

Hybrid AI Systems Grounded on Just-in-time Qualitative Spatio-Temporal Reasoning

Michael Sioutis

Faculty of Information Systems & Applied Computer Sciences
University of Bamberg
Germany

25 October, 2022



Outline

- 1 Bio
- 2 Hybrid AI: Overview
- 3 Hybrid AI with a Spatio-Temporal Flavor
- 4 Fundamentals of Spatio-Temporal Formalisms
- 5 Research Topics on Hybrid Spatio-Temporal AI
- 6 Discussion

Personal Details

- Greek, born in Wickede, Germany
- Languages: Greek \succ English \succ German \succ Spanish \succ French
- Research fellow (Wissenschaftlicher Mitarbeiter) with 5h/week teaching load at Otto-Friedrich-University Bamberg (Germany); Due to start as a Professeur Junior at Université Montpellier (France)




International Outlook

- Postdocs in France, Sweden, Finland, Germany
- Seeking collaboration with Mexico



Research Interests

- Artificial Intelligence
- Data Mining
- Linux System Engineering
- Logic Programming
- Semantic Web



- Ph.D. in Computer Science in 2017 from Université d'Artois (France); brief stay in University of Technology Sydney (Australia)



- M.Sc. in Advanced Information Systems and B.Sc. in Computer Science in 2012 and 2009 respectively from the National and Kapodistrian University of Athens (Greece)



HELLENIC REPUBLIC
National and Kapodistrian
University of Athens

Hybrid AI: An Overview

Part I

Neural-symbolic AI seeks to bring together robust learning in neural networks with reasoning and explainability via symbolic representations for network models

A. d'Avila Garcez and L. C. Lamb¹

¹A. d'Avila Garcez and L. C. Lamb: *Neurosymbolic AI: The 3rd Wave*. (2020)
<https://arxiv.org/abs/2012.05876>

Neuro-Symbolic AI Architectures

- Neural architectures for logical reasoning²
- Logical specification of neural networks³
- *Architectures integrating learning and reasoning*⁴
- ***Our main focus here!***

²E.g., W. W. Cohen et al.: *TensorLog: A Probabilistic Database Implemented Using Deep-Learning Infrastructure*. In: J. Artif. Intell. Res. 67 (2020)

³E.g., G. Soarek et al.: *Lifted Relational Neural Networks: Efficient Learning of Latent Relational Structures*. In: J. Artif. Intell. Res. 62 (2018)

⁴E.g., R. Manhaeve et al.: *Neural probabilistic logic programming in DeepProbLog*. In: Artif. Intell. 298 (2021)

Neuro-Symbolic AI Overview




Figure: Cyclical interaction in Neuro-Symbolic AI; a symbolic system feeds symbolic (partial) knowledge to a neural network system, which can be trained on raw data, and knowledge acquired through machine learning can then be extracted back to the symbolic system, and made available for further processing in symbolic form⁵

⁵S. Bader and P. Hitzler: *Dimensions of Neural-symbolic Integration - A Structured Survey*. In: We Will Show Them! Essays in Honour of Dov Gabbay (1). 2005


Neuro-Symbolic AI Components

Symbolic AI module

- A knowledge base (KB), which stores facts and rules
- A symbolic model, viz., the KB's logic (e.g., modal logic, first-order logic, deductive reasoning, and so on)

Machine learning module

- Any kind of neural network architecture or other connectionist model⁶



- Raw data

⁶M. Minsky: *Logical Versus Analogical or Symbolic Versus Connectionist or Neat Versus Scruffy*. In: AI Mag. 12 (1991)

Alternative Terms for Neuro-Symbolic AI

- Data-driven vs knowledge-driven
- Subsymbolic vs symbolic
- Learners vs solvers
- Thinking fast vs thinking slow (Daniel Kahneman)
- Statistical learning vs symbolic reasoning
- Perception vs logical reasoning
- Connectionism vs computationalism
- Scruffy vs neat

Why Neuro-Symbolic AI?

