On Neuro-Symbolic Challenges in Directional Relation Prediction

Junjie Peng[1], Michael Sioutis[2,*] and Zhiguo Long[1]

[1]School of Computing and Artificial Intelligence, Southwest Jiaotong University, Chengdu 611756, P.R. China
[2]LIRMM UMR 5506, University of Montpellier & CNRS, France

Abstract
In this paper we discuss some neuro-symbolic challenges that exist in combining a machine learning model and a symbolic reasoning framework for directional relation prediction. In particular, we consider a recent machine learning approach that predicts the qualitative directional relations between geographical regions, e.g., *X is north-west of Y*, where each region is a polygon of boundary points, and highlight the challenges of aligning these predicted relations with the inference rules of a well-known qualitative spatial calculus, viz., the Cardinal Direction Calculus.

Keywords
directional relation prediction, neuro-symbolic artificial intelligence, machine learning, qualitative spatial reasoning

1. Introduction

Neuro-Symbolic Artificial Intelligence is a paradigm that deals with the combination of Machine Learning models and Logic-based frameworks; this combination should ideally lead to unified architectures that aim to collaboratively utilize both components to their fullest extent possible. Due to its diverse and human-like nature that involves data-driven inference and logical reasoning, as well as its promise in handling problems that pertain to both large amounts of data and knowledge-based rules, Neuro-Symbolic Artificial Intelligence is an important and re-surging topic of research [1, 2, 3, 4, 5, 6]. A classification of neuro-symbolic approaches is provided in [7, Figure 24]; here, we focus on the class of *architectures integrating learning and reasoning*, a high-level illustration of which is shown in Figure 1.

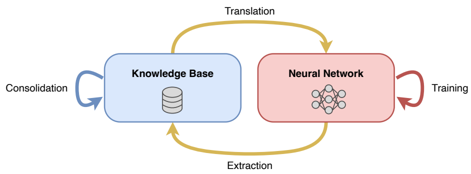

Figure 1: Cyclical interaction in neuro-symbolic AI as depicted in [8]; a symbolic framework provides symbolic (partial) knowledge to a machine learning model, which is trained on raw data, and (any) knowledge acquired through the model can then be extracted and fed back to the symbolic framework for further processing.

In this paper, we discuss some challenges, as well as ways of addressing these challenges, that arise when trying to join together a machine learning model and a symbolic reasoning framework for directional relation prediction, such as X *is north-west of* Y; these challenges pertain to problems that arise during this fusion of the two paradigms in the context of qualitative directional relation prediction. Specifically, we consider a machine learning model for directional relation prediction from the recent literature that views each region as a polygon of boundary points and then treats the task of predicting the directional relation between two regions as a multi-label classification problem [9]; each label corresponds to a subset of the eight specific direction relations, viz., $N, NE, E, SE, S, SW, W, NW$ (see also Figure 3). The reasons to use machine learning instead of direct calculations to get directional relations from geometric data are as follows: (1) direct calculations cannot automatically adapt to varying interpretations of directional relations but machine learning can; (2) the output of a machine learning model for directional relation prediction can be quantitative and thus easier to combine with other tasks that involve quantitative inputs. At the other end, we exploit a well-known calculus, viz., the Cardinal Direction Calculus [10, 11], from the area of Qualitative Spatio-Temporal Reasoning (QSTR) [12, 13]. In sum, QSTR allows one to spatially or temporally relate one object with another object or oneself by using everyday, human-like, natural language descriptions, and perform reasoning with those descriptions. In the case of the Cardinal Directional Calculus these natural language descriptions are

STRL'23: Second International Workshop on Spatio-Temporal Reasoning and Learning, 21 August 2023, Macao, S.A.R
*Corresponding author.
2022212423@my.swjtu.edu.cn (J. Peng); michael.sioutis@lirmm.fr (M. Sioutis); zhiguolong@swjtu.edu.cn (Z. Long)
https://msioutis.gitlab.io/ (M. Sioutis); https://zhiguolong.github.io/ (Z. Long)
0009-0004-9906-8263 (J. Peng); 0000-0001-7562-2443 (M. Sioutis); 0000-0002-2714-3453 (Z. Long)

