TOWARD AN EXPERIMENTAL COGNITIVE ROBOTICS FRAMEWORK

Mehul Bhatt

SFB/TR 8 Spatial Cognition
Universität Bremen, Germany
bhatt@informatik.uni-bremen.de

ABSTRACT

We position our experimental framework for cognitive robotics that is aimed at integrating logic-based and cognitively-driven agent-control approaches, qualitative models of space and the ability to apply these in the form of planning, explanation and simulation in a wide-range of robotic-control platforms and simulation environments. In addition to its primary experimental function, the research proposed herein also has a utility toward pedagogical purposes. We present the overall vision of the project, and discuss ongoing work and present capabilities.

KEYWORDS: reasoning about actions and change, qualitative spatial reasoning, dynamic spatial systems, control, experimental robotics, simulation

1 INTRODUCTION

Research in the field of Reasoning about Actions and Change (RAC), also and increasingly being referred to as Cognitive Robotics, has considerably matured [Levesque and Lakemeyer, 2007]. Over the last decade, some of the theoretical work and the resulting formalisms for representing and reasoning about dynamic domains have evolved into practically applicable high-level agent control languages, the most prominent examples here being the situation calculus based GOLOG [Levesque et al., 1997] family of languages, e.g., CCGOLOG [Grosskreutz and Lakemeyer, 2000], CONGOLOG [De Giacomo et al., 2000], INDIGOLOG [Giacomo and Levesque, 1999], which is an incremental deterministic version of CONGOLOG, and the fluent calculus based language FLUX [Thielscher, 2005]. Differences in the theoretical underpinnings notwithstanding, a common feature of all these languages is the availability of imperative programming style constructs for the domain of robotics/agent-control, i.e., statement in the program correspond to actions, events and properties of the world in which an agent is operating. Parallel to the development in the area of reasoning about actions and change, the field of Qualitative Spatial Reasoning (QSR) has emerged as a sub-division in its own right within knowledge representation [Cohn and Hazarika, 2001]. Research in QSR has focused on the construction of formal methods (i.e., qualitative spatial calculi) for spatial modelling and reasoning. The scope of QSR, at least in so far as the context of qualitative spatial calculi is concerned, has been restricted to representational modes for spatial abstraction and reasoning. Major developments in this regard include: (a) the development of spatial calculi that are representative of distinct spatial domains, (b) constraint-based techniques for ensuring the global consistency of spatial information and (c) the application of conceptual-neighborhoods [Freksa, 1991] for dealing with continuous change and time. Similarly, there have also been considerable advances in the benchmarking of the computational aspects of the planning domain and de facto standardization of domain description languages in the form of the PLANNING DOMAIN DEFINITION LANGUAGE (PDDL) [McDermott et al., 1998] and related initiatives. Recent work even indicates a cross-over of results from the planning domain to the cognitive robotics area. For instance, the work by Claßen et al. [2007b,a] combines reasoning using the GOLOG language with modern PDDL planners by the embedding of state-of-art planning systems within the former. The main objective of this line of approach is that the power of modern efficient planners be exploited whilst preserving the overalll representational semantics of the situation calculus formalism that underlies the GOLOG language.

In this paper, we position our ongoing work toward the development of a framework for cognitive robotics that brings together logic-based and cognitively-driven agent-control approaches in an experimental manner. The proposed framework is designed to integrate diverse control calculi based on mathematical logic, qualitative models of space and the ability to apply these – in the form of spatial planning, explanation and simulation – for dynamic spatial modelling with a wide-range of robotic-control platforms and simulation environments. Indeed, the framework is driven by the need for a workbench that seamlessly brings together different control techniques, both logic-based and otherwise, and a generic (based) domain-description language, and qualitative spatial calculi under one unifying, experimental framework. The main objective here being that it should be possible for a domain-modeller to specify the physics of a particular domain once and exploit more than one control approach thereafter, without the need to dwell on the details of any of the available control approaches or qualitative spatial calculi. In addition, it is also required that the envisaged framework provide easy integration with existing low-level control apparatus such as robot control and simulation interfaces that

