


Accepted Manuscript

Towards a Science of Integrated AI and Robotics

Kanna Rajan, Alessandro Saffiotti

PII: S0004-3702(17)30031-0
DOI: <http://dx.doi.org/10.1016/j.artint.2017.03.003>
Reference: ARTINT 3006

To appear in: *Artificial Intelligence*



Please cite this article in press as: K. Rajan, A. Saffiotti, Towards a Science of Integrated AI and Robotics, *Artif. Intell.* (2017), <http://dx.doi.org/10.1016/j.artint.2017.03.003>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Editorial:
Towards a Science of Integrated AI and Robotics

Kanna Rajan^{a,b}, Alessandro Saffiotti^c

^a*Department of Engineering Cybernetics,
Center for Autonomous Marine Operations and Systems (AMOS)
Norges teknisk-naturvitenskapelige universitet (NTNU), Trondheim, Norway*

^b*Underwater Systems and Technology Laboratory
Faculty of Engineering, University of Porto, Portugal*

^c*AASS Cognitive Robotic Systems Lab
School of Science and Technology, Örebro University, Örebro, Sweden*

Abstract

The early promise of the impact of machine intelligence did not involve the partitioning of the nascent field of Artificial Intelligence. The founders of AI envisioned the notion of embedded intelligence as being conjoined between perception, reasoning and actuation. Yet over the years the fields of AI and Robotics drifted apart. Practitioners of AI focused on problems and algorithms abstracted from the real world. Roboticists, generally with a background in mechanical and electrical engineering, concentrated on sensorimotor functions. That divergence is slowly being bridged with the maturity of both fields and with the growing interest in autonomous systems. This special issue brings together the state of the art and practice of the emergent field of integrated AI and Robotics, and highlights the key areas along which this current evolution of machine intelligence is heading.

Keywords: intelligent robots, cognitive robotics, autonomous systems, embodied AI, integrated systems

Email addresses: Kanna.Rajan@ntnu.no (Kanna Rajan), asaffio@aass.oru.se (Alessandro Saffiotti)

1. Introduction

The fields of Artificial Intelligence (AI) and Robotics were strongly connected in the early days of AI, but have since diverged. One of the early goals of AI consisted of building embodied intelligent systems [19]. Such a goal however has shown to be quite challenging, and researchers have isolated its many distinct facets and focused on making progress on each facet separately; this has resulted in AI and Robotics developing as divergent research lines, with little in the way of cross-pollination of ideas. With advancements in both fields, there is now a renewed interest in bringing the two disciplines closer together, towards the establishment of a distinct trend in *integrated AI and Robotics*.

In AI, three factors have contributed to making this field more ready than ever to be applied to robotics. First, the rapid progress in hardware has led to more computational power in smaller form-factor devices, thus allowing embedded devices to run sophisticated algorithms. Second, and partly enabled by the first trend, is the exponential increase in data because of increasing growth of digital content from the Internet. Machine learning techniques, popularly grafted onto “big data” as a result, have become a burgeoning field of applied research. Third, having gone through an ‘AI Winter’, researchers have not only become more adept at divining what techniques and representations are likely to be promising, they have also made substantial inroads into the science of both AI and Robotics — contributions to this special issue clearly showcase this.

Robotics on its side has matured enormously in the last two decades. Common robotic platforms are now available, together with reliable techniques and shared tools to solve basic perception, navigation, and manipulation tasks; the world of open source software and hardware has only pushed this further along. From a socio-economic perspective, robotics is predicted to be one of the fastest growing markets in the next 15 years, with a double-digit growth caused by the pervasive introduction of robots in production, service and home-care sectors [20]. Policy makers in Europe and the US agree that, to enter these new markets, future robots will need to be more versatile and address a range of tasks in less engineered and more open environments [21, 22]. To do so, they will need to rely on cognitive capabilities such as knowledge representation, planning, learning, adaptation, and natural human-robot interaction. These are precisely the capabilities that have formed the focus of AI research for 60 years. Thus, future robots will need to

increasingly incorporate techniques pioneered by the AI community. Perhaps more importantly, the integrated use of AI and Robotics technologies is expected to enable disruptive innovation in many products and services beyond robotics, including domestic appliances, intelligent vehicles, assistive home care, as well as in the strategic areas of Autonomous Systems, Cyber-Physical Systems and the Internet of Things.

In the face of this growing interest, we lack a clearly defined field of *integrated* AI and Robotics, which articulates methods, representations and mechanisms that lead to enablement, while laying out challenges and solutions. With few exceptions, AI and Robotics are today regarded as separate fields, that belong to different academic partitions (computer science for AI, mechanical or electronic engineering for robotics), and have their own respective communities and scientific traditions. Typical graduate student curricula concentrate on AI or on Robotics, but rarely on both, and students in one area are seldom aware of the concepts and achievements in the other. This situation is not unlike the one described in 1948 by Norbert Wiener [23, p. 2]:

There are fields of scientific work . . . in which . . . important work is delayed by the unavailability in one field of results that may have already become classical in the next field.

