Including Qualitative Spatial Knowledge in the Sense-Plan-Act Loop

Masoumeh Mansouri (PhD student) and Federico Pecora (supervisor)

Center for Applied Autonomous Sensor Systems, Örebro University, SE-70182 Sweden
{mmi, fpa}@aass.oru.se

Abstract. In this paper we present ongoing work on integrating qualitative and metric spatial reasoning into planning for robots. We propose a knowledge representation and reasoning technique, grounded on well-established constraint-based spatial calculi, for combining qualitative and metric knowledge and obtaining plans expressed in actionable metric terms.

Keywords: Qualitative spatial constraints, Combining AI and Robotics

1 Introduction

When we plan to achieve activities we rely on several, different abstractions of the world around us. These abstractions, many of which are learned through experience, are often qualitative: we know that knives should be put to the right of dishes and forks to the left when setting a table. When performing the actions to set the table, this qualitative knowledge is used to instantiate precise placing actions in metric space.

Whereas bridging this gap is natural for humans, it is not at all evident how to endow robots with this capability. A robot's world is entirely metric: it perceives events and carries out actions in metric time, it can localize, displace itself and perceive objects in a reference frame. Yet specifying sophisticated robot behavior in purely metric terms is difficult, as the specification would have to be long and overly specialized to the particular setting in which the robot operates. Conversely, qualitative constraint-based representations facilitate modeling by humans — although they often fail to capture the details that are necessary for proper execution. Integrating temporal, spatial, and other reasoning capabilities in the sense-plan-act loop is an important step towards building general purpose robots (see Related Work).

The present work focuses on spatial knowledge. We attach metric semantics to qualitative spatial constraints and prove formal properties of the obtained calculus. These properties enable reasoning in three important phases of the sense-plan-act loop, namely (1) matching perceived context with general knowledge about the environment, (2) instantiating plans into the metric space of the real world, and (3) detecting and reacting to contingencies. In all three processes, different levels of abstraction are used: perception generates metric spatial knowledge, while general knowledge about the environment can be qualitative in nature; similarly, plans may call for the achievement of qualitative spatial relations, but actions must be precisely instantiated in metric space.

2 Related Work

Bridging the gap between the robot's metric world and its symbolic knowledge is studied in automated planning.However, the focus is primarily on metric time and resources [1, 2]. Work on combining qualitative spatial knowledge with perception, planning and actuation is sparse. Loutfi et al. [3] look at this problem in the context of perceptual anchoring to provide qualitative relations inferred from observed metric relations (partially addressing point (1) above). Work in Cognitive Vision addresses this issue [4] by focusing on scene understanding. Work on structural pattern recognition [5] provides techniques for matching qualitative spatial knowledge (representing a specified structure) to perceived context. In all of the above, qualitative relations do not belong to a well-defined calculus, which would facilitate logical reasoning, rather they are tailored to capture specific features which are useful for pattern specification and recognition in the particular application.

In robotics, many focus on geometric reasoning. Guitton and Farges [6] employ metric constraints in combination with planning, thus partially addressing point (2), while Jang et al. [7] propose ad-hoc metric spatial reasoning for analyzing perceived context. Others leverage the richness of qualitative spatial calculi predominantly for representation rather than reasoning (e.g., robot navigation and self-localization [8], motion planning [9]).

In Robotics, qualitative knowledge can be modeled through domain-specific predicates which allow to perform metric spatial reasoning though procedural attachment (point (2) above, see [10, 11]). To the best of our knowledge, no work employs well-founded qualitative spatial calculi in conjunction with metric spatial reasoning.

3 Combining Qualitative and Metric Relations

Most existing qualitative calculi represent spatial relations as constraints. Each of these constraint-based calculi focuses on one or more categories of spatial concepts— e.g., Region Connection Calculus (RCC) [12], Cardinal Direction Calculus (CDC) [13, 14] and Rectangle Algebra (RA) [15]. In this work, we chose RA for the specific properties of this calculus that is combining topology and orientation.

