Towards Robust Qualitative Spatio-Temporal Reasoning for Hybrid AI Systems

1st Michael Sioutis
Faculty of Information Systems and Applied Computer Sciences
Bamberg University
Bamberg, Germany
michael.sioutis@uni-bamberg.de

2nd Hua Meng
School of Mathematics
Southwest Jiaotong University
Chengdu, China
menghua@swjtu.edu.cn

***Abstract*—In this short challenge paper, we argue for the need of having *robust* Qualitative Spatio-Temporal Reasoning (QSTR), as a means to develop hybrid AI systems involving computations with spatio-temporal information. In short, QSTR is a Symbolic AI framework for representing and reasoning about spatial and temporal information via the use of disjunctive natural relations, e.g., "Task A is scheduled *after* or *during* Task C", and *robustness* entails a notion of resistance to the possible future alterations of a spatio-temporal configuration. So far, robustness for QSTR has been defined in the literature in a very rigid manner, in that a configuration may either resist a perturbation or not. Here, we propose an alternative formulation based on probabilities, that should allow a spatio-temporal configuration to also be pliable and adapt to enforced changes. We close the paper by giving some examples of how robust QSTR can be the backbone of hybrid AI systems, emphasizing on abductive learning and AI planning in particular.**

***Index Terms*—Qualitative constraints, spatial and temporal reasoning, robustness, hybrid AI**

I. Introduction

Qualitative Spatial and Temporal Reasoning, QSTR for short, is a major area of research in AI that deals with the fundamental cognitive concepts of space and time in an abstract, human-like manner, ranging from theoretical computer science and logic to practical algorithms and applications [1]. In brief, QSTR simplifies complex mathematical theories that revolve around spatial and temporal entities to manageable qualitative constraint languages (calculi), which can in turn give rise to interpretable spatio-temporal representations, typically viewed as constraint networks of disjunctions of *atoms* like *inside*, *precedes*, or *north of* [1]; see Figure 1 for an illustration of some standard QSTR terminology. Thus, QSTR forms a concise paradigm for dealing with entities pertaining to space and time with the potential to boost research in a plethora of domains such as dynamic GIS [2], cognitive robotics [3], deep learning [4], qualitative model generation from video [5], visual sensemaking [6], data mining [7], and qualitative case-based reasoning and learning [8]. The interested reader can obtain a more descriptive review of the emergent applications, the trends, and the possible future directions of QSTR in [9] and [10]. In addition, qualitative spatial and temporal calculi are surveyed in detail in [11]; the related fundamental reasoning tasks of such calculi, such as constraint network satisfiability checking, are typically NP-hard [11].

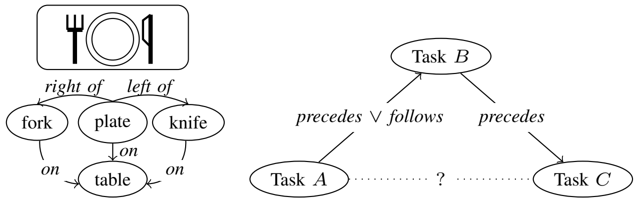

Fig. 1: A qualitative description/interpretation of a concrete spatial scene, and a simplified qualitative constraint network involving 3 temporal tasks; ? stands for complete uncertainty

As a first illustration, Allen proposed Interval Algebra in [12] as a constraint language for representing and reasoning about time in a qualitative manner. Allen's motivation at the time was to have a framework in the context of natural language processing to reliably and efficiently enough deal with the occurring temporal information in everyday communication; thus, he considered the qualitative relations between time intervals (e.g., *before*). In particular, Interval Algebra uses intervals on the real line to represent entities corresponding to events, actions, or tasks, and also encodes the possible relations between these intervals. Interval Algebra is considered today as one of the most prominent and well-known qualitative constraint languages, by virtue of its extensive use in various and diverse applications. Typical examples of such applications involve planning [13], temporal databases [14], [15], natural language processing [16], molecular biology (e.g., arranging DNA fragments along a linear chain involves certain temporal-like problems) [17], workflow [18], and intensive care medicine [19].

