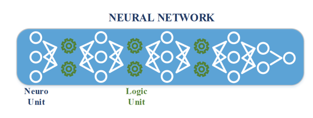

Example : Satnet

Imitating logical reasoning with tensor calculus to learn the execution of symbolic operations through neural networks. Satnet : A layer that enables end-to-end learning of both the constraints and so- lutions of logic problems and a smoothed differentiable (maximum) satisfiability solver that can be integrated into the loop of deep learning systems.(Wang et al, 2019) [a]

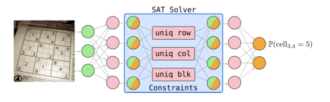

Hybrid AI (or Informed ML)

Principle (Rueden et al, 2021)

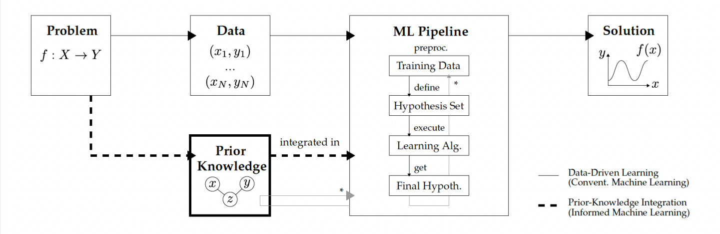

1. **Source** : Which source of knowledge is integrated ?
2. **Representation** : How is the knowledge represented ?
3. **Integration** : Where in the learning pipeline is it integrated ?
4. **Task** : What is the task ?
5. **Expected benefits** of the integration : explainability, performance, frugality ?

Hybrid AI (or Informed ML)

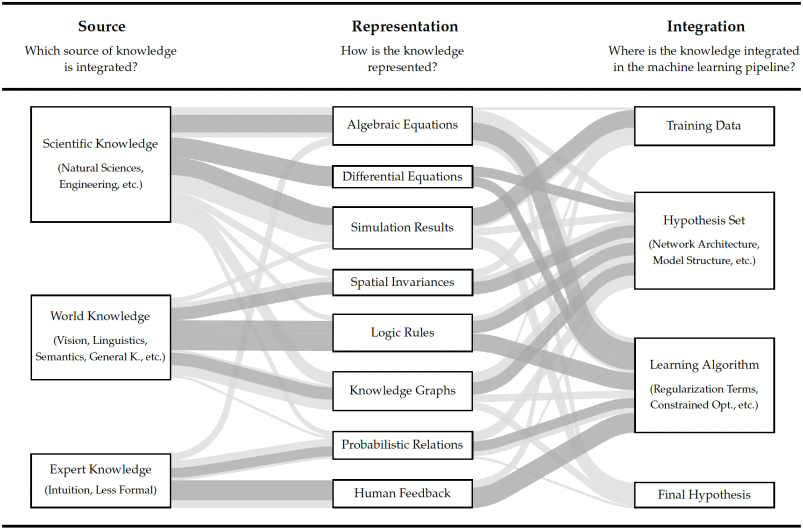

[Rueden et al, 2021]

Hybrid AI (or Informed ML)

Knowledge

Algebraic Equations	Differential Equations	Simulation Results	Spatial Invariances	Logic Rules	Knowledge Graphs	Probabilistic Relations	Human Feedback
$E = m \cdot c^2$ $v \leqslant c$	$\frac{\partial u}{\partial t} = \alpha \frac{\partial^2 u}{\partial x^2}$ $F(x) = m \frac{d^2 x}{dt^2}$	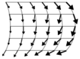	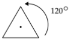	$A \wedge B \Rightarrow C$	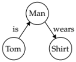	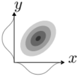	

[Rueden et al, 2021]

Plan

1. Introduction : An overview of Neuro-Symbolic AI

2. Some works on image interpretation
   - Classification with explanations and spatial relations
   - Knowledge integration into the learning algorithm