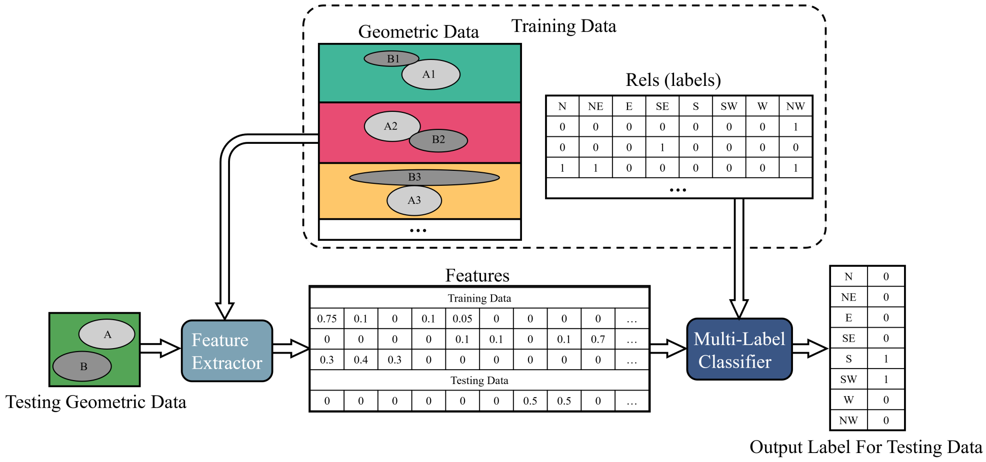

Figure 2: The training and testing pipeline of the machine learning model for directional relation prediction of [9].

nothing more than the eight directional relations that we mentioned earlier, viz., $N, NE, E, SE, S, SW, W, NW$. However, unlike the machine learning model, which performs statistical inference, the Cardinal Direction Calculus comes with its own logic-based inference rules, and aligning the two is part of the discussion in the sequel.

The rest of the paper is organized as follows. In Section 2 we summarize the machine learning model for direction relation prediction of [9] and introduce the theory behind QSTR and, in particular, the Cardinal Direction Calculus. Then, in Section 3 we introduce and expand on the neuro-symbolic research opportunities / challenges that exist when trying to align the statistical inference of the machine learning model with the logic-based inference of the symbolic one. Finally, in Section 4 we conclude the paper and provide a discussion about possible future directions of work.

2. Background

2.1. Machine Learning-based Directional Relation Prediction

In [9], the authors discuss how to predict qualitative directional relations between geographical regions by using machine learning techniques, where each region is represented as a polygon formed by a sequence of boundary points. Figure 2 introduces the overall idea of the model, including its training and testing process.

In particular, the authors of [9] model the problem of predicting qualitative directional relations between regions as a multi-label classification problem. Each label correspond to one of the eight specific directions, i.e., $N, NE, E, SE, S, SW, W, NW$, and there might be multiple directional relations between two regions, e.g., $\{NW, N\}$ is encoded as $(1, 0, 0, 0, 0, 0, 0, 1)$. The training geometric data are formed of pairs of polygons (A_i, B_i), where each pair represents a reference region and a target region. Geometric data are pre-processed with a hand-crafted feature extractor to extract quantitative features, such as angles, areas, intersections with regions of acceptance, etc. The labels are from binary encoding of given qualitative directional relations from Wikipedia, e.g., NW is encoded as $(0, 0, 0, 0, 0, 0, 0, 1)$. These training data are then used to train the machine learning (ML) model. For a new pair of polygons (A, B), the qualitative directional relation between A and B can be predicted by the ML model via feeding to the trained ML model their quantitative features obtained using the same feature extractor.