- High explainability and high accuracy performance
 - Explainability stems from Symbolic AI
 - Accuracy⁷ stems from ML model
- Mimics human problem-solving
 - Perception, mostly accurate and fast, can be seen as a data-driven process
 - Reasoning, mostly complex and slow, can be realized by logical reasoning
- Hot research topic and major AI challenge⁸

⁷ML models can be designed and trained with relatively much less effort compared to their accuracy performance

⁸<https://ai100.stanford.edu/2021-report/gathering-strength-gathering-storms-one-hundred-year-study-artificial-intelligence>

Ok, but Why Not End-to-end Learning?

Representation learning only includes statistical information—it does not capture perturbations (interventions), reasoning, planning

*ML models should be tied to assumptions that align with **physics** and **human cognition** to allow for generalization*

Modular approach: decompose knowledge into independent and recomposable pieces

B. Schölkopf et al.⁹

⁹B. Schölkopf et al.: *Towards Causal Representation Learning*. (2021)
<https://arxiv.org/abs/2102.11107>

ML Caveats: Statistics and Causality (1/4)




Figure: Nicolas Cage causing deaths?¹⁰

¹⁰<http://www.tylervigen.com/spurious-correlations>

ML Caveats: Perturbations/Interventions (2/4)




Figure: An adversarial example¹¹

- Relation to causality: conditional probabilities cannot reliably predict the outcome of an intervention
 - “Seeing people with open umbrellas suggests that it is raining” → “Closing umbrellas does not stop the rain”

¹¹https://gradientscience.org/intro_adversarial/

ML Caveats: Counterfactuals (3/4)

- Counterfactual problems involve:
 - reasoning about why things happened;
 - imagining the consequences of different actions in hindsight;
 - and determining which actions would have achieved a desired outcome
- “Would a given patient have suffered heart failure if they had started exercising a year earlier?”

ML Caveats: Data Diversity (4/4)

- Data diversity is an untestable condition a priori



Figure: Thanks to machine learning models the robot apocalypse was short lived¹²

¹²<https://www.smbc-comics.com/comic/rise-of-the-machines>

Ok, but Why Not Just Symbolic AI Then?

- Knowledge acquisition bottleneck
 - Never-ending need for carefully crafted rules
 - Too much data vs too few domain experts to classify and structure it
 - Semantic alignment issues
- Symbols can be largely used when the input is definite and falls under certainty

Neuro-Symbolic AI Sum-up So Far

- Better than standalone Machine Learning or Symbolic AI
- Statistics (correlations) + logic \approx causality
- Raises many open questions / opportunities for research
- Resembles human reasoning, with respect to some rationale
- Still far away from dealing with “insight problems”¹³

¹³<https://thereader.mitpress.mit.edu/ai-insight-problems-quirks-human-intelligence/>

Useful Links

- Browse the following Neuro-Symbolic AI website
<http://www.neural-symbolic.org/>
- Watch the latest Neural-Symbolic AI tutorials
<https://dtai.cs.kuleuven.be/tutorials/nesytutorial/>
- Watch the lecture by Prof. Bernhard Schölkopf on Symbolic, Statistical, and Causal Representations
<https://youtu.be/gSfntg5vpUUx>

Hybrid AI with a Spatio-Temporal Flavor

Part II

We recall the previous quote

*ML models should be tied to assumptions that align with **physics** and **human cognition** to allow for generalization*

B. Schölkopf et al.¹⁴

¹⁴B. Schölkopf et al.: *Towards Causal Representation Learning*. (2021)
<https://arxiv.org/abs/2102.11107>

Two Logical Rules: Which is Better? Why?

- “if I accelerate faster than the vehicle directly in front of me, then I will overtake it”
- “if I accelerate faster than the vehicle directly in front of me, then I will bump into it”

Injecting Causality via Spatio-Temporal Calculi

- Qualitative Spatial & Temporal Reasoning (QSTR) is a major field of study in KR, and Symbolic AI in general¹⁵
- QSTR abstracts from numerical quantities of space and time by using natural descriptions instead (e.g., *precedes*, *contains*, *is left of*), grounded on **physics** and **human cognition**




Figure: Left: Qualitative abstraction of a spatial configuration, Right: Temporal constraint network of 3 variables (tasks); ? denotes complete uncertainty

¹⁵G. Ligozat.: *Qualitative Spatial and Temporal Reasoning*. ISTE Series. Wiley, 2011

Example Calculus: RCC8




Figure: The base relations of RCC8; $\cdot i$ denotes the inverse of \cdot .