3. Conclusion

Classification or annotation with explanations

Coll J.P. Poli (CEA LIST), PhD of Regis Pierrard

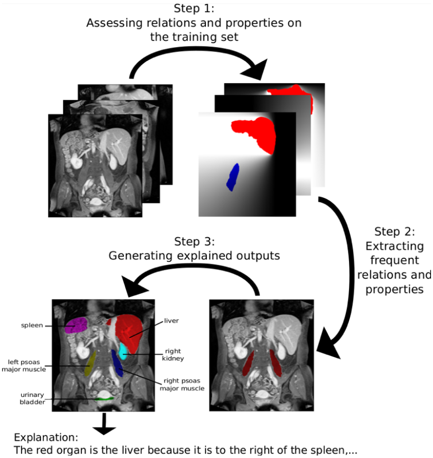

Classification or annotation with explanations

Main characteristics

- **Neuro;Symbolic** approach : Neural and symbolic parts focus on different but complementary tasks in a big pipeline.
- **Knowledge representation** : constraints or rules
- **Task** : classification or annotation
- **Expected benefit** : explainability

Step 1 : Evaluation of relations

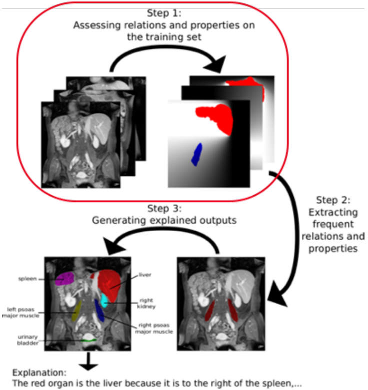

- Catalogue of relations : **Vocabulary**
- All n-ary relations are computed for each possible n-tuple of entities

Step 2 : learning relevant relations

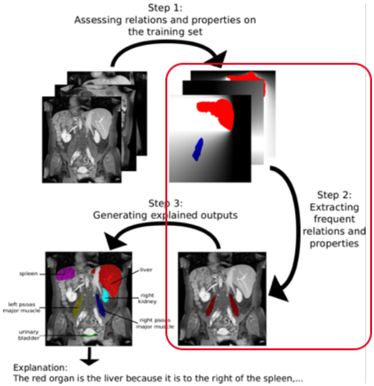

- Extracting relevant sets of relations : **descriptors**
- Postulate : relevant = frequent
- Class by class learning
- Frequent itemset mining with a fuzzification of the CLOSE Algorithm (Pierrard et al, 2028)

Step 3 : performing the task and generating explanations

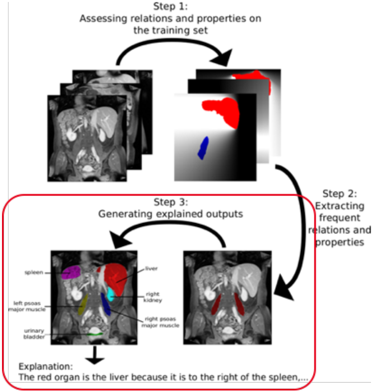

- Generating rules (classification) or constraints (annotation).
- Automatically generating explained outputs

Step 3 : Rule generation

- We build rules from the frequent descriptors.
- A frequent itemset I is maximal iff there exists no frequent itemset I' such as $I \subseteq I'$.
- Rules are built from descriptors that are both maximal and closed.
- Since the algorithm is performed class by class, rules are class-specific.
- Let $\mathcal{X}$ be the space on which instances are defined and let $\mathcal{Y}$ be a set of labels, $y \in \mathcal{Y}$ and MFC_y the set of descriptors corresponding to this class.
  - Remove from the descriptors all the relations that are common to several classes to get the set of discriminative descriptors MFC_y^*
  - For each $y \in \mathcal{Y}$, we can build a rule for each descriptor I in MFC_y^* such as $IF \bigvee_{R \in I} R$ then class $=y$
  - There may be several descriptors for a given class, rules are aggregated

set of maximal closed itemsets

$$MFC_y^* = \left\{ I^* \mid I^* = I \setminus \overbrace{\bigcup_{y' \in \mathcal{Y} \setminus \{y\}} \left(\bigcup_{J \in MFC_{y'}} J \right)}^{\text{sets of items describing other classes}}, I \in MFC_y \right\}$$

$$\forall y \in \mathcal{Y}, \mathbf{x} \in \mathcal{X}, \mu_y(\mathbf{x}) = \bigvee_{I^* \in MCF_y^*} \left(\text{support}(I^*) \times \bigwedge_{\mathcal{R} \in I^*} R\right)$$