One important facet of this problem is that we do not have a major scientific venue where the two communities can discuss research questions and findings, and in doing so share a common vernacular. Such a venue is critical to foster a much needed conceptual integration and dialogue among researchers. They need to be familiar with the methods and limitations in the two fields, to do work across their traditional boundaries, and to combine theoretical insights from both areas with practical understanding of physical robotic systems while converging towards a common *lingua-franca*.

This special issue on “AI and Robotics” provides one such venue, presenting a carefully selected set of successful integration efforts. It is our hope that it will constitute an important stepping stone in building a community of AI and Robotics researchers working across their traditional boundaries towards an ultimately shared scientific goal, thus setting the stage for a new science of integrated AI and Robotics.

2. AI & Robotics: A Historical Perspective

The early pioneers of Artificial Intelligence had envisioned the use of computers as tools for encapsulating cognition within machines. While they did not explicitly articulate the need for a robotic embodiment of a form of human intelligence, this objective as a means to study cognition was foremost on their minds [19].

Conversely, mechanical tools were conceived early in the history of humanity as a way to help humans to conquer nature and as aids in doing “work”. Mechanical automata from China and Japan were designed more as artifacts of human curiosity than of actual practical labor saving devices [24]. With the advent of the modern computer, the transition of such devices from mechanical to electronic control led in turn to the birth of mobile automata. The term ‘robot’ itself, coined by a playwright and popularized as a field by a science-fiction writer, shows the broad-based interest of what we now consider as embedded automated cognition. It was only natural that these three areas, AI/Cognition, Robotics and modern electronics, come together with the advent of micro-electronics, inexpensive sensors, and the increasing digitalization of our world.

An early use of AI was in automatic problem solving. If one could formalize both a problem and its solutions in a suitable symbolic language, it was reasoned, then computer-based symbol manipulation could be used to go from the former to the latter. The main challenge was believed to lie in finding algorithms that realize this transformation. Once these are available, it would suffice to automatically generate the problem formulation from perception (e.g., `box(b1)`, `banana(b2)`, ... in a canonical monkey-banana problem) and to automatically enact the solution by actuation (e.g., `push(b1,p1)`, `climb(b1)`, ...), to have a software agent solving problems in the physical world.

The vision of automated problem solving in the physical world was first realized in concrete form in the Shakey project, arguably the first robot that used AI methods [25]. Shakey was a success, in that it demonstrated substantial reasoning capabilities, including the ability to plan its action and to react to unexpected events during execution. However, while embedded in a physical world, its environment was entirely artificial, made of simple geometric shapes in a well lit laboratory. In moving to more natural surroundings, the challenges of going from unreliable sensor data to a formal, symbolic problem description turned out to be much more substantial than

expected. The same was true for going from a formal, symbolic description of the solution to physical actuation. These were a major driving force to break the problem down into subproblems in order to solve them separately. Reasoning, perception and physical action were carved up, consequently.

Most AI researchers were driven by top-down (and abstract) reasoning, rarely connecting with the physical world. Reasoning ignored the physicality and embodiment in automation, assuming that a symbolic description of the problem was given, and that a symbolic description of the solution could actuate in the physical world, leading to decades of “disembodied” AI. Conversely, researchers in the field of robotics took a more bottom-up approach, mostly ignoring the reasoning aspects and focusing instead on the embodiment. Each community focused on their own problems and methods. Often however, the overlaps between the two became apparent; a telling example is in ‘planning algorithms’ which while sounding relevant from both ends of AI and Robotics, were a source of confusion to the other. The textbooks by Ghallab, Nau and Traverso [26] and by Lavelle [27] are a case in point. Both delved in-depth into ‘planning’, yet were referenced by the two communities for understanding two different classes of problems and methods: one about task planning (mostly discrete) and the other about motion planning (mostly continuous). (See [28] for an early effort to sort out this confusion.)

At the same time, the AI community held the practicality of problem solving in the real-world as hard, challenging or counter-productive from an algorithmic development perspective. The disconnect with real world problem solving in turn led the community down a variety of paths which have had little long-term impact; a case in point are the Planning competitions,¹ where there is substantial algorithmic overfitting to the rules of the competition, while much remains to be desired in terms of the connect to real-world problem solving. Similar issues plague other sub-fields [29] of AI.

Decoupling the two fields was of course highly effective: enormous progress has been made in AI and in Robotics, albeit separately. Sensor fusion, SLAM, machine learning, and to some extent automated planning, for instance, are now considered mature sub-fields of the two communities with significant bodies of work and wide applicability. It was in this context that a range of activities were pushed towards reuniting AI and Robotics at the end of the 90’s, bringing physical embodiments and reasoning together. Early signs of

¹<http://icaps-conference.org/index.php/main/competitions>

this were visible in the robot Flakey [30], in the museum tour robots Rhino and Minerva [31, 32] and in NASA's Remote Agent experiment on the Deep Space One spacecraft [33, 34]. At about the same time the international community set up the RoboCup challenge [35] with the explicit long-term goal to advance robotics and AI, and the CogRob workshop series on cognitive robotics was initiated.