As an additional illustration, Randell et al. in [20] proposed the Region Connection Calculus (RCC) for qualitatively represening and reasoning about space in a mereotopological sense. Specifically, this theory considers regions in any arbitrary topological space, as well as the possible relations among such regions, and is grounded on the primitive relation of connection. For example, the base relation *partially overlaps* suggests

This work was supported by the National Natural Science Foundation of China (61806170), the Humanities and Social Sciences Fund of Ministry of Education (18XJC72040001), and the National Key Research and Development Program of China (2019YFB1706104).

that part of some region x connects with part of some other region y. Two derived calculi of RCC, namely, RCC-8 and RCC-5 (a simplification of RCC-8 where region boundaries are not accounted for), have been used in various real-life applications. In particular, in [21] RCC-5 has been used in smart cars, in [22] RCC-8 has been used in autonomous unmanned aerial vehicles (or, simply, drones), and in [23] a discrete domain counterpart of RCC-8 (called *discrete mereotopology*) has been used in medical image processing. Other typical applications involve mobile robot navigation [24], natural language processing [25], and computer vision [26].

Current QSTR tools are not suitable for dynamic settings

With regard to algorithms for tackling fundamental reasoning tasks, current work revolves around methods that are primarily static in nature, i.e., they operate on fixed input data [27]–[32]; see also [33] and cited works therein. With respect to decomposability and tractability properties that can be exploited to boost the reasoning efficiency and enable parallelization, to the best of our knowledge there is no *practical* published research aside from the works utilizing tree decompositions for qualitative constraint networks [31], [34] (notions such as k-consistency can be leveraged theoretically [35], but are hardly practical or suitable for applications)—see also [36] and cited works therein. However, this approach is almost fully graph-based and, hence, does not adequately consider the semantics of the relations of a network during the decomposition phase. In relation to this research direction, Sioutis and Janhunen recently suggested to identify and make use of certain structural features in qualitative constraint networks (viz., *backdoors* and *backbones*) [37], as a means to define collaborative frameworks involving SAT, CP, and native reasoning tools and inspire novel decomposition and parallelization techniques; as a result, adaptive constraint propagators having a better insight than the state of the art into the particulars of real-world datasets could be developed. In addition, the authors recently led a preliminary work on a novel proactive approach that introduced a notion of *resistance* (called *robustness* in that work) to the possible future alterations of a spatio-temporal configuration [38], which is an important concept in the context of fluctuant and dynamic situations, such as real-life settings. This work allows for establishing frameworks that will be most likely to withstand perturbation (robustness) and that can be easily repaired when they did not suceed in doing so (stability/flexibility). Notably, this work also gave rise to a recent approach where certain spatio-temporal relations are preferred to others when trying to refine a spatio-temporal configuration, e.g., refining the relation between Tasks A and B to *precedes* in the network of Figure 1, depending on the count of local satisfiable atomic configurations that the relations are involved with [39].

Contribution

In this paper, we build upon the aforementioned preliminary work about robustness, and present a roadmap on how it can be

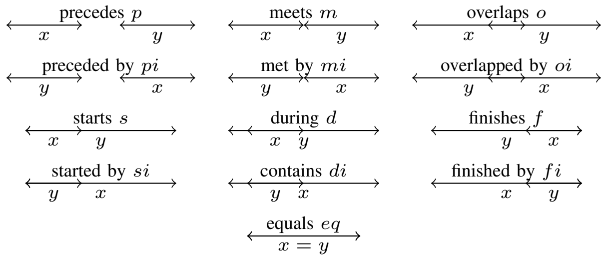

Fig. 2: A visualization of the 13 base relations b of IA, where each base relation relates two potential intervals x and y as in $x\ b\ y$; here, bi denotes the converse of b (viz., b^{-1} formally)

further extended with probabilities, and aid in the development of hybrid AI systems as a consequence, i.e., systems that can combine the strengths of machine learning and symbolic AI techniques.