2.2. Qualitative Spatio-Temporal Reasoning

To facilitate discussion, we first recall the formal definition of a *qualitative constraint language*, which is a constraint language that is used to represent and reason about qualitative information. A binary qualitative spatial or temporal constraint language is based on a finite set B of *jointly exhaustive and pairwise disjoint* relations, called *base relations* [13] and defined over an infinite domain D. The base relations of a particular qualitative constraint language can be used to represent the definite knowledge between any two of its entities with respect to the level of granularity provided by the domain D. The set B contains the identity relation Id, and is closed under the *converse* operation ($^{-1}$). Indefinite knowledge can be specified by a union of possible

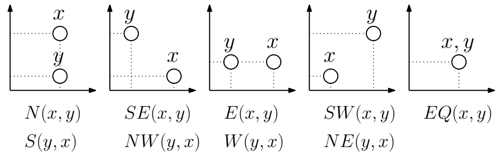

Figure 3: The base relations of Cardinal Direction Calculus.

base relations, and is represented by the set containing them. Hence, 2^{B} represents the total set of relations. The set 2^{B} is equipped with the usual set-theoretic operations of union and intersection, the converse operation, and the *weak composition* operation denoted by the symbol $\diamond$ [13]. For all $r \in 2^{\mathsf{B}}$, we have that $r^{-1} = \bigcup\{b^{-1} \mid b \in r\}$. The weak composition ($\diamond$) of two base relations $b, b' \in \mathsf{B}$ is defined as the smallest (i.e., most restrictive) relation $r \in 2^{\mathsf{B}}$ that includes $b \circ b'$, or, formally, $b \diamond b' = \{b'' \in \mathsf{B} \mid b'' \cap (b \circ b') \neq \emptyset\}$, where $b \circ b' = \{(x, y) \in \mathsf{D} \times \mathsf{D} \mid \exists z \in \mathsf{D}$ such that $(x, z) \in b \wedge (z, y) \in b'\}$ is the (true) composition of b and b'. For all $r, r' \in 2^{\mathsf{B}}$, we have that $r \diamond r' = \bigcup\{b \diamond b' \mid b \in r,\ b' \in r'\}$. We note that the weak composition operation is of particular importance as it is the core operation used to perform some basic inference/reasoning, e.g., inside$(x, z) \diamond$ inside$(z, y) \rightarrow$ inside(x, y), and consequently check the admissibility/validity of large complex networks defined over disjunctions of base relations [14, 15].

Cardinal Direction Calculus

Let us first introduce the qualitative temporal constraint language of Point Algebra (PA) [16, 17, 18], which uses points to represent temporal entities (e.g., events) and the following three base relations to reason about the relative position of those temporal entities in the timeline: *precedes* ($<$), *equals* ($=$), and *follows* ($>$). These three base relations considered by Point Algebra are interpreted on a set with a linear ordering relation. In particular, considering the points on the line of rational numbers and the usual ordering relation $<$, the three base relations of Point Algebra are defined in the following manner: *precedes* $= \{(x, y) \in \mathbb{Q} \times \mathbb{Q} \mid x < y\}$, *follows* $= \{(x, y) \in \mathbb{Q} \times \mathbb{Q} \mid y < x\}$, and *equals* $= \{(x, y) \in \mathbb{Q} \times \mathbb{Q} \mid x = y\}$. Based on these three base relations, we can define eight relations of Point Algebra in total that correspond to the set $2^{\mathsf{B}} = \{\{<, =, >\}, \{<, >\}, \{<, =\}, \{=, >\}, \{<\}, \{>\}, \{=\}, \emptyset\}$. As an example, relation $\{<, >\}$ allows us to represent the knowledge that an event occurs before or after another event, but not at the same time. Further, two events $x, y \in \mathbb{Q}$ satisfy relation $\{<, >\}$ if and only if $x \neq y$.

Now, the Cardinal Direction Calculus (CDC) [10, 11] is a qualitative constraint language with a spatial aspect and can be seen as an extension of the qualitative constraint language of Point Algebra discussed earlier. The entities of the domain D are points in the Euclidean plane and are equipped with an orthogonal reference. In particular, the relative position between two entities is determined by the Point Algebra base relations that are derived from projections of those points on the two axes. As such, we obtain nine possible base relations between two given entities x and y in Cardinal Direction Calculus, namely, east (E), north (N), south (S), west (W), northeast (NE), northwest (NW), southwest (SW), southeast (SE), and equals (EQ). These base relations can be viewed in Figure 3.

Remark Generally, one can even use a rectangle, e.g., denoting a 2-dimensional minimum bounding box, to approximate a region, instead of a single point, and thus perform reasoning on directional relations about rectangles [19, 20]. As the main purpose of this paper is to demonstrate the challenges of integrating learning and reasoning, we chose the simpler approximation to better convey the essential idea.