It was, however, not until the beginning of the current decade that the diverse initiatives working on the reunion of AI and Robotics started to come together as a critical mass. This has since been sustained by events at major AI conferences with an outreach to the robotics field, and vice versa, like robotic competitions and tracks at IJCAI, AAI and ICAPS, and special sessions on cognitive abilities at ICRA and IROS. Today, we are witnessing an increasing commercial interest in products that require a combination of robotics and AI technologies, like autonomous cars and industrial vehicles. Silicon Valley in California, the hot-bed of innovation, is now in the throes of a major push to commercialize products in AI and Robotics [36].

Our own initiatives (and consequently this Special Issue) were initiated at a meeting at Örebro, Sweden, in December 2012.² This brought together researchers from Europe and the US to initiate a joint activity to channel a long-term effort in bringing these two disciplines together. Not only has this meeting led to planning this Special Issue, but it has also regularly brought together students and academics across the world to a series of Winter Schools,³ to study the state of the art and practice in integrated AI and Robotics. With these events and this Special Issue, we hope a vibrant community of researchers will take on the challenge of pulling the diverse threads across these disciplines and confront new challenges.

3. The Research Landscape of AI and Robotics

Since AI and Robotics have been disjoint for so long, and since they developed within different academic compartments, integrating results from these fields is especially challenging. Topics and solutions in one field are often unknown to researchers in the other. Even when they are known, these may be framed in a language which is unfamiliar to researchers in the other field, or they may assume away aspects that are instead crucial in that

²<http://aass.oru.se/Agora/Lucia2012/people.html>.

³<http://aass.oru.se/Agora/Lucia/>.

field. As a consequence, researchers in the latter field may perceive those techniques as inadequate or irrelevant to their problems, and instead develop new solutions.

These difficulties are best seen by looking at an example. *Semantic maps* constitute an active research area in robotics today. Semantic maps enrich spatial representations of the environment with information about the type and functionality of the elements in it, and enable the possibility to reason about this information [37, 38]. One would expect that this research line leverages the rich AI tradition in knowledge representation and ontologies. Instead, much of the current work on semantic maps in robotics makes little or no use of this tradition, and has sometimes developed solutions that are entirely independent from the concepts and techniques in AI. Part of the reason is insufficient mutual awareness between these two fields. But a deeper reason is that a great deal of the work in knowledge representation and ontologies is not readily usable in robotics *as is*. For example, much of that work assumes crisp and consistent knowledge; but uncertainty and contradictions are inevitable when information comes from sensors and physical embodiment. Also, much of that work assumes discrete models where objects and events have clear boundaries; but robots live in a continuous world, where metrics matter, and where the boundaries among objects or events may be blurred or difficult to extract.

What this example shows us is that it is not enough to be familiar with the results of both fields to be able to integrate them. Instead one must identify the key factors that hinder integration, and address them as fundamental research questions. We claim that most of these factors have a common origin: Robots operate in a closed loop with the physical environment, and must necessarily deal with aspects like strict timing, continuous values, complex dependencies and uncertainty; but tools and techniques developed in AI often assume those aspects away. The key question, then, is: what are the critical aspects that AI techniques must consider in order to be applicable to robotics?

To discuss this question, we find it convenient to group AI techniques according to the type of *cognitive ability* that they enable, and that we want our future robots to possess. We consider here three families of (non-disjoint) cognitive abilities, that we label CA1–CA3.

CA1: Robots that know, addresses the representation, creation and maintenance of knowledge in robotic systems; roughly put, CA1 refers to the

abilities to gain access to the relevant models; these include knowledge acquisition, learning, and adaptation.

CA2: Robots that reason, addresses reasoning in robotic systems; reasoning can be about past and present states (perception and understanding), future states (deliberation and planning), or both; reasoning typically requires models, so CA2 depends on CA1.

CA3: Robots that cooperate, addresses the use of knowledge and reasoning in multi-actor systems; this CA includes all forms of human-robot interaction, multi-robot cooperation, and knowledge sharing.

This grouping of cognitive abilities is deliberately conservative in the context of AI, even if it is comprehensive: the above CAs cover traditional techniques (e.g., classical planning in CA2) as well as more recent methods (e.g., deep learning in CA1). Traditional methods in the field of robotics do not have much to say about these abilities, while they have always been in focus for AI researchers. AI methods that address these abilities, however, often rely on assumptions that make them unsuitable to address challenges that are crucial for robotic systems. For the purpose of this discussion, we put forward four such *robotic challenges*, RC1–RC4.

RC1: Uncertainty Robots sense and change the physical world, but algorithms read and write internal variables. One cannot expect that these variables accurately reflect the state of the physical environment. In general, a robot’s primary source of knowledge about the environment, sensor data, is noisy, limited in range and error-prone, thus leading to information that is uncertain, incomplete, and possibly contradictory. The same is true for actuation and its effects on the environment.