Step 3 : Constraint generation

- Constraints are generated from the maximal closed descriptor of largest cardinality

$$I_y^M = \underset{I \in \mathcal{J}}{\operatorname{argmax}} \left[\text{support}(I)\right] \text{ such as } \overbrace{\mathcal{J} = \{J \in MFC_y \mid |J| = \max_{P \in MFC_y} |P|\}}^{\text{maximal closed descriptor of largest cardinality}}$$

- All frequent relations in I_y^M, are translated into constraints, such as

$$C_i = ((\nu_1, \nu_2), R_i)$$

- Fuzzy constraint satisfaction problem (FCSP) [Dubois, 1996] :

$$X = \{\nu_1, \nu_2, \ldots \nu_{n_X}\}$$

$$D = \{d_1, d_2, \ldots d_{n_X}\}$$

$$C = \{c_1, c_2, \ldots c_{n_C}\}$$

- Solving a FCSP consists in two steps :
  - Backtracking
  - FAC-3 [Dubois, 1996 ; Vanegas, 2016]
- Generation of the explanation from the constraints

Explainable automatic image annotation

The dataset (images taken from the VISCERAL project)

- CT and MR images
- Whole body and abdomen
- 35 images.
- 9 organs to label on each image

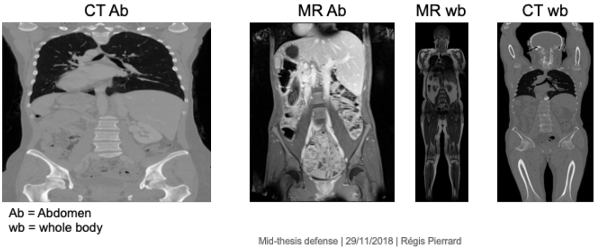

Explainable automatic image annotation

CSP

- Variables X : the organs to label
- Domains D : regions of the image obtained by segmentation
- Constraints C : fuzzy spatial relations between organs

Fuzzy spatial relations

- (completely) Above
- (completely) Below
- (completely) To the left of
- (completely) To the right of
- Line symmetry [Colliot, 2004]

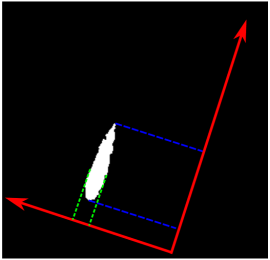

Explainable automatic image annotation

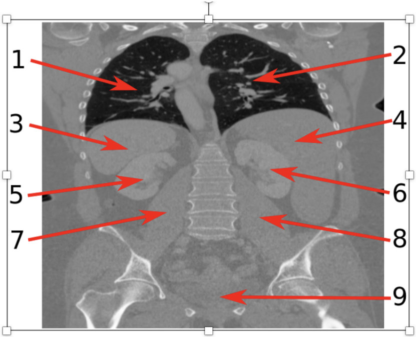

- **Entity 4** is annotated as the *liver* with a very high confidence **because**
  - it is below entity 2 (right lung),
  - it is completely to the right of entity 3 (spleen),
  - entity 3 (spleen) is completely to its left,
  - entity 8 (right psoas) is below it,
  - entity 6 (right lung) is completely below it.
- **Entity**…

Explainable automatic image classification : other ongoing work

Ph.D thesis of Dao Thauvin with Stephane Herbin (ONERA) and Wassila Ouerdane (MICS) on the use of dialogue and argumentation for explainable image classification.

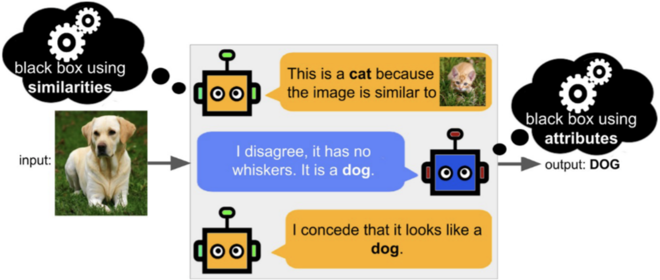

Figure 1: Exemple de Dialogue pour réaliser une classification d'images avec 2 agents, l'un raisonne par similarité (en jaune), l'autre par attributs (en bleu).