II. Preliminaries

A qualitative spatial or temporal constraint language is based on a finite set B of binary relations, called *base relations*, that are *jointly exhaustive and pairwise disjoint* and defined over an infinite domain D (e.g., the real line or some topological space) [1]. Given a qualitative calculus over a domain D, each one of its base relations, alone, represents the definite knowledge between any two of the entities it may constrain. A union of base relations specifies indefinite knowledge, and such knowledge is represented by the set containing the involved base relations; thus, 2^{B} represents the total set of possible relations. Further, the set B is closed under the *converse* operation ($^{-1}$) and contains the identity relation Id, and its powerset, viz., 2^{B}, employs the union and intersection set-theoretic operations, the special *weak composition* operation ($\diamond$) [1], and the converse operation. For all $r \in 2^{\mathsf{B}}$, we have that $r^{-1} = \bigcup\{b^{-1} \mid b \in r\}$. The weak composition of two base relations $b, b' \in \mathsf{B}$ is characterized as the most restrictive, i.e., smallest, relation $r \in 2^{\mathsf{B}}$ that includes $b \circ b'$, i.e., $b \diamond b'=\{b'' \in \mathsf{B} \mid (b \circ b') \cap b'' \neq \emptyset\}$, 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 usual composition of b and b'. Finally, for all $r, r' \in 2^{\mathsf{B}}$, we have that $r \diamond r' = \bigcup\{b \diamond b' \mid b \in r, b' \in r'\}$.

As an illustration, let us revisit the Interval Algebra (IA) qualitative temporal constraint language, proposed by Allen [12]. Temporal entities in IA are intervals on the real line, and its set of base relations $\mathsf{B} = \{eq, p, m, o, s, d, f, pi, mi, oi, si, di, fi\}$ encodes knowledge about the temporal relations between such intervals, as described in Figure 2.

The challenge of representing and reasoning about qualitative information can be facilitated by a *qualitative constraint network* (QCN), for which we recall the following definition:

Definition 1. *A* QCN *is a tuple* (V, C) *where:*

- $V = \{v_1, \ldots, v_n\}$ *is a finite set of variables over some infinite domain* D *(e.g., space- or time-based);*
- *and* C *is a mapping* $C : V \times V \rightarrow 2^{\mathsf{B}}$ *associating a relation (set of base relations) with each pair of variables.*

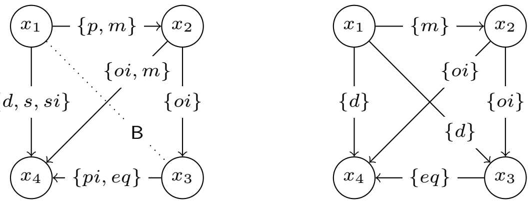

(a) A QCN $\mathcal{N} = (V, C)$ (b) A scenario $\mathcal{S}$ of $\mathcal{N}$

Fig. 3: Illustrative examples of QCN terminology using IA; here, $C(v, v) = \{\mathsf{Id}\}$, and $(C(v', v))^{-1} = C(v, v')\ \forall v, v' \in V$

An example of a QCN appears in Figure 3a; loops or converse relations do not appear in the figure for clarity.

Definition 2. *Let $\mathcal{N} = (V, C)$ be a* QCN*, then:*

- *a* solution *of $\mathcal{N}$ is a mapping $\sigma : V \to \mathsf{D}$ such that $\forall (v, v') \in V \times V$, $\exists b \in C(v, v')$ such that $(\sigma(v), \sigma(v')) \in b$; and $\mathcal{N}$ is* satisfiable *iff it yields a solution;*
- *a* refinement *$\mathcal{N}'$ of $\mathcal{N}$ is a* QCN *(V, C') such that $C'(v, v') \subseteq C(v, v')\ \forall v, v' \in V$;*
- *$\mathcal{N}$ is* atomic *iff $\forall v, v' \in V$, $|C(v, v')| = 1$;*
- *a* scenario *$\mathcal{S}$ of $\mathcal{N}$ is a satisfiable atomic refinement of $\mathcal{N}$, should $\mathcal{N}$ be satisfiable (see Figure 3b);*
- *the* size *of $\mathcal{N}$, denoted by $|\mathcal{N}|$, is $|\{(v, v') \mid v, v' \in V\}|$.*

For convenience, given a QCN $\mathcal{N} = (V, C)$, we can use $N[v, v']$ to represent relation $C(v, v')$.