An abundance of calculi exist for many types of spatial and temporal information (e.g., trajectories, occlusion, intervals)¹⁶

¹⁶F. Dylla et al.: *A Survey of Qualitative Spatial and Temporal Calculi: Algebraic and Computational Properties*. In: ACM Comput. Surv. 50 (2017)

Example Calculus: Allen's Interval Algebra




Figure: The base relations of Interval Algebra; $\cdot i$ denotes the inverse of \cdot .

Applications: Geospatial Semantic Segmentation




Figure: A semantic referee reasons about the mistakes made by the classifier based on ontological concepts and provides additional information back to the classifier that prevents the classifier from making the same misclassifications¹⁷

¹⁷M. Alirezaie et al.: *Semantic referee: A neural-symbolic framework for enhancing geospatial semantic segmentation*. In: Semantic Web 10 (2019)

Applications: Medicine / Image Processing



Figure: Left: segmented cell bodies (green), lobulated cell nuclei (yellow and red) and background (black), Middle: segmented cell nucleus extending outside border of host cell (red pixels), Right: the result of applying a morphological erosion operator; in this case the original *partially overlaps* relation changes to *proper part*¹⁸

¹⁸M. Sioutis et al.: *Ordering Spatio-Temporal Sequences to Meet Transition Constraints: Complexity and Framework*. In: AIAI. 2015

Applications: Region Approximation / Emergency Response




Figure: Illustration of locating a region by natural language descriptions, e.g., “Bushfire burning about 5km northwest of Yallourn North” and “I saw fire about 10km northeast from Moe”, with the help of a region approximation method¹⁹

¹⁹Ongoing work with Southwest Jiaotong University (China) (submitted to a journal)

Applications: Drone Monitoring



Figure: “Never fly over an urban area for more than 3 minutes”: $\forall r \in \text{UrbanRegion},$
 $\Box(PO \vee TPP \vee NTPP(\text{Drone}, r) \rightarrow \Diamond_{[0,180s]} DC(\text{Drone}, r))^{20}$

²⁰F. Heintz and D. de Leng: *Spatio-Temporal Stream Reasoning with Incomplete Spatial Information*. In: ECAI. 2014

Fundamentals of Spatio-Temporal Formalisms

Part III

Qualitative Constraint Language

A binary qualitative constraint language is based on a finite set B of base relations such that:

- its base relations are defined on an infinite domain D

for example, D can be the real line;

- its base relations are *jointly exhaustive and pairwise disjoint*

for example, $X \ b \ Y$, where $b \in \{<, >, =\}$;

- B contains the identity relation Id

for example, if $B = \{<, >, =\}$, then Id is $=$;

- B is closed under the converse operation ($^{-1}$)

for example, $<^{-1}$ is $>$;

2^B expresses all relations (definite and indefinite knowledge).

Example Calculus: RCC8 (Reminder)




Figure: The base relations of RCC8; $\cdot i$ denotes the inverse of \cdot

An abundance of calculi exist for many types of spatial and temporal information (e.g., trajectories, occlusion, intervals)²¹

²¹F. Dylla et al.: *A Survey of Qualitative Spatial and Temporal Calculi: Algebraic and Computational Properties*. In: ACM Comput. Surv. 50 (2017)

Qualitative Constraint Network

Spatial or temporal information for a set of entities can be represented by a qualitative constraint network (QCN).

Definition

A QCN is a pair $\mathcal{N} = (V, C)$ where V is a non-empty finite set of variables, and C a mapping $C : V \times V \rightarrow 2^{\mathcal{B}}$.




Figure: A QCN of RCC8 along with a solution

An atomic satisfiable sub-QCN of the original QCN is a *scenario* of it.