Plan

1. Introduction : An overview of Neuro-Symbolic AI
2. Some works on image interpretation
   - Classification with explanations and spatial relations
   - Knowledge integration into the learning algorithm
3. Conclusion

Semantic regularization, projection and conditioning at inference

- Integrating logical constraints into multi-label classification models
- Work in collaboration with IRT System X, SMD project, PhD of Arthur Ledaguenel with M. KHOUADJIA
- **Neuro :Symbolic** or **Neuro$_{\textbf{Symbolic}}$** approach : symbolic knowledge is compiled into the structure of Neural models
- **Knowledge representation** : propositional logics
- **Task** : multi-label classification
- **Expected benefit** : frugality

Binary multilabel classification with semantic constraints

Thèse de A. Ledaguenel, IRT System X, SMD

- **Model** $\mathcal{M}_\theta(X)$
- **Dataset** $(X^i, y^i)_{1 \leq i \leq n}$ with $X^i \in \mathbb{R}^d$ and $y^i \in \{0,1\}^k$
- **Categories** $\{Y_j\}_{1 \leq j \leq k}$
- **Logical sentence** (or propositional formula) α over the categories

Example : $\alpha = (Y_1 \vee \neg Y_2)$ implies that any instantiation $y = (0, 1, ...)$ is **invalid**.

Task

Predict an instantiation over a set of categories

Binary multilabel classification with semantic constraints

Thèse de A. Ledaguenel, IRT System X, SMD

- **Model** $\mathcal{M}_\theta(X)$
- **Dataset** $(X^i, y^i)_{1 \leq i \leq n}$ with $X^i \in \mathbb{R}^d$ and $y^i \in \{0,1\}^k$
- **Categories** $\{Y_j\}_{1 \leq j \leq k}$
- **Logical sentence** (or propositional formula) α over the categories

Example : $\alpha = (Y_1 \vee \neg Y_2)$ implies that any instantiation $y = (0, 1, ...)$ is **invalid**.

Task

Predict an instantiation over a set of categories

Binary multilabel classification with semantic constraints

Thèse de A. Ledaguenel, IRT System X, SMD

- **Model** $\mathcal{M}_\theta(X)$
- **Dataset** $(X^i, y^i)_{1 \leq i \leq n}$ with $X^i \in \mathbb{R}^d$ and $y^i \in \{0,1\}^k$
- **Categories** $\{Y_j\}_{1 \leq j \leq k}$
- **Logical sentence** (or propositional formula) α over the categories

Example : $\alpha = (Y_1 \vee \neg Y_2)$ implies that any instantiation $y = (0, 1, ...)$ is **invalid**.

Task

Predict an instantiation over a set of categories

Binary multilabel classification with semantic constraints

Thèse de A. Ledaguenel, IRT System X, SMD

- **Model** $\mathcal{M}_\theta(X)$
- **Dataset** $(X^i, y^i)_{1 \leq i \leq n}$ with $X^i \in \mathbb{R}^d$ and $y^i \in \{0,1\}^k$
- **Categories** $\{Y_j\}_{1 \leq j \leq k}$
- **Logical sentence** (or propositional formula) α over the categories

Example : $\alpha = (Y_1 \vee \neg Y_2)$ implies that any instantiation $y = (0, 1, ...)$ is **invalid**.

Task

Predict an instantiation over a set of categories

Binary multilabel classification with semantic constraints

Thèse de A. Ledaguenel, IRT System X, SMD

- **Model** $\mathcal{M}_\theta(X)$
- **Dataset** $(X^i, y^i)_{1\leq i\leq n}$ with $X^i \in \mathbb{R}^d$ and $y^i \in \{0,1\}^k$
- **Categories** $\{Y_j\}_{1\leq j\leq k}$
- **Logical sentence** (or propositional formula) α over the categories

Example : $\alpha = (Y_1 \vee \neg Y_2)$ implies that any instantiation $y = (0, 1, ...)$ is **invalid**.