Inspired by the work on temporal reasoning [16], we augment qualitative RA relations by attaching specified bounds to them, e.g., $A\langle \text{Before}[5, 13], \text{Before}[0, +\infty)\rangle B$, where A and B are axis-parallel rectangles, and $\langle \text{Before}, \text{Before}\rangle$ is an atomic qualitative RA relation. Similarly to augmented Allen Interval Algebra, specified bounds maintain the semantics of the qualitative relations unaltered. Since qualitative relations subsume metric relations [17], we employ simple distance constraints to represent the metric semantics of qualitative relations, as well as the metric bounds of these relations. This allows to process both qualitative constraints and their metric bounds uniformly at the metric level, i.e., reasoning at the metric level computes the consequence of both qualitative and metric relations among rectangles. The result of reasoning, which is tractable and complete, is a set of admissible bounds on the placement of rectangles, a metric solution which is directly understandable by the robot.

A similar work addressing the problem of combining qualitative and metric spatial constraints is proposed by [18]. This approach relies on several constraint networks: one is a RA network concerned only with qualitative relations; and two networks of simple distance constraints are used to represent metric bounds among the extreme points of Allen's intervals along the first and second axes. Qualitative and metric relations are thus de-coupled, and their mutual consistency must be verified through both qualitative and metric inference. On the basis of the achievement of [18], it is straightforward to prove that the consistency of a convex rectangle constraint network[1] augmented by specified bounds can be decided by the consistency of two underlying simple distance constraint networks corresponding to the projection of augmented RA relations onto the first and second axes. These results will be shown in a forthcoming paper.

4 Modeling Perceived Context with RA

In order to model perceived context, we introduce the unary constraint *Size*, which bounds the distances between two points of the same rectangle along one axis. Augmented RA constraints together with *Size* constraints express a robot's knowledge about orientation and topology of spatial entities. Furthermore, we wish to model bounds on the placement in the scene of bounding boxes. We thus define the constraint *At* which bounds the absolute placement of bounding boxes. We can prove that the consistency of a convex RA constraint network with specified bounds enriched with $Size$ and At constraints is decided by the path consistency of two underlying simple constraint networks. In the context of this work, the computational core of reasoning is planning, which can be PSPACE-hard. Therefore, the polynomial complexity of path consistency enforcement does not contribute significantly to the overall complexity of reasoning.

Overall, a RA network (enriched with *Size* and *At* constraints) can be used to represent uniformly both a desired spatial layout of objects and the observed spatial layout. We reduce the problem of matching observed spatial relations and perceived context (1) to consistency checking as follows: a set of rectangles is created to model the observed objects, each of which is constrained with an *At* constraint reflecting its position; all knowledge is encoded as further rectangles and qualitative spatial relations (with optional specified bounds); a binary RA constraint $\{\mathrm{eq}, \mathrm{eq}\}$ is added to unify each observed object with its counterpart in the model (see Figure 1). If this network is consistent, then the observed state of the world adheres to the the robot's spatial knowledge.

5 Solution extraction

Maintaining a network representing both qualitative and metric relations enables to do more than matching perceived context to knowledge. As a constraint network forms a query of our desire, we can use constraint networks to answer queries for instantiating planned actions. For instance, we can construct a constraint network containing *At* relations for all perceived objects, as well as variables for object(s) that are not in the scene and need to be placed. A solution to this spatial constraint satisfaction problem (i.e.,

[1] A convex RA network contains only convex relations. Convex relations are such that they can be "translated" into simple temporal distance constraints [19].

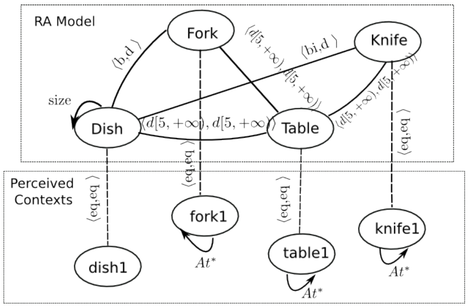

Fig. 1. Unification of perceived context with the RA model.

a substitution of coordinates to points that is consistent) represents a placement of the objects in the scene that is consistent with qualitative and metric knowledge. In order to support this capability (2), we require a way to extract and discern between different solutions in a way that is appropriate for the robotic domain, as shown below.