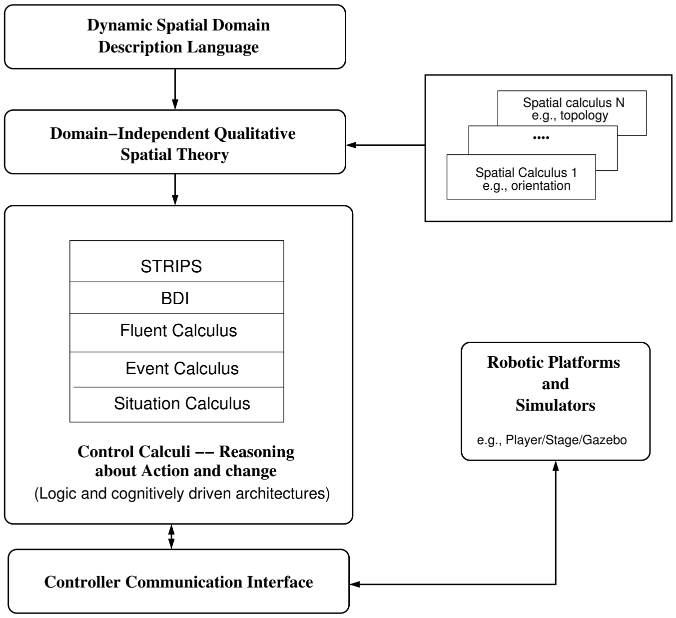

Figure 1: Overview of the Framework

exist in the open-source domain, a prime example of such an environment being the PLAYER/STAGE/GAZEBO project. We envisage that such a framework that supports easy experimentation with different control techniques, provides general modes of spatial information representation and reasoning, and additionally seamlessly integrates low-level control apparatus would, in addition to serving it primary experimental function, also be useful for pedagogical purposes at the tertiary-level.
Section 2 provides a brief overview of the proposed framework and section 3 positions the work in progress in terms of its implementation and other technical details. In section 4, we briefly discuss the immediate directions of the ongoing work.

2 Overview of the Framework

The primary aim of the framework is to provide a workbench of different control approaches that may be used '*independently*' for experiments in cognitive robotics. Toward this end, the framework consists of four key components. In (C1–C4) in the following, a conceptual overview of these components, as they are presently envisioned, and the main motivations thereof are presented with reference to the schematic in Fig. 1:

C1. Uniform Domain Description

A basic requirement within the framework is that it should be possible for modellers to specify the domain theory of their particular scenario, i.e., the underlying physics of the domain, using an uniform representation medium that is independent of the control apparatus that is being utilized. Such a medium requires a unified ontological view that transcends beyond any particular control calculus or non-logical control approach that is available within the framework. Among other things, what is required is that key ontological aspects pertaining to actions, events, effects, fluents and conditions need to be integrated in an ontology that may be utilized by the modeller of a domain. Such an ontology will facilitate the specification of a dynamic domain in a manner that is independent of the precise control mechanism (available within the framework) that is utilized as the basis of modelling and reasoning about change. Consequently, it is also implied that such a generic domain description is usable across all control (or reasoning) approaches that are available within the experimental framework. Although the issue of ontology construction is quite orthogonal to the issue of the precise language or mechanism to be used for instantiating it whilst modelling a domain, we consider the semantics of the basic versions of the PDDL language to be rich enough to cover most general scenarios. However, the use of PDDL as a means for uniform domain description within our framework is a topic of ongoing investigation.