Task

Predict an instantiation over a set of categories

Example : Mutual Exclusion

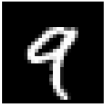

$$\mathbf{y} := (0, 0, 0, 0, 0, 0, 0, 0, 0, 1)$$

Example : Mutual Exclusion

$$\mathbf{y} := (0, 0, 0, 1, 0, 0, 0, 0, 0, 0)$$

Example : Mutual Exclusion

$$\alpha$$

$$X_0 \vee X_1 \vee X_2 \vee X_3 \vee X_4 \vee X_5 \vee X_6 \vee X_7 \vee X_8 \vee X_9$$

$$\bigwedge_{0 \leq i < j \leq 9} (\neg X_i \vee \neg X_j)$$

Example : Hierarchical

Bulbul

$\mathbf{y} \in \{0,1\}^{1000}$

Example : Hierarchical

Bulbul

$$\mathbf{y} \in \{0,1\}^{1000}$$

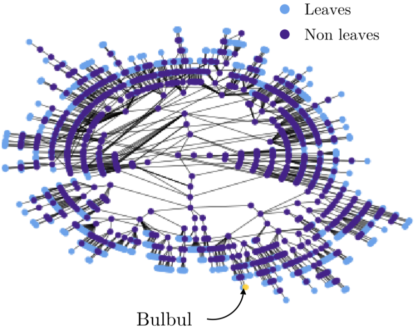

Example : Hierarchical

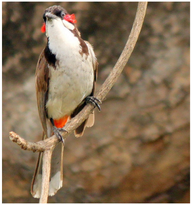

Bulbul

$\mathbf{y} \in \{0,1\}^{1860}$

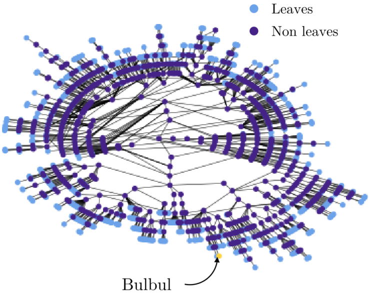

$$\kappa_H := \left(\bigwedge_{(i,j) \in E_h} Y_i \vee \neg Y_j \right) \wedge \left(\bigwedge_{(i,j) \in E_e} (\neg Y_i \vee \neg Y_j) \right) \tag{1}$$

the first part ensures that a son node cannot be *true* if its father node is not and the second part prevents two mutually exclusive nodes to be *true* simultaneously.

Explosion of dimensionality

Dimensionality of label space

Given categories $\{Y_j\}_{1 \leq j \leq k}$, the number of instantiations is $\mathbf{2^k}$

This implies that the label space becomes **quickly intractable** for **standard mutually exclusive classification**

We have *a priori* **no control** over the **growth** of the space of **valid instantiations**.

Explosion of dimensionality

Dimensionality of label space

Given categories $\{Y_j\}_{1\leq j\leq k}$, the number of instantiations is $\mathbf{2^k}$

This implies that the label space becomes **quickly intractable** for **standard mutually exclusive classification**

We have *a priori* **no control** over the **growth** of the space of **valid instantiations**.

Explosion of dimensionality

Dimensionality of label space

Given categories $\{Y_j\}_{1 \leq j \leq k}$, the number of instantiations is $\mathbf{2^k}$

This implies that the label space becomes **quickly intractable** for **standard mutually exclusive classification**

We have *a priori* **no control** over the **growth** of the space of **valid instantiations**.

Task

First let's analyse how we deal with standard **binary multilabel classification** :

- **Model** $\mathcal{M}_\theta(X)$
- **Dataset** $(X^i, y^i)_{1 \le i \le n}$ with $X^i \in \mathbb{R}^d$ and $y^i \in \{0, 1\}^k$
- **Categories** $\{Y_j\}_{1 \le j \le k}$
- ~~**Logical sentence**~~

Task

First let's analyse how we deal with standard **binary multilabel classification** :

- **Model** $\mathcal{M}_\theta(X)$
- **Dataset** $(X^i, y^i)_{1\leq i\leq n}$ with $X^i \in \mathbb{R}^d$ and $y^i \in \{0,1\}^k$
- **Categories** $\{Y_j\}_{1\leq j\leq k}$
- ~~**Logical sentence**~~

Binary Multilabel Classification

Objective

Add structural prior to enforce the instantiation problem

Classical approaches

- **Strictly mutually exclusive categories** : by adding a softmax layer on the last activation scores produced by the model
- **Independent multilabel regression** : all combinations of variables are possible, the probability of a category being activated depends only on the activation score of that category : implementation by a sigmoid layer on the last activation scores

Our objective : find a formalism that enables to enforce a more complex range of structural prior.