III. Robustness, stability, and flexibility properties in QSTR

As noted in Section I, the authors recently led a preliminary work on a novel proactive approach that introduces a notion of *resistance* (called *robustness* in that work) to the possible future alterations of a spatio-temporal configuration, and in particular a spatio-temporal scenario [38]. Such a concept of resistance can be important in the context of fluctuant and dynamic situations, such as real-life settings. Let us revisited the simplified QCN of Figure 1. Given such a QCN, we are looking into the problem of obtaining a satisfiable scenario of it that is more resistant to perturbation than any other scenario, or, equivalently, a satisfiable scenario that is more likely than any other scenario to retain its validity (satisfiability) after the perturbation occurs; we call such a scenario a *robust scenario*. Here, choosing Task A to *precede* Task B forces us to have Task A *precede* Task C too. However, if we choose Task A to *follow* Task B instead, then we keep all the possibilities between Tasks A and C available (e.g., *precedes*, *overlaps*, *during*, or *follows*). Consequently, regardless of the change that may (or may not) occur in the relation between Tasks A and C, placing Task A after Task B yields a satisfiable scenario that will be able to maintain its satisfiability against this possible pertubation and, hence, that is more robust compared to the other ones in this case. In sum, a robust scenario can be considered as one that is best in terms of perturbation tolerance, and as a proactive measure limiting to the extent possible any need for successive repairs.

A. Theoretical study of robustness and its related properties

Robustness is a concept in problem solving that has been a topic of research arguably ever since the first various techniques for problem solving started to appear. Indeed, as soon as there were different solutions to some problem, it became important to have a way to grade/rank those solutions based on some (often robustness-related) criteria; see for example [40], and [41] and cited works therein. Robustness has been studied quite thoroughly in traditional constraint programming, where also other related notios such as *stability* and *flexibility* were introduced [42]–[44]. These related notions are well worth exploring in qualitative spatial and temporal reasoning too. In brief, a stable solution is one that can be repaired with a minimum number of corrections in case a perturbation invalidates it, and a flexible solution is a more relaxed version of a stable one, where the number of repairs need not be minimal but under some fixed (subjectively-defined) threshold. We remind the reader again that, in contrast to a stable or a flexible solution, a robust solution is simply more likely to remain valid after a perturbation occurs; the difference is fine, but it should be clear nonetheless. In the context of CSPs with infinite domains, viz., QCNs, and in particular in the introductory work of [38], a notion of robustness for QCNs was defined and it was left as an *open question* whether the problem of verifying a robust scenario is in NP (given of course that the satisfiability problem in the considered calculus is NP-complete). As a consequence, no interesting tractable cases were identified (should they exist). Furthermore, that notion of robustness assumed an equal probability among all possible perturbations, which may not be realistic enough in real-life settings, and the related properties of stability and flexibility were not introduced and studied at all. Although we will draw inspiration from the aforementioned works in traditional constraint programming to obtain theoretical results and generalized frameworks for robustness and its related properties in the context of QSTR, it should be mentioned that the methods that exist for CSPs are not readily applicable to QCNs, as the latter involve infinite domains and have a distinct set of exploitable properties [11]. (In support of the aforementioned statement, see also [45] for an overview and evaluation of different paradigms for solving QCNs.)

B. Practical tools for obtaining robust scenarios

In the work of [38] an implementation of a novel algorithm was proposed that, although very useful in assessing the differences that appear in the scenario-space of a given QCN in a proof-of-concept setting, is far from being scalable, as it can only handle up to few tens of variables. Therefore, here we argue for implementing tools that will exploit the theoretical results of the previous study in order to achieve scalability. Furthermore, and in line with the proposed theoretical study, we argue that these tools should incorporate different probabilities for different kinds of perturbation, and should be able to provide both exact and approximate solutions to a given problem (depending on the application, an approximate solution—which can be computed much faster than an exact one—might already be sufficient). Regarding the modeling of perturbations themselves, we suggest looking into the concept of *generalized neighbourhood graphs* [46], which allow certain perturbations to be modeled in a more dynamic setting; of course, an extension with probabilities will be required here too.