3. Challenges

Previously, a machine learning model [9] was proposed to predict the directional relations between geographic regions (see Section 2.1. However, the machine learning model did not consider the semantic connections between different relations and between different pairs of regions, and may have issues such as missing or conflicting directional relations in these predictions, as illustrated in the forthcoming examples, which are taken from the actual testing data of [9].

This article proposes to use qualitative spatial reasoning to identify, add, or modify directional relation networks that have already been obtained in order to enrich the network information and make the network more complete and accurate.

In what follows, the universal constraint of a calculus, which corresponds to the entire set of base relations B of that calculus, will be denoted by $\star$ to avoid ambiguity between what is a constraint and what is the signature of the calculus, respectively (even though they are the exact same relation).

3.1. Information Refining

Filling missing relations

In the prediction results of the machine learning model in Section 2.1, sometimes there only exists the prediction of the directional relation from region j to region i, but the relation from region i to region j is missing, i.e., $R_{ij} \neq \star$ and $R_{ji} = \star$. In this case, we can obtain an approximation of R_{ji} by taking the inverse of R_{ij}:

$$R_{ji} \leftarrow R_{ij}^{-1}.$$

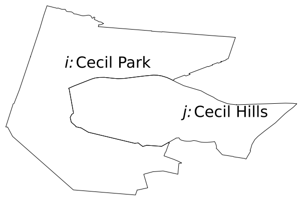

Figure 4: An actual example illustrating filling missing relations in predictions by taking inverse.

Figure 4 gives a such example. In this figure, the predicted R_{ij} is $\{E\}$, i.e., region j is on east of region i. However, the machine learning model did not predict R_{ji}. By taking the inverse of R_{ij}, we can directly get $R_{ji} = R_{ij}^{-1} = \{SW, W, NW\}$, meaning that the relation of j w.r.t. i can be SW, W, or NW.

Sometimes there exists the prediction of the directional relation from region i to region j and the directional relation from region j to region k, but the relation from region i to region k is missing, i.e., $R_{ij} \neq \star$ and $R_{jk} \neq \star$ and $R_{ik} = \star$. In this case, we can obtain an approximation of R_{ik} by composing R_{ij} and R_{jk}:

$$R_{ik} \leftarrow R_{ij} \diamond R_{jk}.$$

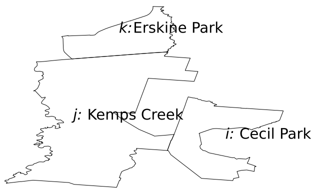

Figure 5: An actual example illustrating filling missing relations in predictions by taking composition.

For instance, in Figure 5, the predicted R_{ij} is $\{W\}$ and R_{jk} is $\{N\}$. However, the machine learning model did not predict R_{ik}. By composing R_{ij} and R_{jk}, we can directly get $R_{ik} = R_{ij} \diamond R_{jk} = \{N, W, NW\}$.

Removing unfeasible relations

Sometimes there is no absence of a relation, but there is a contradiction between R_{ij} and the composition of R_{ik} and R_{kj}, i.e., $R_{ij} \neq \star$ and $R_{ik} \neq \star$ and $R_{kj} \neq \star$ and $R_{ij} \neq R_{ij} \cap (R_{ik} \diamond R_{kj})$. In this case, we can obtain an approximation of R_{ij} by taking the intersection of $R_{ik} \diamond R_{kj}$ and R_{ij}:

$$R_{ij} \leftarrow R_{ij} \cap (R_{ik} \diamond R_{kj}).$$

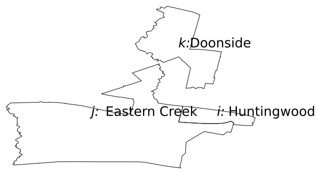

Figure 6: An actual example illustrating removing impossible relations from predictions.

As an illustration, in Figure 6, the predicted R_{ij} is $\{S, SW, W\}$ and R_{ik} is $\{NW\}$ and R_{kj} is $\{SW\}$. However, the composition of R_{ik} and R_{kj} is $\{SW, W, NW\} \neq R_{ij}$, so we can update $R_{ij} = R_{ij} \cap (R_{ik} \diamond R_{kj}) = \{SW, W\}$.