Standard Deep Learning Model

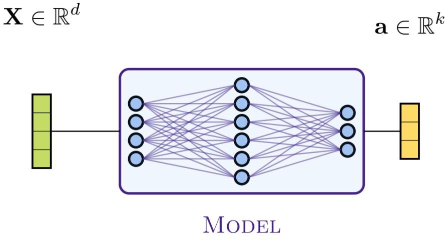

Neural Classification System

Pseudo-functional framework

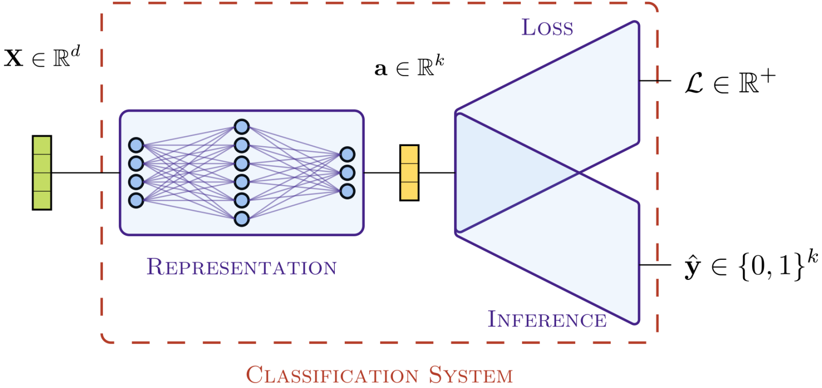

Standard multi-label modules

TRAINING	INFERENCE
cross-entropy	*mode*
$\mathcal{L}_\perp = -\log(\mathbf{P}_\perp(\mathbf{y}^t \mid \mathbf{a}))$	$\hat{\mathbf{y}} = \arg\max_{\mathbf{y}\in\{0,1\}^k} \mathbf{P}_\perp(\mathbf{y} \mid \mathbf{a})$

TABLE – Inference and Training

with $\mathbf{y}^t$ the ground truth
Probabilistic interpretation : The model produces a conditional distribution of the output given the input, the loss module computes the cross entropy of that distribution with a given label and the inference module computes the most probable output given the learned distribution.

Standard multi-label modules

TRAINING	INFERENCE
loss	*mode*
$\mathcal{L} = -\sum_{1\leq i\leq k} \mathbf{y}_i^t . \log(p_i) + (1-\mathbf{y}_i^t) . \log(1-p_i)$	$\hat{\mathbf{y}} = 1[\mathbf{p} \geq 0.5]$

TABLE – Inference and Training

with $\mathbf{y}^t$ the ground truth, $\mathbf{p} \in [0,1]^k$ and $\forall i, p_i = s(a_i) = \frac{e^{a_i}}{1+e^{a_i}}$

Objective :find a **unified** formalism that enables to enforce a more complex range of structural prior.

- **Inference** : enforce a certain probability distribution over the instanciation space $\{0, 1\}^k$ from which the most probable instanciation will be predicted
- **Training** : compute a loss functional to be optimized through backpropagation given label samples.

Derive appropriate loss and inference modules from the propositional background knowledge of the task.