Fundamental Reasoning Problems of QCNs (1/3)

Definition

The satisfiability checking problem of a QCN \mathcal{N} is deciding whether \mathcal{N} admits a solution.

- The satisfiability checking problem is **NP-hard** in general.

Fundamental Reasoning Problems of QCNs (2/3)

Definition

The minimal labeling problem (MLP) of a QCN \mathcal{N} is finding the strongest implied constraints of \mathcal{N} .

- The MLP is polynomial-time Turing reducible to the satisfiability checking problem




Figure: A RCC8 network (left) and its minimal network (right)

Fundamental Reasoning Problems of QCNs (3/3)

Definition

The redundancy problem of a QCN \mathcal{N} is determining if a particular constraint of \mathcal{N} is entailed by the rest of the constraints of \mathcal{N}

- Like the MLP, the redundancy problem is polynomial-time Turing reducible to the satisfiability checking problem




Figure: A RCC8 network (left) and its prime network (right)

Local Consistencies

- Approximate satisfiability
- Simplify a QCN / prune search space
- Realize forward-checking in a backtracking algorithm

Weak Composition Operation (\diamond)

The weak composition $EC \diamond NTPP$ yields the set of base relations $\{NTPP, TPP, PO\}$.




Figure: Three possible configurations for regions x, y, z when we have $EC(x, y) \diamond NTPP(y, z)$

We can precompute and store all weak composition outputs in a table in memory.

\diamond_G -Consistency: Definition

Given a QCN $\mathcal{N} = (V, C)$ and a graph $G = (V, E)$, \diamond_G -consistency enforces consistency on all triples of variables in \mathcal{N} that correspond to triangles in G .²²

Definition

Given a QCN $\mathcal{N} = (V, C)$ and a graph $G = (V, E)$, \mathcal{N} is \diamond_G -consistent if and only if $\forall \{v_i, v_j\}, \{v_i, v_k\}, \{v_k, v_j\} \in E$, we have that $C(v_i, v_j) \subseteq C(v_i, v_k) \diamond C(v_k, v_j)$.




Figure: A QCN of RCC8 that is not \diamond -consistent

²²M. Sioutis et al.: An Efficient Approach for Tackling Large Real World Qualitative Spatial Networks. Int. J. Artif. Intell. Tools 25 (2016)

\diamond_G -Consistency: Results

Proposition

Let $\mathcal{N} = (V, C)$ be a QCN defined over a subclass of relations that has patchwork for \diamond -consistent QCNs, and $G = (V, E)$ a chordal graph such that $G(\mathcal{N}) \subseteq G$. If \mathcal{N} is not trivially inconsistent and \diamond_G -consistent, then \mathcal{N} is satisfiable.

\diamond -consistency is \diamond_G -consistency where G is a complete graph




Figure: Triangulation of the underlying constraint graph of a QCN

\diamond_G -Consistency: Algorithm

Algorithm 1: $\text{PWC}(\mathcal{N}, G)$

in : A QCN $\mathcal{N} = (V, C)$ and a graph $G = (V, E)$.

output : $\diamond_G(\mathcal{N})$.

```
1 begin
2    $Q \leftarrow E$ ;
3   while  $Q \neq \emptyset$  do
4      $\{v, v'\} \leftarrow Q.\text{pop}()$ ;
5     foreach  $v'' \in V \mid \{v, v''\}, \{v', v''\} \in E$  do
6        $r \leftarrow C(v, v'') \cap (C(v, v') \diamond C(v', v''))$ ;
7       if  $r \subset C(v, v'')$  then
8          $C(v, v'') \leftarrow r$ ;
9          $C(v'', v) \leftarrow r^{-1}$ ;
10       $Q \leftarrow Q \cup \{\{v, v''\}\}$ ;
11       $r \leftarrow C(v'', v') \cap (C(v'', v) \diamond C(v, v'))$ ;
12      if  $r \subset C(v'', v')$  then
13         $C(v'', v') \leftarrow r$ ;
14         $C(v', v'') \leftarrow r^{-1}$ ;
15       $Q \leftarrow Q \cup \{\{v'', v'\}\}$ ;
16 return  $\mathcal{N}$ ;
```