C2. Multiple Control Approaches

As aforementioned, the primary aim of the framework is to provide a suite of different control approaches that may be used for representing and reasoning about dynamic environments. The suite of control approaches available with the framework also constitutes the most important (functional) component of the overalll experimental framework. It consists of a collection of several different formal techniques, both logic-based and cognitively-driven models, that can be used as control mechanisms in robotic domains or be used to reason about changing spatial environments in general. Control approaches based on the following formalisms will be available for use in an *independent* manner within the proposed framework:

1. A basic STRIPS like planning system [Fikes and Nilsson, 1990]
2. Belief-Desire-Intention (BDI) Approach [Bratman, 1987]
3. Event calculus [Kowalski and Sergot, 1986]
4. Fluent calculus [Thielscher, 1998]
5. Situation calculus [McCarthy and Hayes, 1969]

Several high-level languages that are directly based on the above mentioned formal approaches are already available, e.g., the GOLOG family of languages based on the situation calculus [Levesque et al., 1997, Lakemeyer and Grosskreutz, 2001], FLUX – the fluent calculus based language [Thielscher, 2005] and a discrete event calculus based reasoner [Mueller, 2007]. The utility of the aforementioned control calculi and the high-level languages that are based on these calculi for the modelling of dynamics cannot be taken granted – rather fundamental problems (e.g., frame, ramification, qualification) relevant to modelling changing environments have been thoroughly investigated in the context of the class of formalisms aforementioned [Shanahan, 1997]. This has also resulted in several non-monotonic extensions to classical symbolic approaches that are better suited for modelling dynamically changing systems [Bhatt and Loke, 2008, Bhatt, 2008] and representing cognitively adequate (human-like) common-sense reasoning with incomplete information. By including these diverse control approaches within the framework, the objective is to facilitate and promote the experimental, pedagogical and other potential uses of the framework. Finally, it should be noted that the proposed framework directly embeds such high-level languages in a manner that completely abstracts from the control approach specific details by the use of the generic domain description component in (C1).

C3. Underlying Qualitative Physics

Of key interest to this work is to operationalize the notion of a domain-independent qualitative spatial theory, which is representative of an underlying 'qualitative physics' that is applicable for a wide-range of dynamic spatial domains/systems. Here, by a dynamic spatial system, we refer to a specialization of the dynamic systems [Sandewall, 1994] concept for the case where a domain theory consists of changing qualitative spatial relationships pertaining to arbitrary spatial aspects such as the orientational [Freksa, 1992, Moratz et al., 2000], directional [Frank, 1996, Ligozat, 1998] and topological [Randell et al., 1992] spatial dimensions. Basically, what this implies is that spatial relationships are modelled as time-dependent properties (i.e., fluents) and the manner in which they change are strictly governed by the rules (conceptual neighborhoods, compositional consistency and so forth) that are intrinsic to a particular spatial calculus (e.g., topological or orientation calculi) that is being modelled as a part of the underlying qualitative physics. This notion of a domain-independent qualitative spatial theory within the framework is primarily used as a means to demonstrate the applicability of (existing) qualitative spatial models relevant to different aspects of space in realistic dynamic spatial scenarios. In addition, such a theory has the advantage of being general and re-usable in a wide-range of spatial domains. In [Bhatt and Loke, 2008], we have presented an in-depth study of realizing such a domain-independent spatial theory in the context of the situation calculus formalism and presently, work is is progress to extend the approach therein to event calculus and fluent calculus.

C4. Application Platform Independence

It is necessary that the framework be independent of any particular robotic system/platform or agent simulation environment thereby ensuring applicability in a wide-range of real or simulated environments. Basically, an adequate level of abstraction between the experimental framework and robotic hardware and simulated systems is necessary. Toward this end, the framework consists of a Controller Communication Interface' (CCI) that provides the necessary abstraction between robotic or simulation platforms and the experimental framework. This independence is achieved by the generic CCI by explicitly defining all possible modes of communication (e.g., by way of serializing control actions to the robot's actuators and the inflow of sensing information) between the framework and the external world, which the framework is being interfaced with. Other details are included in section 3 (T3).

3 Technical Overview and Progress

The discussion in section 2 is intended to provide an overview of the complete framework, as it is presently envisaged. In this section, we report the preliminary progress in that direction and highlight our working exemplar that implements parts of the proposed framework. Because of the work-in-progress nature of the proposal, we only discuss aspects where conclusive implementations have been realized.