C. Robust QSTR as the backbone of hybrid AI systems

In this section we describe how robust QSTR can be the backbone of hybrid AI systems, presenting some examples in the context of abductive learning and AI planning specifically.

Abductive learning: In [9] the authors make the case for neurally-enhanced QSTR, by proposing to use probabilistic values to construct a *bidirectional feedback loop* between the machine learning model and the symbolic framework, referred to in the recent literature as *abductive learning* [47, Figure 1]; see also [48]. Specifically, they propose to annotate the variables in a QCN with the probabilistic values produced by the ML model, and assign probabilistic values to the constraints of the QCN too. A neurally-enhanced symbolic relation could then be as follows ($\%$s denote likelihood):

$$X^{(90\%\text{ yolk})} \text{ is } \textit{inside}^{(40\%\text{ true})} \text{ or } \textit{overlaps}^{(60\%\text{ true})}\ Y^{(95\%\text{ egg})}.$$

In a sense, logic is used to compose (partial) knowledge learned by machine learning methods, and this learned knowledge is allowed to influence that logical composition too. To this end, it is important to establish probabilistic robustness measures relating to QSTR frameworks, and also construct adaptive and dynamic algorithms for verifying neural network-based components at runtime. We take the next step here and propose such a measure as follows:

Definition 3. *(Robustness Measure) Given a* QCN $\mathcal{N} = (V, C)$*, an atomic refinement* $\mathcal{N}'$ *of* $\mathcal{N}$*, and the probability* $p_{ij}(\mathcal{N}'[v_i, v_j])$ *of a relation* $\mathcal{N}'[v_i, v_j]$ *appearing in any scenario of* $\mathcal{N}$*, the* robustness measure *of* $\mathcal{N}'$*, denoted by* $\mathsf{robustness}(\mathcal{N}')$*, is defined to be:*

$$\mathsf{robustness}(\mathcal{N}') = \frac{\sum\limits_{v_i, v_j \in V} p_{ij}(\mathcal{N}'[v_i, v_j])}{|\mathcal{N}'|}.$$

For example, given an atomic refinement of some QCN with a robustness measure of 0.9, this would mean that on average each of its relations would have 90% chance to appear in a given scenario of that QCN, out of the set of all scenarios of that QCN. In that sense, it can be also stated that the aforementioned atomic refinement is around 90% compatible with the scenario space of the QCN. As noted earlier, the probability $p_{ij}(\mathcal{N}'[v_i, v_j])$ can be approximated by machine learning techniques; hence, we think that this scheme can adequately support abductive learning once implemented. Clearly, based on the definition above, the most robust atomic refinement (where the probabilities are based on the output of a classifier for example) might not even be satisfiable. Therefore, we also argue for developing probability-based methods towards maximizing satisfiability (or, in other words, minimizing inconsistency), either via belief revision or traditional constraint propagation frameworks.

AI planning: To further detail how robustness and dynamic reasoning can play a role in hybdid AI systems, let us consider the following simplified rule in the context of a system that involves AI planning, and in particular train schedulling:

$$\begin{gathered}(\text{Train } X\ \{\textit{uses}\}\ (\text{Track } A \oplus \text{Track } B))\wedge\\ (\text{Train } Y\ \{\textit{uses}\}\ \text{Track } B)\wedge\\ (\text{departInt}(\text{Train } X)\ \{\text{p,m,o}\}\ \text{departInt}(\text{Train } Y)).\end{gathered}$$

It is specified that Train X, which uses either Track A or B, has a departure slot that *may* overlap (o) that of Train Y, which uses Track B; it is assumed here that all outcomes are equally probable. Naturally, some safety properties must hold too, like the following one:

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

From a resource allocation perspective, it may be preferred to only use Track B, and from a collision avoidance perspective, it may be preferred to use both tracks by allocating Track A to Train X. In any case, a plan can be devised based on these preferences. However, one day new information becomes available via an ML system that predicts that Track A will be subject to buckling due to a heatwave [49]:

$$\Diamond\text{blocked}(\text{Track } A).$$

Clearly, we now need to take this new information into account and revise our perceptions of robustness pertaining to resource allocation and collision avoidance, as well as update our original configuration. Although we have used here a simple example involving linear temporal logic for demonstrative purposes, more expressive temporal logics can be considered, like metric temporal logic (see for example [50]), and more complex plans can be envisioned too. A related discussion on robust planning appears in [51].

IV. Conclusion

In this short challenge paper, we argued for the need to have *robust* Qualitative Spatio-Temporal Reasoning (QSTR), as the backbone of hybrid AI systems involving also computations with spatio-temporal information. QSTR is a well-established Symbolic AI framework for representing and reasoning about spatial and temporal information via the use of disjunctive natural relations, e.g., "Region X is *inside* or *overlaps* Region Y". Here, *robustness* entails a notion of resistance to the possible future alterations of a spatio-temporal configuration. In this paper, we propose a formulation of robustness based on probabilities, that should allow a spatio-temporal configuration to also be pliable and adapt to enforced changes, and gave some examples of how robust QSTR can be the backbone of hybrid AI systems, emphasizing on abductive learning and AI planning specifically. On the one hand, our paper is meant to pose some challenges to the scientific community, and the QSTR community in particular, that would allow QSTR to be better integrated into hybrid AI systems; on the other hand, it is also meant to inspire discussions and feedback.

ACKNOWLEDGMENT

We would like to thank Prof. Dr. Christoph Schlieder from Bamberg University, Germany, and Dr. Zhiguo Long from Southwest Jiaotong University, China, for their availability in discussing our ideas with them, and for providing contructive feedback as a result of these discussions.

REFERENCES

[1] G. Ligozat, *Qualitative Spatial and Temporal Reasoning*, ser. ISTE Series. Wiley, 2011.

[2] M. Bhatt and J. O. Wallgrün, "Geospatial Narratives and Their Spatio-Temporal Dynamics: Commonsense Reasoning for High-Level Analyses in Geographic Information Systems," *ISPRS Int. J. Geo-Information*, vol. 3, pp. 166–205, 2014.

[3] F. Dylla and J. O. Wallgrün, "Qualitative Spatial Reasoning with Conceptual Neighborhoods for Agent Control," *J. Intell. Robotic Syst.*, vol. 48, pp. 55–78, 2007.

[4] N. Krishnaswamy, S. Friedman, and J. Pustejovsky, "Combining Deep Learning and Qualitative Spatial Reasoning to Learn Complex Structures from Sparse Examples with Noise," in *AAAI*, 2019.

[5] K. S. R. Dubba, A. G. Cohn, D. C. Hogg, M. Bhatt, and F. Dylla, "Learning Relational Event Models from Video," *J. Artif. Intell. Res.*, vol. 53, pp. 41–90, 2015.

[6] J. Suchan, M. Bhatt, and S. Varadarajan, "Out of Sight But Not Out of Mind: An Answer Set Programming Based Online Abduction Framework for Visual Sensemaking in Autonomous Driving," in *IJCAI*, 2019.

[7] R. Moskovitch and Y. Shahar, "Classification of multivariate time series via temporal abstraction and time intervals mining," *Knowl. Inf. Syst.*, vol. 45, pp. 35–74, 2015.

[8] T. P. D. Homem, P. E. Santos, A. H. R. Costa, R. A. da Costa Bianchi, and R. L. de Mántaras, "Qualitative case-based reasoning and learning," *Artif. Intell.*, vol. 283, p. 103258, 2020.

[9] M. Sioutis and D. Wolter, "Qualitative Spatial and Temporal Reasoning: Current Status and Future Challenges," in *IJCAI*, 2021.