Runtime: $O(\Delta(G)|E||B|)$

$\overleftarrow{\diamond}$ -Consistency: Definition

Given a QCN $\mathcal{N} = (V, C)$ and some ordering of its variables, $\overleftarrow{\diamond}$ -consistency entails consistency on all triples of variables in \mathcal{N} along that ordering.²³

Definition

Given a QCN $\mathcal{N} = (V = \{v_0, v_1, \dots, v_{n-1}\}, C)$, \mathcal{N} is $\overleftarrow{\diamond}$ -consistent iff for all $v_i, v_k, v_j \in V$ with $i, j < k$ we have that $C(v_i, v_j) \subseteq C(v_i, v_k) \diamond C(v_k, v_j)$.




Figure: A QCN of RCC8 that is not $\overleftarrow{\diamond}$ -consistent w.r.t. the ordering (z, y, x)

²³M. Sioutis et al.: Leveraging Variable Elimination for Efficiently Reasoning about Qualitative Constraints. Int. J. Artif. Intell. Tools 27 (2018)

\leftarrow -Consistency: Results (1/2)

Proposition

Let $\mathcal{N} = (V = \{v_0, v_1, \dots, v_{n-1}\}, C)$ be a QCN defined over a distributive subclass of relations over which every \diamond -consistent atomic QCN is satisfiable. If \mathcal{N} is \leftarrow -consistent and not trivially inconsistent, then \mathcal{N} is satisfiable.




Figure: Eliminating variables one by one and propagating constraints

$\leftarrow\Diamond$ -Consistency: Results (2/2)

Proposition

Let $\mathcal{N} = (V = \{v_0, v_1, \dots, v_{n-1}\}, C)$ be a $\leftarrow\Diamond$ -consistent and not trivially inconsistent QCN defined over a distributive subclass of relations over which every \Diamond -consistent atomic QCN is satisfiable. To refine \mathcal{N} to a scenario S , for each k from 1 to $n - 1$ and for each $i \in \{0, \dots, k - 1\}$ do:

- $C(v_k, v_i) \leftarrow \bigcap_{j=0}^{k-1} C(v_k, v_j) \Diamond C(v_j, v_i)$;
- $C(v_k, v_i) \leftarrow \{b\}$ for some $b \in C(v_k, v_i)$;
- $C(v_i, v_k) \leftarrow (C(v_k, v_i))^{-1}$.




Figure: Adding variables one by one and propagating constraints

\diamond -Consistency: Algorithm

Algorithm 2: $\text{DWC}(\mathcal{N}, \alpha)$

in : A QCN $\mathcal{N} = (V, C)$ with $|V| = n$, and a bijection α of V onto $\{0, 1, \dots, n-1\}$.
output : $\diamond(\mathcal{N})$.

```
1 begin
2    $G \leftarrow (V, E = E(G(\mathcal{N})))$ ;
3   for  $x$  from  $n-1$  to 1 do
4      $v \leftarrow \alpha^{-1}(x)$ ;
5      $\text{adj} \leftarrow \{v' \mid \{v', v\} \in E \wedge \alpha(v') < \alpha(v)\}$ ;
6     foreach  $v', v'' \in \text{adj}$  do
7       if  $\alpha(v') < \alpha(v'')$  then
8         if  $\{v', v''\} \notin E$  then
9            $E \leftarrow E \cup \{\{v', v''\}\}$ ;
10         $t \leftarrow C(v', v'') \cap (C(v', v) \diamond C(v, v''))$ ;
11        if  $t \neq C(v', v'')$  then
12           $C(v', v'') \leftarrow t$ ;
13           $C(v'', v') \leftarrow t^{-1}$ ;
14 return  $\mathcal{N}$ ;
```

Runtime: $O(\Delta(G)^2|V|)$

Efficient Algorithm for the Satisfiability Problem of QCNs

Given a QCN $\mathcal{N} = (V, C)$, a chordal graph $G = (V, E) \supseteq G(\mathcal{N})$, and a subclass \mathcal{A} for which \diamond_G -consistency is complete, we aim to characterize a QCN \mathcal{N}' such that:

$$\mathcal{N}' \subseteq \mathcal{N} \text{ and } \mathcal{N}' = \diamond_G(\mathcal{A}(\mathcal{N}'))$$

Solving via backtracking search

Let us consider a QCN \mathcal{N} , we tackle it as follows.