Enforcing path consistency on the networks of simple distance constraints updates the bounds of all points in the network. Hence, the assignment of all lower bounds or of all upper bounds after path consistency enforcement are both valid solutions [20]. Other possible solutions can be obtained through incremental propagation.

In a robotic context, assignments other than the lower and upper bound solutions are preferable. Specifically, we are interested in obtaining the solution that has maximum distance from these two solutions, as the region that is given to a robot to place an object should tolerate the inaccuracy of manipulation. In other words, if the robot does not place an object exactly within the region, the spatial layout should still be consistent. For this reason, we prefer assignments that are close to the center of the solution space. Obtaining the exactly centered solution is an optimization problem that is too computationally demanding to solve on-line, therefore we sacrifice the optimality of the solution in favor of efficiency. In order to do so, we compute an approximation of the most centered solution for each rectangle by leveraging the concept of 2D representation of an interval introduced by [21].

6 Culprit Detection, Recovery Recommendation

We now turn our attention to spatial reasoning for detecting and reacting to contingencies (3). Suppose a robot has to place a dish on a table in which several other dishware have already been placed, and the robot is given a set of spatial relations which describe a well-set table (e.g., that a dish is to be placed between fork and knife). As described above, the robot instantiates the placing action by computing the solution to the query constraint network. However, the network turns out to be inconsistent, indicating that it cannot achieve the goal of placing a dish.

It may be possible to recover from the failure by finding the source(s) of inconsistency in the RA network. The source(s) of inconsistency are constraints, specifically *At* constraints deriving from observation (e.g., the fork and knife are so close that a dish cannot fit between them). Eliminating the inconsistency thus means to relax one or more *At* constraints. The power set of *At* constraints in the network contains all possible *culprit sets*. If there exists at least one culprit set such that its removal from the network makes the network consistent, then a solution of this network contains new positions for the objects related to constraints in the culprit set. These new positions are now consistent with the original placing action(s), and can be used to instantiate new actions that bring about this consistent situation. Finding a desirable culprit set is a search process. The search employs a heuristics which favors small sets and moves which least affect the *spatial rigidity* of the network. The result of the search is a set of new coordinates (i.e., new *At* constraints) for rectangle(s) involved in the culprit set. The rationale behind choosing small sets is that moving fewer objects is less prone to failure than moving many. To define spatial rigidity, we employ a measure known as the root mean square (RMS) rigidity of a network of simple distance constraints [22]. Low rigidity entails that the admissible bounds of objects are such that there is significant slack in determining new placements, whereas high rigidity means that the constraints in the network afford placement options which are close to failure, and therefore manipulation must be more precise. The complexity of culprit detection is exponential in the number of *At* constraints. In practice, the exponent is often small as it corresponds to the number of objects in the scene.

7 Ongoing and Future Work

We have presented a knowledge representation formalism for integrated metric and qualitative spatial reasoning. By uniformly representing metric bounds alongside qualitative relations, the calculus can be used both to specify spatial knowledge in human-accessible terms and to represent perceived context. Its formal properties can be leveraged to realize three important reasoning tasks in the sense-plan-act loop.

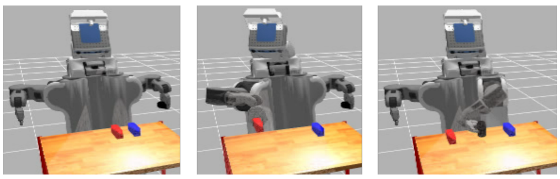

(a) (b) (c)

Fig. 2. A simulated PR2 robot employing spatial knowledge to set a table.

In ongoing work, we are implementing the proposed approach on a PR2 robot. In the experiment shown in Figure 2, the robot is given the knowledge that knives should be put to the right of dishes and forks to the left when setting a table. The robot finds an inconsistent layout when reaching the table (a), as the knife and fork are too close; it identifies the knife as a possible culprit and replaces it (b); finally, it places the cup (c).

As future work, we will investigate the consequence of adding specified bounds to the disjunction of RA relations. Furthermore, we will focus on integrating spatial reasoning with temporal, resource and causal reasoning in a constraint based framework for a robot platform.