RC2: Gaps in knowledge The results of reasoning, planning and acting can only be as good as the models which they are based on. But in any realistic, open domain it is unlikely that a robot has complete factual and operational knowledge. Robots must be able to tolerate these knowledge gaps, and possibly fill the missing knowledge when needed.

RC3: Complexity The diversity, size, dynamics and interactivity of robot domains implies great complexity of knowledge and processing. Moreover, a robot’s decisions and actions must comply to the timescale of

its environment. This challenges systems to find ways to filter their knowledge and to prune or restrict their reasoning.

RC4: Hybrid reasoning Most AI techniques address a single cognitive ability and rely on a single type of knowledge. But robots must deal with a diversity of cognitive abilities and of types of knowledge, and must do so in an integrated way. Ontological, temporal, spatial, and causal knowledge need to be used together, as must continuous and discrete information.

The applicability of a given AI technique, designed to provide a certain cognitive ability, to robotics may depend on its ability to address one or more of the above challenges. In general, cognitive abilities and robotic challenges can be organized in the matrix shown in Figure 1. Each cell at the intersection of a cognitive ability and a robotic challenge can host a family of open *research questions*: how to extend the existing AI techniques that implement that ability, so that they can cope with that challenge. Each such research question would enable, if solved, a step forward toward the realization of intelligent robotic systems.

Consider again the semantic map example. Many ontology representation formalisms (which are part of the CA1 family) are not directly usable in semantic maps because of their inability to deal with uncertain and conflicting information (RC1), or with the combination of symbolic and geometric information (RC4). In order to use those techniques in semantic maps, then,

| | | ROBOTIC CHALLENGES | | | |
|---------------------|-----------------------|--------------------|-------------------|------------|------------------|
| | | Uncertainty | Gaps in knowledge | Complexity | Hybrid reasoning |
| COGNITIVE ABILITIES | Robots that know | | | | |
| | Robots that reason | | | | |
| | Robots that cooperate | | | | |

Figure 1: The research landscape of the field of integrated AI and Robotics. Key research questions in this field lie at the intersections of rows and columns.

one would need to extend them to address those challenges. In doing so, non-trivial research questions which lie at the intersection of CA1 and RC1, or at the intersection of CA1 and RC4, need to be addressed. This is graphically shown in Figure 2.

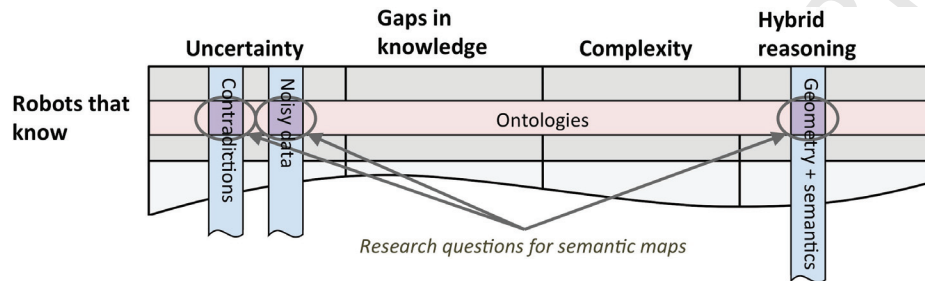


Figure 2: The semantic map example in the above landscape. The circles indicate some areas where key research questions need to be addressed in order to use ontology formalisms in robotics.

The matrix in Figure 1 gives structure to the research landscape in the field of integrated AI and Robotics. We believe that it is the research questions at the intersections of the rows and columns in this matrix that characterize this field. Coming with a well-defined set of such questions would be an important first step to define a much needed science of integrated AI and Robotics. The papers in this Special Issue contribute to this process, by addressing an interesting sample of questions across such a landscape.

4. About this Special Issue

This Special Issue is based on an open call to both the AI and Robotics communities.⁴ Fifty submissions received were deemed appropriate for the Special Issue; the nominal and systematic peer-review standards of the AI Journal brought the final count of papers down to the 18 contributions collected here. Each paper was typically reviewed by three individuals in addition to one of the guest editors. A majority of the papers went through two rounds of revisions; a few had to go through additional rounds. In most cases, the same individuals reviewed the papers over the multiple stages, ensuring a high-quality and consistent feedback to the authors. A total of upwards of

⁴See <http://aass.oru.se/Agora/AIJ2014/>.

200 reviewers were recruited. The entire process took considerable time and effort, delaying the final version of this Special Issue till now.

The geographical distribution⁵ of papers, graphically represented in Figure 3, shows an interesting pattern. A majority of submissions (29) came in from Europe — what has traditionally been seen as a foundation of AI and Robotics research, the United States and Canada, sent in substantially fewer papers (12). While difficult to point to a precise cause, one possible explanation is that integrated AI and Robotics work is indeed going on within the Americas, but perhaps mostly in the commercial realm where there are substantially fewer incentives to publish in peer-reviewed publications. This would also be consistent with increased competition and the lag in funding from US funding agencies in comparison to the output of PhD’s. Conversely, the European Union’s Framework Programs have had a substantial up-tick in AI and Robotics funding within the FP6, FP7 and H2020 programs. Thirteen of the contributions in this Special Issue acknowledge support from these programs.