- Every relation r forming a constraint in \mathcal{N} is split into subrelations $r' \subseteq r$.
- These subrelations r' belong to a set of relations \mathcal{A} over which the QCN becomes tractable.
- After every refinement of a relation r into one of its subrelations r' , a validity check is performed:
 - if the refinement is valid, we proceed with the next one;
 - if the refinement is not valid, we backtrack to the previous one.
- When \mathcal{N} has been refined into relations of \mathcal{A} , we stop.

The Consistency Algorithm

Algorithm 3: Consistency(\mathcal{N} , G , \mathcal{A})

in : A QCN $\mathcal{N} = (V, C)$, a graph $G = (V, E)$, and a subclass \mathcal{A} of $2^{\mathbf{B}}$.

output : Null, or a refinement of network \mathcal{N} over \mathcal{A} .

```
1 begin
2    $\mathcal{N} \leftarrow \text{PWC}(\mathcal{N}, G)$ ;
3   if  $\exists \{v_i, v_j\} \in E$  such that  $C(v_i, v_j) = \emptyset$  then
4     return Null;
5   if  $\forall \{v_i, v_j\} \in E$  we have that  $C(v_i, v_j) \in \mathcal{A}$  then
6     return  $\mathcal{N}$ ;
7   choose a constraint  $C(v_i, v_j)$  with  $\{v_i, v_j\} \in E$  such that  $C(v_i, v_j) \notin \mathcal{A}$ ;
8   split  $C(v_i, v_j)$  into  $r_1, \dots, r_k \in \mathcal{A}$ :  $r_1 \cup \dots \cup r_k = C(v_i, v_j)$ ;
9   foreach  $r \in \{r_l \mid 1 \leq l \leq k\}$  do
10     $\mathcal{N}[v_i, v_j] \leftarrow r$ ;  $\mathcal{N}[v_j, v_i] \leftarrow r^{-1}$ ;
11    result  $\leftarrow \text{Consistency}(\mathcal{N}, G, \mathcal{A})$ ;
12    if result  $\neq$  Null then
13      return result;
14  return Null;
```

A QCN as a SAT instance

Let $\mathcal{N} = (V, C)$ be a QCN instance, make variable x_{ij}^b for every $v_i, v_j \in V$ and $b \in C(v_i, v_j)$, then make encoding as follows²⁴

at least one constraint: $x_{ij}^{b_1} \vee x_{ij}^{b_2} \dots \vee x_{ij}^{b_{|C(v_i, v_j)|}}$

at most one constraint: $\neg(x_{ij}^{b_k} \wedge x_{ij}^{b_l}), \forall$ distinct $b_k, b_l \in C(v_i, v_j)$

weak composition: $(x_{ij}^{b_m} \wedge x_{ij}^{b_n}) \rightarrow (x_{ij}^{b_1} \vee x_{ij}^{b_2} \dots \vee x_{ij}^{b_w})$

²⁴D. N. Pham et al.: *Towards an efficient SAT encoding for temporal reasoning*. In: CP. 2006

A QCN as an ASP instance

Let $\mathcal{N} = (V, C)$ be a QCN instance over the set of base relation B , then make encoding as follows²⁵

base relation declaration: $\{rel(R).\}, \forall R \in B$

choice rule: $1\{label(U, V, L) : rel(L)\}1 :- variable1(U), variable2(V), U < V.$

QCN size: $variable1(0..|V| - 1). variable2(0..|V| - 1).$

weak composition: $\{:- label(X, Y, R_1), label(Y, Z, R_2), label(X, Z, R_3).\}$

integrity constraint: $\{:- label(X, Y, R).\} \quad // R \notin C(X, Y)$

²⁵J. J. Li: *Qualitative Spatial and Temporal Reasoning with Answer Set Programming*. In: ICTAI. 2012

A Case of Robustness in AI

- We want to have some measure of the quality of a spatial or temporal configuration, viz., *scenario*.²⁶
- We focus on scenarios that can withstand perturbations, i.e., that limit as much as possible the need for successive repairs.




