I Have a Robot, and I'm Not Afraid to Use It!

Gal A. Kaminka

■ *Robots (and roboticists) increasingly appear at the Autonomous Agents and Multiagent Systems (AAMAS) conferences because the community uses robots both to inspire AAMAS research as well as to conduct it. In this article, I submit that the growing success of robotics at AAMAS is due not only to the nurturing efforts of the AAMAS community, but mainly to the increasing recognition of an important, deeper, truth: it is scientifically useful to roboticists and agent researchers to think of robots as agents.*

Robots, and therefore roboticists, have been a part of the agents community from its auspicious beginnings in the Autonomous Agents series of conferences, and continuing with the merger into the Autonomous Agents and Multiagent Systems (AAMAS) conferences. Today, there is a resurgent interest and recognition of the importance of robotics research framed within areas of research familiar to autonomous agents and multiagent systems researchers.

Robots (and roboticists) increasingly appear at the AAMAS conferences, and for a good reason. The AAMAS community is investing efforts to encourage robotics research within itself. An annual robotics special track, an associated robotics workshop (Autonomous Robots and Multirobot Systems), and a series of exciting AAMAS-sponsored plenary speakers and awards over a number of years are drawing roboticists in. The number of robotics papers is increasing. There are fruitful interactions with the other communities within AAMAS, such as virtual agents, game theory, and machine learning. Robots are being used both to inspire AAMAS research as well as to conduct it.

I posit that the growing success of robotics at AAMAS is due not only to the nurturing efforts of the AAMAS community, but mainly to the increasing recognition of an important, deeper, truth: robots are agents. In other words, there is growing recognition that it is scientifically useful to roboticists and agenticists (agent researchers) in other AAMAS areas to think of robots as agents.

Indeed, this follows in the footsteps of similar moves in the past. Some 20 years ago, artificial intelligence researchers began to recognize that thinking about agents as such is useful, and this ultimately lead to the birth of the AAMAS community. Research in agents raises challenges in integrated capabilities for intelligence, such as planning and execution, learning exploration and exploitation, strategic decision making in multiagent settings, and more. It pushes us to close the loop[1] from sensing through thinking to acting and back to sensing. It requires us to consider deeply and critically what we mean by calling a system autonomous. It raises challenges in the software architectures needed for such capabilities, and as a result raises challenges involving software engineering and programming languages appropriate for building agents. It is a field that calls on its practitioners to invest in thinking (also) about the system, rather than (only) the component.

This systemwide view is one that many roboticists share. Robotics, by nature of the research, requires its practitioners to evaluate not only a component (for example, a new vision system), but also its use within the system (the robot), and its contribution toward a design goal. Moreover, many roboticists are increasingly setting their goals higher than what we in AI (sometimes arrogantly) refer to as low-level control. The cost and stability of platforms have made it possible for roboticists to examine complex tasks, in which there is need for intelligence and knowledge, and for considering multiple robots. For example, the use of robots in assisted living, in defense and homeland security, in automation, all present significant challenges to anyone seeking to deploy robots with even partial autonomy. Agenticists have a wide variety of tools and techniques that can be brought to bear in facing both single and multiple robot challenges.

Similarly, agenticists are increasingly realizing that it is useful for them to think of robots as agent exemplars. To agent researchers, working with real robots (made of plastic, metal, electronics, and the sweat of graduate students) brings out important challenges to our current theory and practice. Robots make us fail in interesting ways and give opportunity for gaining insights otherwise unattainable. They extend the system perspective to go beyond the conceptual perception and actuation to consider sensors and motors (with their uncertainties, faults, and latencies), imperfect communications, and multiple bodies (each with its two- or three-dimensional geometry). They challenge us to understand better the concept of an environment with which the agent interacts: environments that obey certain laws (in most robots' cases, the laws of physics), and often have complex, unknown structure. Roboticists know much about these challenges and can greatly influence intellectual development within agents.

To support my argument, I report from the trenches of ongoing robotics work within the AAMAS community, highlighting success stories in which robotics research benefited from AAMAS research, and vice versa. I therefore admit in advance to a bias toward work appearing in AAMAS conferences and journals. This bias is intended to highlight robotics in the context of research areas appearing in AAMAS, as was requested of me. However, the unfortunate result of this bias is that groundbreaking work in AI and robotics appearing elsewhere (for example, Thrun, Burgard, and Fox's game-changing work on probabilistic robotics [Thrun, Burgard, and Fox 2005]) will not receive proper treatment here. This, despite such work being excellent evidence for the generality of my argument as to the usefulness of AI to robotics, and vice versa.