3.2. Inconsistency handling

Sometimes there is a contradiction between the predicted R_{ij} and the composition of R_{ik} and R_{kj}, i.e., $R_{ij} \neq \star$ and $R_{ik} \neq \star$ and $R_{kj} \neq \star$ and $R_{ij} \cap R_{ik} \diamond R_{kj} = \emptyset$. In this case, we can resolve contradiction by replacing R_{ij} with the composition of R_{ik} and R_{kj}:

$$R_{ij} \leftarrow R_{ik} \diamond R_{kj}.$$

Figure 7: An actual example illustrating resolving inconsistency by taking composition.

For example, in Figure 7, the predicted R_{ik} is $\{W\}$ and R_{kj} is $\{S\}$ and R_{ij} is $\{SE\}$ and $R_{ij} \cap R_{ik} \diamond R_{kj} = \{SE\} \cap \{SW\} = \emptyset$. So we can resolve the inconsistency by setting $R_{ij} = R_{ik} \diamond R_{kj} = \{SW\}$.

There can also be a contradiction between the predicted R_{ij} and R_{ji}, i.e., $R_{ij} \neq \star$ and $R_{ji} \neq \star$ and $R_{ij} \cup R_{ji}^{-1} \neq R_{ji}^{-1}$ and R_{ij} looks more reasonable (probably based on criteria including the area in regions of acceptance, the angle of the line connecting center points, etc.). In this case, we can obtain an approximation of R_{ij} by taking the inverse of R_{ij}:

$$R_{ji} \leftarrow R_{ij}^{-1}.$$

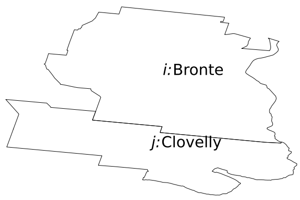

Figure 8: An actual example illustrating resolving inconsistency in predictions by taking inverse.

In Figure 8, the predicted R_{ji} is $\{N\}$ and R_{ij} is $\{W\}$ and R_{ji} is more reasonable. So we can directly get $R_{ij} = R_{ji}^{-1} = \{SW, S, SE\}$.

4. Discussion

In this paper we discussed some neuro-symbolic challenges that arise when trying to combine a machine learning model and a symbolic reasoning framework for directional relation prediction. Specifically, on one hand, we considered the machine learning approach of [9] that predicts the qualitative directional relations between geographical regions, e.g., *X is north-west of Y*, and, on the other hand, we employed the symbolic framework of the Cardinal Direction Calculus to capture and reason with those predicted relations [10, 11].

It is important to note that we just initiated the discussions by presenting several example cases where a symbolic reasoning framework can help with the predictions of a machine learning model. Much more work can be done in the future, e.g., when inconsistency is detected by composition or inverse, how to determine which predicted relations are more plausible is an important yet insufficiently researched topic. How to exploit the predictions of a machine learning model to perform symbolic reasoning better is also very interesting. For the problem considered in this paper, an implicit assumption is that the semantics of symbolic reasoning matches the semantics of predicted relations, which in real-world applications is seldom the case. As have been discussed in [8], machine learning predictions can help symbolic reasoning frameworks build reasoning rules that are consistent with the observations in real-world. Automatically discovering conceptual neighbourhood graphs (CNGs) in [21] is a good start, but there is still a huge gap between reasoning and CNGs. Finally, the type of integration between the machine learning model and the logical component remains open to discussion; in the future, we would like to tackle this via abductive reasoning, utilizing the neuro-symbolic framework proposed in [22].

Acknowledgements

The work was partially funded by the Agence Nationale de la Recherche (ANR) for the "Hybrid AI" project that is tied to the chair of Dr. Sioutis, and the I-SITE program of excellence of Université de Montpellier that complements the ANR funding.

References

[1] L. G. Valiant, Knowledge Infusion: In Pursuit of Robustness in Artificial Intelligence, in: IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, 2008.