Figure: Tasks scheduled in a factory; as there might be unpredictable incidents concerning resource availability (e.g., power outage), how should the factory schedule production so that it may be least disturbed?

²⁶M. Sioutis et al.: *On Robustness in Qualitative Constraint Networks*. In: IJCAI. 2020

Definition

Given a network \mathcal{N} , a scenario \mathcal{S} of \mathcal{N} is robust iff:

$$\mathcal{S} \in \arg \max_{\mathcal{S}' \in [[\mathcal{N}]]} \text{similarity}(\mathcal{S}', [[\mathcal{N}]])$$

where

$$\text{similarity}(\mathcal{N}, \mathfrak{M}) = \frac{\sum_{\mathcal{N}' \in \mathfrak{M}} \# \text{sameConstraints}(\mathcal{N}, \mathcal{N}') / |\mathcal{N}'|}{|\mathfrak{M}|}$$

Intuitively, a robust scenario of a network \mathcal{N} has the largest number of common constraints on average with each scenario of \mathcal{N} , i.e., with each satisfiable atomic network in $[[\mathcal{N}]]$.

Revisiting the Task Scheduling Example




Figure: Tasks scheduled in a factory, revisited

- Choosing *follows* over *precedes* between Tasks A and B, maintains the whole range of possibilities between Tasks A and C.
- Choosing *precedes* over any other relation between Tasks A and C, can resist the change of *follows* switching back to *precedes* between Tasks A and B.

Evaluation: Perturbation Tolerance Comparison

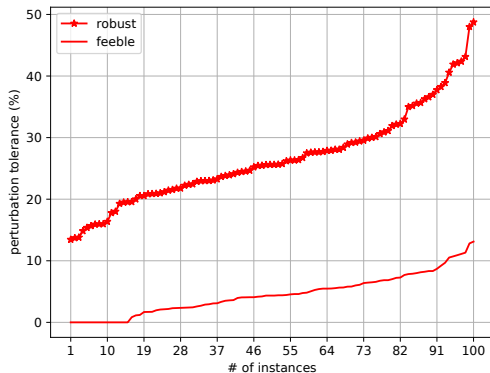


Figure: Perturbation tolerance comparison with job-shop dataset

Useful Links

- Python Qualitative Constraint-based Tools

<https://msioutis.gitlab.io/software/>

- GQR: C++ Qualitative Constraint-based Reasoner

<https://www.sfbtr8.spatial-cognition.de/en/project/reasoning/r4-logospace/research-tools/gqr/index.html>

- SparQ: A Toolbox for Qualitative Spatial and Temporal Reasoning

<https://www.uni-bamberg.de/en/sme/research/sparq/>

Future Research Topics

Part IV

Enabling Abductive Learning

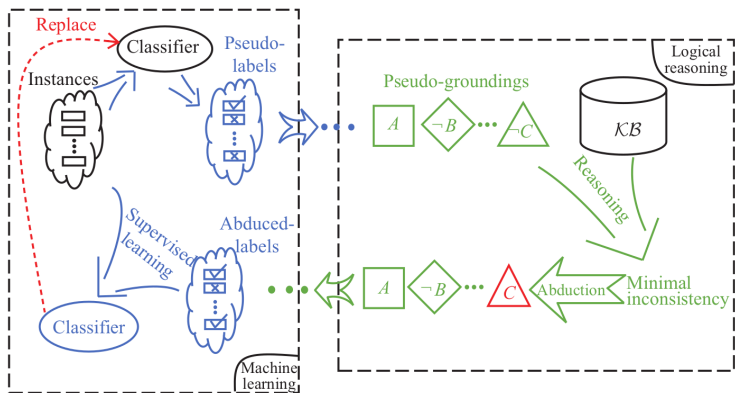


Figure: Abductive Learning framework²⁷

²⁷Z.-H. Zhou: *Abductive learning: towards bridging machine learning and logical reasoning*. In: Sci. China Inf. Sci. (2019)