Regarding the coverage of topics, while no outcome or expectations was designed by the guest editors, we were delighted to note that the end result is both broad and deep with respect to the landscape introduced in the previous section, as shown in Figure 4.⁶ The 18 contributions to this special issue illustrate how AI methods and techniques addressing a given cognitive




Figure 3: The overall geographical distribution of papers submitted to this Special Issue.

⁵Based on author affiliations when papers were submitted.

⁶Each contribution in this volume is represented by a single rounded rectangle corresponding to its main focus. This is arguably a oversimplification made for presentation convenience, since many papers also touch on other abilities and/or challenges.

ability can be extended to address some of our four robotic challenges. While such a categorization is subjective, we feel it reflects the cross-disciplinary nature of this Special Issue and provides the AI generalist with a broad overview of what is clearly a re-emergent and symbiotic field.

The individual papers in this Special Issue are summarized below with the perspective of this landscape.

4.1. Papers focusing on CA1: Robots that know

Six of the contributions in this Special Issue have the acquisition or representation of knowledge as their center of gravity. These are highlighted in blue in Figure 4.

[2] and [7] use learning to let the robot acquire new knowledge, and therefore directly address challenge RC2. In [2], the robot gains general knowledge of a domain prior to being given any specific task. The paper articulates an algorithm for intrinsically motivated learning, that learns models of the transition dynamics of a domain using random forests while combining intrinsic with external task rewards. [7] proposes an online algorithm that enables a humanoid robotic agent to incrementally learn new skills in the order of increasing learning difficulty, from its onboard high-dimensional camera inputs and low-level kinematic joint maps, driven purely by its intrinsic motivation.

Papers [13] and [16] also deal with learning, but put a strong emphasis on managing the complexity of doing so, thus addressing RC3. [13] introduces a scalable methodology to learn and transfer knowledge of the transition (and

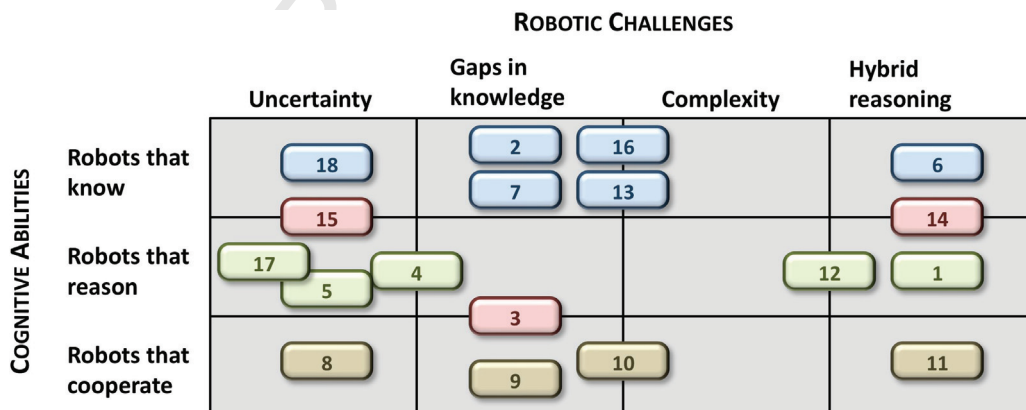


Figure 4: The contributions in this Special Issue as mapped to the above research landscape. Colors refer to the subsections below.

reward) models for model-based reinforcement learning in a complex world. The authors use a formulation of Markov decision processes that support efficient online-learning of relevant problem features in order to approximate world dynamics. [16] presents a reinforcement learning approach for aerial cargo delivery tasks in environments with static obstacles. The authors plan and create swing-free trajectories with bounded load displacements. To deal with the very large action space, learning occurs in a simplified problem space and then transferred to the full space.

[18] shows how learning can be done when the context changes between executions, thus addressing RC1. It proposes a model-based contextual policy search algorithm to deal with such contextual changes between task executions. It does so by utilizing a hierarchical approach, learning an upper-level policy that generalizes the lower-level controllers to new contexts.

Finally, [6] discusses how one can structure a knowledge base for robots for goal achievement that combines different knowledge types, thus addressing RC4. The authors propose a combination of different knowledge areas, different knowledge sources and different inference mechanisms to cover the breadth and depth of required knowledge and inferences. Various methods have been developed in the AI community for representing and reasoning about temporal relations, action effects and changing situations, most of them focusing on individual inference problems. On the other end, representations developed in robotics have largely focused on special-purpose probabilistic models that usually lack clear semantics. This paper attempts to combine parts of these approaches in a common framework to provide robots with comprehensive knowledge and inference capabilities.

4.2. Papers focusing on CA2: Robots that reason

Five contributions focus on reasoning techniques. These are highlighted in green in Figure 4.