Building Architectures

Robotics research today must address increasingly complex missions that the robots should carry out. Previously the problem of controlling a robot could be addressed by a carefully designed controller out of the seemingly endless variety that mechanical engineers have developed over the years. Relatively basic (and very useful) tasks, such as finding a path, avoiding obstacles, and navigating to a goal location, are now mature areas of research. Instead, robots are now expected to go beyond reaching a goal to carrying out missions in which there are multiple (changing) goals and multiple tasks that should be carried out (sometimes concurrently, sometimes in sequence). Such complex missions require planning, managing resources, and in particular making decisions.

Well, this is something that agenticists in particular, and AI researchers in general, know about. Beginning with Gat's ATLANTIS architecture (Gat 1992), AI researchers have begun to integrate planners into the robot's control architecture. Following ATLANTIS's three-tier design (affectionally called 3-T architecture), a standard approach was to integrate a high-level planner together with controllers, mediated by an executive module whose role is to issue plan requests, schedule, and monitor the controller's successful execution. The key idea in such hierarchical layering is that tasks are planning problems, to be solved by a planner operating in discrete atomic steps. An executive module works at a different, finer resolution to carry out the task. Guided by the plan, it translates its discrete steps into controller instantiations. In a sense, this is a solution approach based on Newell and Simon's problem-space hypothesis (Newell 1990).

Around the same time, agents researchers have begun to advocate the idea that agents are not just planners (Pollack 1990; Pollack and Horty 1999). Rather, agents should reason about plans, generating them, adapting them, and contrasting them, to make decisions about carrying them out in service of various goals.

Agent researchers developing robots with integrated capabilities have focused on integrating planning and execution (sometimes also with learning) in a way that reflects such reasoning about plans. Here, execution and planning work at the same temporal and task resolution. The approach, called plan-based control by Beetz and colleagues (Beetz et al. 2001), relies on utilizing a plan representation as a central, first-class object, which is reasoned about, generated, adapted, revised, and managed through the lifetime of the robot. Planners are used not only to generate plans (and replan), but also to provide predictions (including of resource use and execution time). Separate processes estimate the state of the robot and the world, address perception (symbol grounding, sensor fusion, and so on), and make decisions as to alternative courses of action.

One approach built on hierarchical layering of a planner and execution modules. For instance, Haigh and Veloso (Haigh and Veloso 1997, 1998), and Simmons and colleagues reported on experiences with an office-delivery robot, Xavier (Simmons et al. 1997a, 1997b). In particular, they report on a layered architecture (Rogue) for controlling the robot. Somewhat similarly to ATLANTIS, a planner is used to plan multistep tasks. Steps are carried out by a path planner and motion controller executive module, which is also responsible for triggering replanning when necessary. However, in Rogue, plans are reasoned about and manipulated using learning. For instance, Xavier learned to avoid hallways that are crowded during specific times. Beetz and colleagues discuss plan-based control in depth and argue for a specific implementation called Structured Reactive Controllers (Beetz 2001), which allow for revising plans as opportunities and failures occur. To discover such opportunities and failures, the state of the robot and the environment are estimated using probabilistic reasoning. Similar robots, cited by Beetz et al., include Minerva (Thrun et al. 2000) (museum tour guide), and Remote Agent (Pell et al. 1997) (an experimental autonomous NASA spacecraft agent).

Thus the need to reconsider the design of the agent architecture was lead both by theorists as well as robotics researchers working in the context of agents research. The challenge of how to integrate different capabilities was met, within the agents community, with an already existing body of knowledge and significant fascination with general agent architectures: with how, in general, agents should be built. Indeed, one of the early conferences of the field was Agent Theories, Architectures, and Languages (ATAL), which focused on the theory and practice of building agents for various domains, and allowed researchers working in very different domains to discuss commonalities and differences in their designs.

Over the years, research into agent architectures that work across a wide variety of agent types and environments (including robots in various applications) has resulted in greater understanding of the architecture components and their operation. Some specific areas of research—still continuing today—are discussed in this article.