Enabling Abductive Learning: Underlying Challenges (1/2)

- Adaptive and dynamic algorithms for runtime verification of neural network-based components
- Parsimonious local consistencies (e.g., restricted to neighbourhood²⁸)
- Dynamic algorithms for repair (revision), inconsistency minimization

²⁸M. Sioutis et al.: *On neighbourhood singleton-style consistencies for qualitative spatial and temporal reasoning*. In: Inf. Comput. (2020)

Enabling Abductive Learning: Underlying Challenges (2/2)

- Relation-aware and/or lazy tree decompositions
- Backbone-based decoupling and backdoor-based search space splitting approaches²⁹ (reasoning shortcuts)
- Exploitation of AI hardware (e.g., GPUs)

²⁹M. Sioutis and T. Janhunen. *Towards Leveraging Backdoors in Qualitative Constraint Networks*. In: KI. Short paper. 2019

Hybrid AI Systems Grounded on QSTR (1/3)

Neuro-symbolic AI

- Neuro-symbolic formulae:³⁰

$X^{(95\% \text{ yolk})}$ is contained in^(45% true) or overlaps^(55% true) $Y^{(90\% \text{ egg})}$

$(\text{Train } X \{canUse\} (\text{Track } A \oplus \text{Track } B)) \wedge (\text{Train } Y \{canUse\} \text{Track } B) \wedge$
 $(\text{departInt}(\text{Train } X) \{precedes 10\%, \text{meets } 30\%, \text{overlaps } 60\%\} \text{departInt}(\text{Train } Y))$

$\square(\neg \text{blocked}(\text{Track } A) \wedge \neg \text{blocked}(\text{Track } B))$

Update: $(\Diamond \text{blocked}(\text{Track } A))$ 90%

³⁰M. Sioutis and D. Wolter: *Qualitative Spatial and Temporal Reasoning: Current Status and Future Challenges*. In: IJCAI. Survey paper. 2021

Hybrid AI Systems Grounded on QSTR (2/3)

Hierarchical portfolio of solvers

- QSTR information can be encoded/tackled in different ways (aside from native QSTR approaches), e.g.,:
 - as SAT or even traditional CSP instances³¹;
 - and as Answer Set Programming (ASP) instances too^{32,33}
- Automate choice/combination of tools for a given task

³¹M. Westphal and S. Wöfl.: *Qualitative CSP, Finite CSP, and SAT: Comparing Methods for Qualitative Constraint-based Reasoning*. In: IJCAI. 2009

³²G. Baryannis et al.: *A Generalised Approach for Encoding and Reasoning with Qualitative Theories in Answer Set Programming*. In: Theory Pract. Log. Program. 20 (2020)

³³T. Janhunen and M. Sioutis.: *Allen's Interval Algebra Makes the Difference*. In: DECLARE. 2019

Hybrid AI Systems Grounded on QSTR (3/3)

Data mining

- Knowledge discovery in spatio-temporal data
- Removal of redundancy from spatio-temporal KBs
- Pattern discovery / identification at run-time

Final Remarks

- QSTR, within Neuro-Symbolic AI especially, is a research area relevant for many AI application domains (e.g., in defence or medicine)
- It extends from Theoretical Computer Science to Practical Applications, and vice versa
- Important to consider modular architectures for AI that are grounded on *physics* and *human cognition*

Acknowledgements

- Dr. Zhiguo Long (Southwest Jiaotong University, Chengdu, China)

<https://zhiguolong.github.io/>

- IJCAI Spatio-Temporal Reasoning and Learning Workshop

<https://str12022.github.io/>

Thank you for your interest and attention!

`http://msioutis.gitlab.io`

The purpose of abstraction is not to be vague, but to create a new semantic level in which one can be absolutely precise

Dijkstra