[1] provides an overview of deliberation in AI and Robotics, a key piece of enablement to demonstrate embodied machine intelligence. Deliberation functions identified and analyzed are: planning, acting, monitoring, observing, and learning. The paper discusses their main characteristics, design choices and constraints. Being a survey, this contribution touches all the four robotic challenges, but has a special concern on integrating several different deliberation functions (RC4)

[12] also deals with the integration of different deliberation functions (RC4) but considers in addition the challenge of tackling the complexity that

is intrinsic in hybrid reasoning (RC3). The authors propose an approach to hybrid planning for complex robotic platforms in which state-based forward-chaining task planning is tightly coupled with motion planning and with other forms of geometric reasoning. Hybrid states are represented as a symbolic component coupled to a geometric component, with the former used for causal and logical reasoning and the latter for geometric reasoning.

[5] and [17] propose reasoning methods that take uncertainty into account, thus addressing RC1. [5] investigates the manipulation of multiple unknown objects in a crowded environment with incomplete knowledge, including object occlusions. Object observations are imperfect and action success is uncertain, making planning challenging. The authors use a POMDP method that optimizes a policy graph using particle filters and takes uncertainty in temporal evolution and partial observations into account. [17] presents a methodology for allowing flexibility in task execution using qualitative approaches which support representation of spatial and temporal flexibility with respect to tasks. The authors extend compilation approaches developed for temporally flexible execution of discrete activity plans to work with hybrid discrete/continuous systems, and to determine optimal control policies for feasible state regions which can deal with plan disturbance.

Paper [4] also deals with uncertain information (RC1) as well as the lack of abstract models of physical processes (RC2). The idea is to project the physical effects of robot manipulation actions by translating a qualitative physics problem formalization into a parameterized simulation problem. In doing so, it performs a detailed physics-based simulation of a robot plan, logging the state evolution into appropriate data structures and then translating these sub-symbolic data structures into interval-based timelines.

4.3. Papers focusing on CA3: Robots that cooperate

Four contributions focus on human-robot and multi-robot interaction. These are highlighted in brown in Figure 4.

Paper [8] presents a framework of anticipatory action selection for human-robot interaction that is capable of handling nonlinear and stochastic human behaviors such as table tennis strokes, and allows the robot to choose the optimal action based on prediction of the human's intention with uncertainty. The framework thus addresses challenge RC1.

[9] shows how planning operators can be learned from human teachers rather than hand-coded, thus addressing challenge RC2. The paper proposes an approach to automatically learn and execute human-like tasks by

exploiting natural teacher-learner interactions involving a human. Whenever a robotic plan halts, the human teaches an appropriate action and the robot then completes the task by itself.

[10] also deals with learning from teacher demonstrations, but puts emphasis on reducing the complexity of the learning process (RC3). The paper extends relational Reinforcement Learning algorithms by integrating the option of requesting teacher demonstrations, to learn new domains with fewer action executions and no prior knowledge. This technique shows that adding such teacher demonstrations improves the success ratio for execution while significantly reducing the number of exploration actions required.

Finally, [11] identifies key decisional issues for robots that cooperate with humans, including needed cognitive skills: situation assessment based on perspective-taking and affordance analysis; acquisition and representation of knowledge models for humans and robots, each with their specificities; situated, natural and multi-modal dialogue; human-aware task planning; human-robot joint task achievement. It shows that combining such abilities in a deliberative architecture, by pushing human-level semantics within the robot's deliberative system, leads to richer human-robot interactions. This paper therefore addresses challenge RC4.

4.4. Papers spanning multiple cognitive abilities

While most of the above papers touch at multiple cognitive abilities to some extent, they have a clear center of gravity in one; the following three contributions cover broader ground and are highlighted in red in Figure 4.

[3] presents a framework to infer human activities from observations using semantic representations, which can be utilized to transfer tasks and skills from humans to humanoid robots. This work spans across abilities CA2 and CA3, and relies on the creation of new knowledge to address challenge RC2. The framework allows robots to determine a higher-level understanding of a demonstrator's behavior via semantic representations that capture the essence of the activity and thereby indicate which aspect of the demonstrator's actions should be executed for goal achievement.

[14] presents an approach to creating a semantic map of an indoor environment from RGBD data, which is needed for object search or place recognition by robots. It therefore belongs to CA1. It also belongs to CA2, since the authors use a knowledge representation and reasoning formalism for reasoning about perceived objects during the top-down part of closed-loop semantic

mapping. The approach presented combines geometric information and semantic knowledge, thus addressing RC4.

Finally, paper [15] describes an ambitious EU project that deals with representing knowledge (CA1) and reasoning with it (CA2) under strong uncertainty (RC1). The paper describes an embedded robot platform along with a theory of how it can plan information gathering tasks and explain failure in environments that have uncertain sensing, uncertain actions, incomplete information and epistemic goals. Central to the ideas in this paper is how the robot’s knowledge is organized as well as how the robot should represent what it knows or believes.