References

1. Fratini, S., Pecora, F., Cesta, A.: Unifying Planning and Scheduling as Timelines in a Component-Based Perspective. Archives of Control Sciences **18**(2) (2008) 231–271
2. Ghallab, M., Laruelle, H.: Representation and Control in IxTeT, a Temporal Planner. In: Proc. of the 2nd Int'l Conf. on Artif. Intell. Planning and Scheduling. (1994)
3. Loutfi, A., Coradeschi, S., Daoutis, M., Melchert, J.: Using knowledge representation for perceptual anchoring in a robotic system. Journal on Artif. Intell. Tools **17**(5) (2008) 925–944
4. Hois, J., Wünstel, M., Bateman, J., Röfer, T.: Dialog-based 3d-image recognition using a domain ontology. In: Proc. of the Int'l Conf. on Spatial Cognition V: Reasoning, Action, Interaction. (2007)
5. Colliot, O., Camara, O., Bloch, I.: Integration of fuzzy spatial relations in deformable models-application to brain mri segmentation. Pattern Recogn. **39**(8) (2006) 1401–1414
6. Guitton, J., Farges, J.L.: Taking into account geometric constraints for task-oriented motion planning. In: Workshop on Bridging the Gap between Task and Motion Planning. (2009)
7. Jang, H.Y., Moradi, H., Hong, S., Lee, S., Han, J.: Spatial reasoning for real-time robotic manipulation. In: IEEE/RSJ Int'l Conference on Intelligent Robots and Systems. (oct. 2006) 2632 –2637
8. Wolter, D., Wallgrün, J.: Qualitative spatial reasoning for applications: New challenges and the sparq toolbox. In: Qualitative Spatio-Temporal Representation and Reasoning: Trends and Future Directions. IGI Global (2010)
9. Westphal, M., Dornhege, C., Wölfl, S., Gissler, M., Nebel, B.: Guiding the generation of manipulation plans by qualitative spatial reasoning. Spatial Cognition and Computation: An Interdisciplinary Journal **11**(1) (2011) 75–102
10. Mosenlechner, L., Beetz, M.: Parameterizing actions to have the appropriate effects. In: IEEE/RSJ Int'l Conference on Intelligent Robots and Systems. (2011)
11. Lagriffoul, F., Dimitrov, D., Saffiotti, A., Karlsson, L.: Constraint propagation on interval bounds for dealing with geometric backtracking. In: IEEE/RSJ Int'l Conference on Intelligent Robots and Systems. (2012)
12. Randell, D., Cui, Z., Cohn, A.: A spatial logic based on regions and connection. In: Proc. of Int'l Conf. on Principles of Knowledge Representation and Reasoning. (1992)
13. Goyal, R.: Similarity assessment for cardinal directions between extended spatial objects. PhD thesis, The University of Maine (2000)
14. Skiadopoulos, S., Koubarakis, M.: Composing cardinal direction relations. Artif. Intell. **152**(2) (2004) 143–171
15. Balbiani, P., Condotta, J.F., Del Cerro, L.F.: A new tractable subclass of the rectangle algebra. In: Proc. of Int'l Joint Conf. on Artif. Intell. Volume 1. (1999) 442–447
16. Meiri, I.: Combining qualitative and quantitative constraints in temporal reasoning. In: Artificial Intelligence. (1996) 260–267
17. van Beek, P.: Exact and approximate reasoning about qualitative temporal relations. PhD thesis, University of Waterloo (1990)
18. Condotta, J.F.: The augmented interval and rectangle networks. In: KR. (2000) 571–579
19. Ligozat, G.: A new proof of tractability for ORD-Horn relations. In: AAAI'96 Workshop on Spatial and Temporal Reasoning. (1996) 395–401
20. Dechter, R., Meiri, I., Pearl, J.: Temporal constraint networks. Artif. Intell. **49**(1-3) (1991) 61–95
21. Rit, J.F.: Propagating temporal constraints for scheduling. In: Nat. Conf. on Artif. Intell. (1986) 383–388
22. Hunsberger, L.: Algorithms for a temporal decoupling problem in multi-agent planning. In: Proc. of the Nat. Conf. on Artif. Intell. (2002)