[10] M. Bhatt, H. Guesgen, S. Wölfl, and S. Hazarika, "Qualitative Spatial and Temporal Reasoning: Emerging Applications, Trends, and Directions," *Spatial Cognition & Computation*, vol. 11, pp. 1–14, 2011.

[11] F. Dylla, J. H. Lee, T. Mossakowski, T. Schneider, A. van Delden, J. van de Ven, and D. Wolter, "A Survey of Qualitative Spatial and Temporal Calculi: Algebraic and Computational Properties," *ACM Comput. Surv.*, vol. 50, pp. 7:1–7:39, 2017.

[12] J. F. Allen, "Maintaining Knowledge about Temporal Intervals," *Commun. ACM*, vol. 26, pp. 832–843, 1983.

[13] L. Mudrová and N. Hawes, "Task scheduling for mobile robots using interval algebra," in *ICRA*, 2015.

[14] R. T. Snodgrass, "The Temporal Query Language TQuel," *ACM Trans. Database Syst.*, vol. 12, pp. 247–298, 1987.

[15] C. X. Chen and C. Zaniolo, "Universal Temporal Data Languages," in *DDLP*, 1998.

[16] P. Denis and P. Muller, "Predicting Globally-Coherent Temporal Structures from Texts via Endpoint Inference and Graph Decomposition," in *IJCAI*, 2011.

[17] M. C. Golumbic and R. Shamir, "Complexity and Algorithms for Reasoning about Time: A Graph-Theoretic Approach," *J. ACM*, vol. 40, pp. 1108–1133, 1993.

[18] R. Lu, S. W. Sadiq, V. Padmanabhan, and G. Governatori, "Using a temporal constraint network for business process execution," in *ADC*, 2006.

[19] Z. Ibrahim, H. Wu, and R. Dobson, "Modeling Rare Interactions in Time Series Data Through Qualitative Change: Application to Outcome Prediction in Intensive Care Units," in *ECAI*, 2020.

[20] D. A. Randell, Z. Cui, and A. Cohn, "A Spatial Logic Based on Regions and Connection," in *KR*, 1992.

[21] A. D. Lattner, I. J. Timm, M. Lorenz, and O. Herzog, "Knowledge-based risk assessment for intelligent vehicles," in *KIMAS*, 2005.

[22] F. Heintz and D. de Leng, "Spatio-Temporal Stream Reasoning with Incomplete Spatial Information," in *ECAI*, 2014.

[23] D. A. Randell, A. Galton, S. Fouad, H. Mehanna, and G. Landini, "Mereotopological Correction of Segmentation Errors in Histological Imaging," *J. Imaging*, vol. 3, p. 63, 2017.

[24] V. Fenelon, P. E. Santos, H. M. Dee, and F. G. Cozman, "Reasoning about shadows in a mobile robot environment," *Appl. Intell.*, vol. 38, pp. 553–565, 2013.

[25] P. Kordjamshidi and M. Moens, "Global machine learning for spatial ontology population," *J. Web Sem.*, vol. 30, pp. 3–21, 2015.

[26] M. Sridhar, A. G. Cohn, and D. C. Hogg, "From Video to RCC8: Exploiting a Distance Based Semantics to Stabilise the Interpretation of Mereotopological Relations," in *COSIT*, 2011.

[27] N. Amaneddine, J. Condotta, and M. Sioutis, "Efficient Approach to Solve the Minimal Labeling Problem of Temporal and Spatial Qualitative Constraints," in *IJCAI*, 2013.

[28] M. Sioutis, S. Li, and J. Condotta, "Efficiently Characterizing Non-Redundant Constraints in Large Real World Qualitative Spatial Networks," in *IJCAI*, 2015.

[29] Z. Long, M. Sioutis, and S. Li, "Efficient Path Consistency Algorithm for Large Qualitative Constraint Networks," in *IJCAI*, 2016.

[30] M. Sioutis and J. Condotta, "Efficiently Enforcing Path Consistency on Qualitative Constraint Networks by Use of Abstraction," in *IJCAI*, 2017.