5. Conclusions

We believe the time is ripe for AI and Robotics to pursue conjoint research lines towards an integrated approach for autonomous systems. Doing so will enable a new cohort of young researchers who have a wider field of view over the two disciplines, and can make progress towards truly embodied agents by leveraging what is now cross-disciplinary into their curriculum and removing the traditional ‘stove piping’ that often occurs in academia. And ideally to do so for societal benefit in the context of understanding how we can alleviate some of today’s problems afflicting humankind.

One grand challenge might be in harnessing and applying research towards helping understand the environment and provide tools and techniques in understanding the changing climate. This would need AI and Robotics to aid researchers in the pure (and social) sciences to obtain a finer level understanding of earthbound environmental processes, e.g., harmful algal blooms in the oceans, air pollution in cities and towns, the carbon life-cycle, etc. Such a task would require embodiment in sensors and robotic platforms exposed to the harsh conditions terrestrially, in the atmosphere, in the oceans and in space while providing novel means of measuring critical scientific variables, all the while interacting with researchers well outside the traditional comfort zone of technologists.

To enable such grand challenges requires focus and concerted action with a well thought out research agenda. One approach that has been suggested is to define a list of *problem statements* for integrated AI and Robotics [39]. A complementary, applied approach could define *competitive challenges* [40] in the style of RoboCup [35] or the DARPA initiatives [41, 42]. Irrespective, we expect that many of the research questions that will arise lie at the

intersection of the rows and columns in the matrix shown in Figure 1. The solutions to these questions would provide a core set of principles, methods and approaches, from which a curriculum can be built that defines the new science of integrated AI and Robotics. These could also be collected in a compendium of methods and techniques developed by this emerging community of scholars and updated continuously with community-driven tools. The periodic publication of a collected ‘Readings’, akin to [43] and [44], could be instrumental in helping young (or even well-established) researchers get a bird’s eye view of this nascent field.

This Special Issue is essentially the first of many steps that we believe are needed to get AI and Robotics to be more aligned and co-mingled with ideas, techniques and methods, formal and otherwise. For these disciplines to make a clear and positive impact to society, we believe there are no boundaries: collective wisdom can only come with collective action. And we believe this is the time to pull together an effective community in this field.

Acknowledgements

More than 200 reviewers provided high quality reviews over multiple rounds, to allow authors to revise and strengthen their technical content while also making their material more accessible to the general AI or Robotics reader. We, as guest editors are grateful to these reviewers who made this special issue possible; we could not have done it without their time, effort and enthusiasm.

When writing this editorial, we have benefited from extensive feedback from Malik Ghallab, Joachim Hertzberg, Maja Mataric, Nils Nilsson and Erik Sandewall, as well as four anonymous referees: their help is gratefully acknowledged. We also thank the participants of the first “Lucia” meeting,⁷ where the idea of this special issue was conceived, for inspiring discussions and motivation. The opinions in this editorial, however, are our own.

The Lucia initiatives were funded by Örebro University, Sweden, through the special strategic research grant RV422/2011. Rajan was supported in part by the United States Office of Naval Research, ONR Grant # N00014-14-1-0536. The authors are grateful to Federico Pecora and Nicola Muscettola for discussions on various topics in this Special Issue over the years.

⁷Listed at <http://aass.oru.se/Agora/Lucia2012/people.html>.

References

- [1] F. Ingrand, M. Ghallab, Deliberation for autonomous robots: A survey, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [2] T. Hester, P. Stone, Intrinsically motivated model learning for developing curious robots, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [3] K. Ramirez-Amaro, M. Beetz, G. Cheng, Transferring skills to humanoid robots by extracting semantic representations from observations of human activities, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [4] L. Kunze, M. Beetz, Envisioning the qualitative effects of robot manipulation actions using simulation-based projections, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [5] J. Pajarinen, V. Kyrki, Robotic manipulation of multiple objects as a POMDP, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [6] M. Tenorth, M. Beetz, Representations for robot knowledge in the KnowRob framework, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [7] V. R. Kompella, M. Stollenga, M. Luciw, J. Schmidhuber, Continual curiosity-driven skill acquisition from high-dimensional video inputs for humanoid robots, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [8] Z. Wang, A. Boularias, K. Mülling, B. Schölkopf, J. Peters, Anticipatory action selection for human robot table tennis, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [9] A. Agostini, C. Torras, F. Wörgötter, Efficient interactive decision-making framework for robotic applications, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [10] D. Martínez, G. Alenyà, C. Torras, Relational reinforcement learning with guided demonstrations, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [11] S. Lemaignan, M. Warnier, E. A. Sisbot, A. Clodic, R. Alami, Artificial cognition for social HumanRobot interaction: An implementation, *Artificial Intelligence XX (this issue)* (2016) xx–xx.