T1. The Control Apparatus – Reasoning about Action and Change

The framework presently embeds control approaches based on the STRIPS, BDI and an existing interpreter for the situation calculus based language INDIGOLOG [Giacomo and Levesque, 1999]. Without going into the details of any of these approaches, we would like to mention that the case of the STRIPS and BDI based control approaches is trivial in comparison to embedding the interpreter for INDIGOLOG. Indigolog supports the incremental execution of high-level agent control programs through the interleaving of planning, sensing and executing actions in the real/simulated world, i.e., sensing affects subsequent computation. The present communication controller interface (see T2) is minimal and has been designed to conform to the requirements of the INDIGOLOG interpreter, namely – serializing primitive actions execution commands to an arbitrary sink that is connected to the control module, reporting of exogenous events from the external world back to the control module and the capability to perform sensing actions to determine the state of certain properties of the world. The inclusion of the control approaches based on event calculus and fluent calculus is subject to further work and the completion of a complete working exemplar consisting of only STRIPS, BDI and INDIGOLOG.

T2. Controller Communication Interface

In the present exemplar, the controller communication interface has been designed to comply with an existing robotic hardware abstraction platform, namely the PLAYER/STAGE/GAZEBO project that available in the open-source domain [Gerkey et al., 2003]. With a network-centric client-server model, PLAYER provides an interface to a variety of robot and sensor hardware and allows for robot control programs to be written in any programming language and to run on any computer with a network connection to the robot. Since it is not an objective of this project to directly investigate the seamless integration of arbitrary real robotic or simulation platform, using the robot control abstractions provided by the PLAYER system within our exemplar is advantageous because of the following reasons:

1. PLAYER uses a generic API to control a wide range of robotic platforms thereby maximizing applicability in realistic applications
2. The accompanying STAGE and GAZEBO projects provide accurate physical simulators for the 2D and 3D case respectively that may be transparently used in conjunction with the PLAYER system, i.e., experiments may directly switch between real robotic and simulated modes

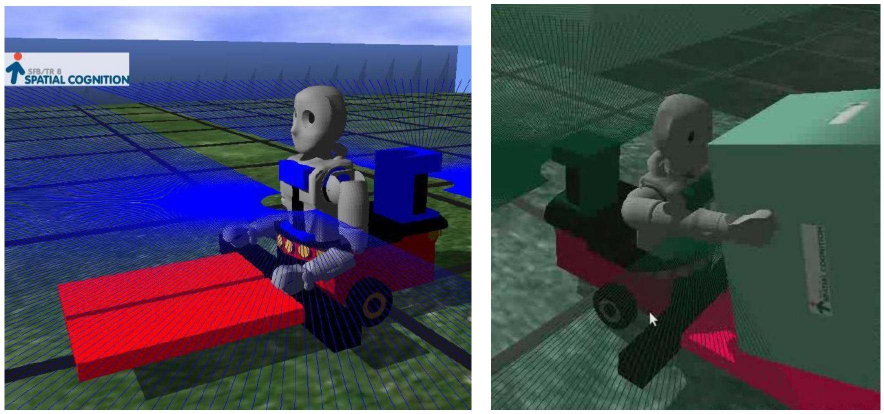

(a) Laser Guided Mobile Humanoid Gripper (b) Humanoid in Action

Figure 2: A Simulated Delivery Robot

without any change in the overall system architecture, and finally

3. The PLAYER/STAGE/GAZEBO project is open-source and continuously being enhanced and updated, which is clearly desirable from a long-term usability viewpoint.

Note that a full-integration of the PLAYER system being an open-ended task is considered beyond the scope of this project. However, working examples for a few PLAYER compatible robotic platforms in the context of the CCI have been demonstrated (see T3 and T4). To re-iterate, the main aim of this project is to develop a framework that can be utilized for high-level control and decision-making. As such, we only implement as much low-level motion control or a primitive skill set as is necessary for us to illustrate the utility of the framework for this purpose.