More complex range of structured prior

Main Neurosymbolic approaches

To take into account **semantic constraints** over the output variables **Y** to enforce a specific probability distribution and-or a loss functional

- **Semantic regularization** : class of techniques that turn logical formulas into regularization terms which are added to the standard cross-entropy loss to steer neural models towards outputs that satisfy the background knowledge, without affecting inference
- **Semantic conditioning** : use probabilistic logics to condition the probability distribution on the output space produced by the neural model on the background knowledge. Both the loss and the inference modules are impacted
- **Semantic conditioning at inference (our)** : inspired from semantic conditioning, but only applies conditioning to the inference module (i.e., infers the most probable output that entails K)

Semantic Conditioning

TRAINING	INFERENCE
cross-entropy	*mode*
$\mathcal{L}_{\vert\alpha} = -\log(\mathbf{P}_\perp(\mathbf{y}^t \vert \mathbf{a}, \alpha))$	$\hat{\mathbf{y}}_{\vert\alpha} = \underset{\mathbf{y}\in\{0,1\}^k}{\arg\max} \mathbf{P}_\perp(\mathbf{y} \vert \mathbf{a}, \alpha)$

TABLE – Inference and Training

Implemented for hierarchical classification by Jia Deng in **Deng2014**, then generalized to arbitrary semantics by Kareem Ahmed in **Ahmed2022spl**

Semantic Regularization

TRAINING	INFERENCE
cross-entropy	*mode*
$\mathcal{L}_r = \mathcal{L}_\perp + \lambda.(-\log(\mathbf{P}_\perp(\alpha\|\mathbf{a})))$	$\hat{\mathbf{y}}_\perp = \arg\max_{\mathbf{y}\in\{0,1\}^k} \mathbf{P}_\perp(\mathbf{y}\|\mathbf{a})$

TABLE – Inference and Training

with $\lambda > 0$, typically $0 < \lambda < 1$

Equivalent to the semantic loss introduced by Xu in **Xu2018**

Semantic Conditioning at Inference

TRAINING	INFERENCE
cross-entropy	*mode*
$\mathcal{L}_\perp = -\log(\mathbf{P}_\perp(\mathbf{y}^t \vert \mathbf{a}))$	$\hat{\mathbf{y}}_{\vert\alpha} = \underset{\mathbf{y}\in\{0,1\}^k}{\arg\max} \mathbf{P}_\perp(\mathbf{y} \vert \mathbf{a}, \alpha)$

TABLE – Inference and Training

Motivations

- A common framework and an evaluation methodology to compare the different approaches including the study of the evolution of the benefits when resources given to the system (e.g. network scale, dataset size, training length, etc.) increase.
- First experimental results on categorical and hierarchical classification lead to several nice observations (outperforming of conditioned inference module (e.g. sc and sci) on those based on independent multilabel classification (e.g. imc and sr), importance of the scale,...)
- sci has interesting computational tractability and efficiency properties.
- sci principle consists in training with the standard imc loss that leads to independent repesentations of the different variables.
  - Enable to change background knowledge after training, for instance if the background knowledge is unavailable at training time or susceptible to evolve, which is not possible when learned representations are tightly entangled through the knowledge.
  - useful property in the era of foundation models

Conclusion

- Obvious motivations.
- A growing literature.
- Many initiatives.
- Many to do.

Conclusion

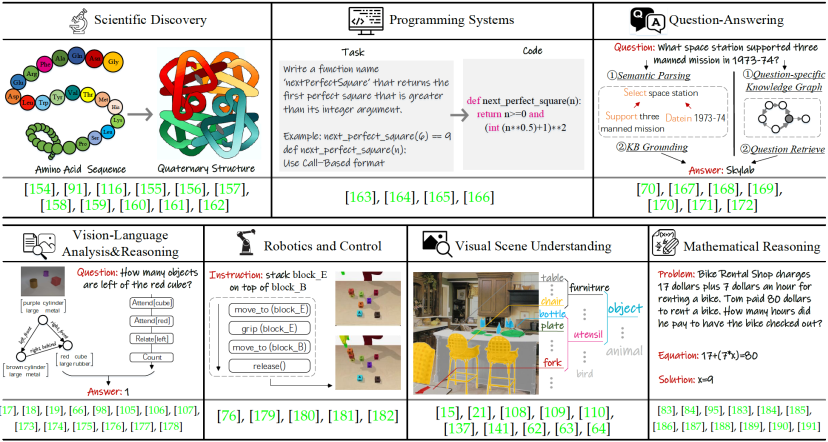

Fig. 10. Summary of typical applications that have been benefited from NeSy (see Sec. 5).

Source : Wang et al, PAMI 2023