- [12] J. Bidot, L. Karlsson, F. Lagriffoul, A. Saffiotti, Geometric backtracking for combined task and motion planning in robotic systems, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [13] T. T. Nguyen, T. Silander, Z. Li, T.-Y. Leong, Scalable transfer learning in heterogeneous, dynamic environments, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [14] M. Günther, T. Wiemann, S. Albrecht, J. Hertzberg, Model-based furniture recognition for building semantic object maps, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [15] Marc Hanheide et al, Robot task planning and explanation in open and uncertain worlds, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [16] A. Faust, I. Palunko, P. Cruz, R. Fierro, L. Tapia, Automated aerial suspended cargo delivery through reinforcement learning, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [17] A. G. Hofmann, B. Williams, Temporally and spatially flexible plan execution for dynamic hybrid systems, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [18] A. Kupcsik, M. P. Deisenroth, J. Peters, L. A. Poh, P. Vadakkepat, G. Neumann, Model-based contextual policy search for data-efficient generalization of robot skills, *Artificial Intelligence XX (this issue)* (2016) xx–xx.
- [19] E. A. Feigenbaum, J. Feldman, et al., *Computers and thought*, AAAI Press, 1963.
- [20] International Federation of Robotics, *World Robotics 2016*, <http://www.worldrobotics.org>, 2016.
- [21] SPARC – The Partnership for Robotics in Europe, *Strategic Research Agenda for Robotics in Europe 2014–2020*, <http://www.eu-robotics.net/sparc>, 2014.
- [22] Robotics VO, *A Roadmap for U.S. Robotics*, <http://www.robotics-vo.us/Roadmap>, 2013.

- [23] N. Wiener, *Cybernetics: Control and communication in the animal and the machine*, Wiley New York, 1948.
- [24] J. Needham, *Science and Civilisation in China. Vol 2: History of Scientific Thought*, Cambridge University Press, 1954.
- [25] N. J. Nilsson, *Shakey The Robot*, Tech. Rep., SRI Artificial Intelligence Center, 1984.
- [26] M. Ghallab and D. Nau and P. Traverso, *Automated Planning Theory and Practice*, Elsevier Science, 2004.
- [27] S. LaValle, *Planning algorithms*, Cambridge university press, 2006.
- [28] T. Dean, M. Wellman, *Planning and control*, Morgan Kaufmann Publishers Inc., 1991.
- [29] K. Wagstaff, *Machine learning that matters*, in: *Proceedings of 29th International Conference on Machine Learning*, Edinburgh, Scotland, 2012.
- [30] A. Saffiotti, K. Konolige, E. H. Ruspini, *A multivalued logic approach to integrating planning and control*, *Artificial intelligence* 76 (1-2) (1995) 481–526.
- [31] W. Burgard, A. B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, S. Thrun, *The interactive museum tour-guide robot*, in: *AAAI/IAAI*, 11–18, 1998.
- [32] S. Thrun, M. Bennewitz, W. Burgard, A. B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, et al., *MINERVA: A second-generation museum tour-guide robot*, in: *IEEE International Conference on Robotics and Automation (ICRA)*, 1999.
- [33] N. Muscettola, P. Nayak, B. Pell, B. Williams, *Remote Agent: To Boldly Go Where No AI System Has Gone Before*, *AI Journal* 103 (1998) 5–48.
- [34] K. Rajan, D. Bernard, G. Dorais, E. Gamble, B. Kanefsky, J. Kurien, W. Millar, N. Muscettola, P. Nayak, N. Rouquette, B. Smith, W. Taylor, Y. Tung, *Remote Agent: An Autonomous Control System for the New Millennium*, in: *Proceedings Prestigious Applications of Intelligent Systems, European Conf. on AI*, Berlin, 2000.

- [35] H. Kitano, M. Tambe, P. Stone, M. Veloso, S. Coradeschi, E. Osawa, H. Matsubara, I. Noda, M. Asada, The RoboCup synthetic agent challenge, in: *RoboCup-97: Robot Soccer World Cup I*, Springer, 62–73, 1998.
- [36] J. Markoff, Artificial Intelligence Swarms Silicon Valley on Wings and Wheels, *The New York Times*, online at <http://nyti.ms/29ZMgCw>, 2016.
- [37] B. Kuipers, Y.-T. Byun, A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations, *Robotics and autonomous systems* 8 (1) (1991) 47–63.
- [38] A. Nüchter, J. Hertzberg, Towards semantic maps for mobile robots, *Robotics and Autonomous Systems* 56 (11) (2008) 915–926.
- [39] J. Hertzberg, personal communication, 2012.
- [40] J. Dias, K. Althoefer, P. U. Lima, Robot Competitions: What Did We Learn?, *IEEE Robotics & Automation Magazine* (2016) 16–18.
- [41] M. Buehler, K. Iagnemma, S. Singh, *The 2005 DARPA Grand Challenge: the great robot race*, vol. 36, Springer Science & Business Media, 2007.
- [42] M. Buehler, K. Iagnemma, S. Singh, *The DARPA urban challenge: autonomous vehicles in city traffic*, vol. 56, springer, 2009.
- [43] R. J. Brachman, H. J. Levesque, *Readings in knowledge representation*, Morgan Kaufmann Publishers Inc., 1985.
- [44] J. Allen, J. A. Hendler, A. Tate, *Readings in planning*, Morgan Kaufmann Pub, 1990.