T3. A MOBILE HUMANOID ROBOT

The BANDIT model available within the GAZEBO simulator has been modified and extended to realize a laser-guided, mobile and grasp capable humanoid robot (see Fig. 2). The humanoid model consists of two simulated SICK lasers for simultaneous forward and backward alignment, a gripper to grasp and lift idealized objects, and a moving platform located at the intrinsic front of the bot on/from which objects may be loaded/unloaded. Indeed, the entire humanoid is mounted on a Pioneer 2DX that is capable of moving around using differential motors. When big objects that obstruct the robot's view are loaded on this platform, the backward laser is used for purposes of movement and alignment (see Fig. 2(b)). The model functions as one unit and is capable of moving around, performing turn actions at varying degrees and picking-up and dropping objects. This primitive skill set is sufficient for our exemplary purposes of realizing one complete exemplar that has high-level, logic-driven reasoning with qualities that is completely abstracted from the precise low-level motion control that occurs within the GAZEBO simulator.

T4. DEMONSTRATIVE SCENARIOS

Two exemplary scenarios have been implemented to realize a fully-functioning system consisting of high-level reasoning on the one hand and low-level motion control with the simulated humanoid robot on the other. One scenario consists of a delivery system that delivers objects from one location to another in an idealized 3D world. Another scenario involves the same simulated robot performing a room-clearing task; here, the objective is to re-arrange a set of objects in a room in a pre-specified manner. Indeed, both scenarios utilize the same set of primitive skills in so far as the movement and object manipulation are concerned. Note that given the outlook of this work (see section 4), we are intentionally focusing on problems that involve qualitative spatial reasoning abilities with orientation and topological information.

4 OUTLOOK

Present work is focused on developing a complete exemplar of the overall framework as proposed in section 2. Toward that end, of primary importance is incorporating a PDDL based domain description language and its mapping to the domain-theory and behavior specification constructs as required by the situation calculus based INDIGOLOG. Preliminary studies show that such a language subsumes similar requirements of the STRIPS and BDI control approaches. Secondarily, albeit purely in the context of the situation calculus, we are also integrating formal spatial (intrinsic orientational and topological) calculi in a way such that qualitative spatial reasoning in the form of consistency and conceptual neighborhood based dynamics may be utilized in arbitrary spatial scenarios. The development and illustration of a test-suite of problems in spatial control, primarily encompassing spatial planning and decision-making in real robotic-control and simulated environments is one of the main aims of this research. The test problems would be used to determine the feasibility of the implemented control approaches and also to perform empirical comparisons amongst them. In addition, they would also be extensively documented from an illustrative viewpoint so as to serve as examples for the utilization of the experimental framework by other users or to be used for pedagogical purposes.

ACKNOWLEDGMENTS

The author acknowledges funding provided by the ALEXANDER VON HUMBOLDT STIFTUNG (GERMANY) toward this project. This work originated and benefited from my previous collaboration with Maurice Pagnucco (UNSW, AUSTRALIA). Practical contributions by Sanchit Sood (IIT BOMBAY, INDIA) and Hemant Agrawal (IIT ROORKEE, INDIA) in implementing the low-level control with PLAYER/GAZEBO are acknowledged.

REFERENCES

M. Bhatt. (Some) Default and non-monotonic aspects of qualitative spatial reasoning. In *AAAI-08 Technical Reports, Workshop on Spatial and Temporal Reasoning*, pages 1–6, 2008. ISBN 978-1-57735-379-9.

M. Bhatt and S. Loke. Modelling dynamic spatial systems in the situation calculus. *Spatial Cognition and Computation*, 8(1):86–130, 2008. ISSN 1387-5868.

M. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.

J. Claßen, P. Eyerich, G. Lakemeyer, and B. Nebel. Towards an integration of golog and planning. In M. M. Veloso, editor, *IJCAI*, pages 1846–1851, 2007a.

J. Claßen, Y. Hu, and G. Lakemeyer. A situation-calculus semantics for an expressive fragment of pddl. In *AAAI*, pages 956–961. AAAI Press, 2007b. ISBN 978-1-57735-323-2.

A. Cohn and S. Hazarika. Qualitative spatial representation and reasoning: An overview. *Fundam. Inf.*, 46(1-2):1–29, 2001. ISSN 0169-2968.