[2] Michael L. Littman, Ifeoma Ajunwa, Guy Berger, Craig Boutilier, Morgan Currie, Finale Doshi-Velez, Gillian Hadfield, Michael C. Horowitz, Charles Isbell, Hiroaki Kitano, Karen Levy, Terah Lyons, Melanie Mitchell, Julie Shah, Steven Sloman, Shannon Vallor, Toby Walsh, Gathering Strength Gathering Storms: The One Hundred Year Study on Artificial Intelligence (AI100) 2021 Study Panel Report, Technical Report, Stanford University, 2021.

[3] P. Hitzler, M. K. Sarker (Eds.), Neuro-Symbolic Artificial Intelligence: The State of the Art, volume 342 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2021.

[4] A. S. d'Avila Garcez, T. R. Besold, L. D. Raedt, P. Földiak, P. Hitzler, T. Icard, K.-U. Kühnberger, L. C. Lamb, R. Miikkulainen, D. L. Silver, Neural-Symbolic Learning and Reasoning: Contributions and Challenges, in: AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches., 2015.

[5] S. Wermter, R. Sun, An Overview of Hybrid Neural Systems, in: Hybrid Neural Systems, 1998.

[6] K. McGarry, S. Wermter, J. MacIntyre, Hybrid Neural Systems: From Simple Coupling to Fully Integrated Neural Networks, Neural Computing Surveys 2 (1999) 62–93.

[7] S. Badreddine, A. d'Avila Garcez, L. Serafini, M. Spranger, Logic Tensor Networks, Artif. Intell. 303 (2022) 103649.

[8] J. H. Lee, M. Sioutis, K. Ahrens, M. Alirezaie, M. Kerzel, S. Wermter, Neuro-Symbolic Spatio-Temporal Reasoning, CoRR abs/2211.15566 (2022). URL: https://doi.org/10.48550/arXiv.2211.15566.

[9] Z. Long, J. Peng, H. Meng, M. Sioutis, H. Kou, Flexible directional relationships, Manuscript under submission to a journal (2023).

[10] G. Ligozat, Reasoning about cardinal directions, J. Vis. Lang. Comput. 9 (1998) 23–44.

[11] A. U. Frank, Qualitative Spatial Reasoning with Cardinal Directions, in: ÖGAI, 1991.

[12] F. Dylla, J. H. Lee, T. Mossakowski, T. Schneider, A. V. Delden, J. V. D. Ven, D. Wolter, A Survey of Qualitative Spatial and Temporal Calculi: Algebraic and Computational Properties, ACM Comput. Surv. 50 (2017) 7:1–7:39.

[13] G. Ligozat, Qualitative Spatial and Temporal Reasoning, ISTE, Wiley, 2013.

[14] Z. Long, M. Sioutis, S. Li, Efficient Path Consistency Algorithm for Large Qualitative Constraint Networks, in: IJCAI, 2016.

[15] M. Sioutis, Z. Long, S. Li, Leveraging Variable Elimination for Efficiently Reasoning about Qualitative Constraints, Int. J. Artif. Intell. Tools 27 (2018) 1860001.

[16] M. Vilain, H. Kautz, P. van Beek, Constraint Propagation Algorithms for Temporal Reasoning: A Revised Report, in: D. S. Weld, J. de Kleer (Eds.), Readings in Qualitative Reasoning About Physical Systems, Morgan Kaufmann, 1990, pp. 373–381.

[17] P. van Beek, R. Cohen, Exact and approximate reasoning about temporal relations, Comput. Intell. 6 (1990) 132–144.

[18] P. van Beek, Reasoning About Qualitative Temporal Information, Artif. Intell. 58 (1992) 297–326.

[19] D. Papadias, Y. Theodoridis, T. K. Sellis, M. J. Egenhofer, Topological Relations in the World of Minimum Bounding Rectangles: A Study with R-trees, in: SIGMOD, 1995.

[20] P. Balbiani, J. Condotta, L. F. del Cerro, A New Tractable Subclass of the Rectangle Algebra, in: IJCAI, 1999.

[21] L. Cai, K. Janowicz, R. Zhu, Automatically Discovering Conceptual Neighborhoods Using Machine Learning Methods, in: COSIT, 2022.

[22] Z. Zhou, Abductive learning: towards bridging machine learning and logical reasoning, Sci. China Inf. Sci. 62 (2019) 76101:1–76101:3.