[31] M. Sioutis, Z. Long, and S. Li, "Leveraging Variable Elimination for Efficiently Reasoning about Qualitative Constraints," *Int. J. Artif. Intell. Tools*, vol. 27, p. 1860001, 2018.

[32] M. Sioutis, A. Paparrizou, and J. Condotta, "Collective singleton-based consistency for qualitative constraint networks: Theory and practice," *Theor. Comput. Sci.*, vol. 797, pp. 17–41, 2019.

[33] M. Sioutis, "Just-In-Time Constraint-Based Inference for Qualitative Spatial and Temporal Reasoning," *Künstliche Intell.*, vol. 34, pp. 259–270, 2020, discussion.

[34] J. Huang, J. J. Li, and J. Renz, "Decomposition and tractability in qualitative spatial and temporal reasoning," *Artif. Intell.*, vol. 195, pp. 140–164, 2013.

[35] M. Bodirsky and S. Wölfl, "RCC8 is polynomial on networks of bounded treewidth," in *IJCAI*, 2011.

[36] M. Sioutis, Y. Salhi, and J. Condotta, "Studying the use and effect of graph decomposition in qualitative spatial and temporal reasoning," *Knowl. Eng. Rev.*, vol. 32, p. e4, 2017.

[37] M. Sioutis and T. Janhunen, "Towards Leveraging Backdoors in Qualitative Constraint Networks," in *KI*, 2019.

[38] M. Sioutis, Z. Long, and T. Janhunen, "On Robustness in Qualitative Constraint Networks," in *IJCAI*, 2020.

[39] M. Sioutis and D. Wolter, "Dynamic Branching in Qualitative Constraint Networks via Counting Local Models," in *TIME*, 2020.

[40] M. L. Ginsberg, A. J. Parkes, and A. Roy, "Supermodels and Robustness," in *AAAI/IAAI*, 1998.

[41] G. Verfaillie and N. Jussien, "Constraint Solving in Uncertain and Dynamic Environments: A Survey," *Constraints*, vol. 10, pp. 253–281, 2005.

[42] L. Climent, "Robustness and stability in dynamic constraint satisfaction problems," *Constraints*, vol. 20, pp. 502–503, 2015.

[43] L. Climent, R. J. Wallace, M. A. Salido, and F. Barber, "Robustness and Stability in Constraint Programming under Dynamism and Uncertainty," *J. Artif. Intell. Res.*, vol. 49, pp. 49–78, 2014.

[44] F. Barber and M. A. Salido, "Robustness, stability, recoverability, and reliability in constraint satisfaction problems," *Knowl. Inf. Syst.*, vol. 44, pp. 719–734, 2015.

[45] M. Westphal and S. Wölfl, "Qualitative CSP, Finite CSP, and SAT: Comparing Methods for Qualitative Constraint-based Reasoning," in *IJCAI*, 2009.

[46] M. Ragni and S. Wölfl, "Temporalizing Spatial Calculi: On Generalized Neighborhood Graphs," in *KI*, 2005.

[47] Z. Zhou, "Abductive learning: towards bridging machine learning and logical reasoning," *Sci. China Inf. Sci.*, vol. 62, pp. 76 101:1–76 101:3, 2019.

[48] W. Dai, Q. Xu, Y. Yu, and Z. Zhou, "Bridging Machine Learning and Logical Reasoning by Abductive Learning," in *NeurIPS*, 2019.

[49] M. N. Nguyen, X. Wang, and C.-H. Wang, "A reliability assessment of railway track buckling during an extreme heatwave," *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, vol. 226, pp. 513–517, 2012.

[50] P. Doherty, J. Kvarnström, and F. Heintz, "A temporal logic-based planning and execution monitoring framework for unmanned aircraft systems," *Auton. Agents Multi Agent Syst.*, vol. 19, pp. 332–377, 2009.

[51] T. Schwartz, M. Sioutis, and D. Wolter, "Towards Robust Constraint Satisfaction in Hybrid Hierarchical Planning," in *ICAPS Workshop on Hierarchical Planning*, 2021.