G. De Giacomo, Y. Lésperance, and H. J. Levesque. ConGolog, A concurrent programming language based on situation calculus. *Artificial Intelligence*, 121(1–2):109–169, 2000.

R. E. Fikes and N. J. Nilsson. Strips: A new approach to the application of theorem proving to problem solving. In J. Allen, J. Hendler, and A. Tate, editors, *Readings in Planning*, pages 88–97. Kaufmann, San Mateo, CA, 1990.

A. U. Frank. Qualitative spatial reasoning: Cardinal directions as an example. *International Journal of Geographical Information Science*, 10(3):269–290, 1996.

C. Freksa. Conceptual neighborhood and its role in temporal and spatial reasoning. In M. Singh and L. Travé-Massuyès, editors, *Decision Support Systems and Qualitative Reasoning*, pages 181–187. North-Holland, Amsterdam, 1991.

C. Freksa. Using orientation information for qualitative spatial reasoning. In *Proceedings of the Intl. Conf. GIS, From Space to Territory: Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, pages 162–178, London, UK, 1992. Springer-Verlag. ISBN 3-540-55966-3.

B. P. Gerkey, R. T. Vaughan, and A. Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *ICAR 2003*, pages 317–323, Coimbra, Portugal, June 2003.

G. D. Giacomo and H. J. Levesque. An incremental interpreter for high-level programs with sensing. In H. J. Levesque and F. Pirri, editors, *Logical Foundation for cognitive agents: contributions in honor of Ray Reiter*, pages 86–102. Springer, Berlin, 1999.

H. Grosskreutz and G. Lakemeyer. cc-Golog: Towards More Realistic Logic-Based Robot Controllers. In *NMR-00*, 2000.

R. Kowalski and M. Sergot. A logic-based calculus of events. *New Gen. Comput.*, 4(1):67–95, 1986. ISSN 0288-3635.

G. Lakemeyer and H. Grosskreutz. On-line execution of cc-golog plans. In *IJCAI*, pages 12–18, 2001.

H. Levesque and G. Lakemeyer. Chapter: Cognitive robotics. In V. Lifschitz, F. van Harmelen, and F. Porter, editors, *Handbook of Knowledge Representation*. Elsevier, 2007.

H. J. Levesque, R. Reiter, Y. Lespérance, F. Lin, and R. B. Scherl. Golog: A logic programming language for dynamic domains. *J. Log. Program.*, 31(1-3):59–83, 1997.

G. Ligozat. Reasoning about cardinal directions. *J. Vis. Lang. Comput.*, 9(1):23–44, 1998.

J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

D. McDermott, M. Ghallab, A. Howe, C. Knoblock, A. Ram, M. Veloso, D. Weld, and D. Wilkins. PDDL – The Planning Domain Definition Language – version 1.2. In *Technical Report CVC TR-98-003, Yale Center for Computational Vision and Control*, 1998.

R. Moratz, J. Renz, and D. Wolter. Qualitative spatial reasoning about line segments. In *ECAI*, pages 234–238, 2000.

E. T. Mueller. Discrete event calculus reasoner. In *System Documentation, IBM Thomas J. Watson Research Center*, 2007. URL `http://decreasoner.sourceforge.net/`.

D. A. Randell, Z. Cui, and A. Cohn. A spatial logic based on regions and connection. In *KR'92. Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference*, pages 165–176. Morgan Kaufmann, San Mateo, California, 1992.

E. Sandewall. *Features and Fluents (Vol. 1): The Representation of Knowledge about Dynamical Systems*. Oxford University Press, Inc., New York, NY, USA, 1994.

M. Shanahan. *Solving the frame problem: a mathematical investigation of the common sense law of inertia*. MIT Press, 1997. ISBN 0-262-19384-1.

M. Thielscher. Introduction to the fluent calculus. *Electron. Trans. Artif. Intell.*, 2:179–192, 1998.

M. Thielscher. Flux: A logic programming method for reasoning agents. *Theory Pract. Log. Program.*, 5(4-5):533–565, 2005. ISSN 1471-0684.