Beliefs, Desires, Intentions, and Other Mental Attitudes

First, it is by now understood that an agent operating in a dynamic environment (the settings for many robot applications) must manage the planning process. It must decide when to plan and when to avoid replanning (as it is computationally infeasible to replan with every change). To do this, the construction of the agents must allow for explicit representation of beliefs, goals, and plans (whether preplanned or dynamically generated). These will be revised, manipulated, contrasted, and reasoned about by the agents' action selection and perception processes. In other words, beliefs, goals, and plans are all first-class objects.

To a large degree, the huge literature on mental attitudes of agents, and in particular on BDI (belief desire intention) theories and architectures (Rao and Georgeff 1995; Georgeff et al. 1998; Padgham and Winikoff 2002; Sardiña and Padgham 2011) is a response to this challenge. Recent years are seeing, side by side, developments in both the theory and practice of plan representation that are amenable to both planning and execution. A variety of academic BDI implementations exists; I mention a few that have been used with robots, such as PRS (Ingrand, Georgeff, and Rao 1992; Lee et al. 1994), RPL (Beetz 2001), and RAPs (Earl and Firby 1997). There are also commercial BDI implementations specifically targeting robotic applications (for example, CogniTeam, Ltd. 2009).[2]

First-Class Plan Representations

In addition to these BDI languages that have been used in robots, there have been of course many plan representations (and sometimes programming languages) that have been tried and tested in robots, but that offer first-class status only to the plan, rather than also beliefs and goals. Nevertheless, they are useful in constructing robots that employ plan-based control. These include finite-state representations (Tousignant, Wyk, and Gini 2011; Risler and von Stryk 2008; Lötzsch, Risler, and Jüngel 2006), Petri net representations (Costelha and Lima 2008; Ziparo et al. 2010), and temporal planning and scheduling languages (for example, T-REX [Py, Rajan, and McGann 2010], which allows for multiple-resolution scheduling of tasks).

No single plan representation has emerged thus far as a clear de facto standard, and in fact the comparison of these representations remains an open challenge. Many of the BDI languages have been develop to address reported failings in finite-state machine representations (such as their lack of a factored state and limited reactivity), but a clear theoretical contrast is still lacking.

Teams of Robots

Perhaps the area in which agents research has had the most impact on robotics research, is in multirobot systems. As can be expected, AAMAS research has resulted in a very significant amount of techniques and tools that are relevant and can be brought to bear in addressing challenges in multirobot systems.

An important part of the success of AAMAS research is due to its principled, domain-independent handling of the combinatorial complexity of multirobot tasks. If multiple robots are to be coordinated in some fashion, the task of making decisions for them is more difficult than that of making decisions for a single robot, since in addition to the individual decisions, one must worry about the combinations of selected actions.

Most multirobot research to date, within the robotics community, focuses on a single task at a time. Some examples of such canonical tasks include moving while maintaining formations (Balch and Arkin 1998; Fredslund and Mataric 2002; Desai, Ostrowski, and Kumar 2001; Carpin and Parker 2002; Inalhan, Busse, and How 2000; Kaminka, Schechter-Glick, and Sadov 2008; Elmaliach and Kaminka 2008), multirobot coverage (Williams and Burdick 2006; Ferranti, Trigoni, and Levene 2007; Rekleitis et al. 2004; Zheng et al. 2005; Rekleitis, Dudek, and Milios 2001; Batalin and Sukhatme 2002; Wagner and Bruckstein 1997; Butler, Rizzi, and Hollis 2000; Hazon and Kaminka 2008; Agmon, Hazon, and Kaminka 2008), foraging (Goldberg and Mataríc 2001; Rybski et al. 1998; Rosenfeld et al. 2008; Zuluaga and Vaughan 2005; Schneider-Fontan and Mataríc 1996; Jager and Nebel 2002; Ostergaard, Sukhatme, and Mataríc 2001; Kaminka, Erusalimchik, and Kraus 2010), and patrolling or surveillance (Elmaliach, Shiloni, and Kaminka 2008; Agmon, Kraus, and Kaminka 2008; Agmon et al. 2008; Jensen et al. 2011; Basilico, Gatti, and Amigoni 2009; Smith, Schwager, and Rus 2011; Agmon, Urieli, and Stone 2011; Marino et al. 2009; Delle Fave et al. 2009). Many of these are approached from the perspective of distributed control. In other words, a controller is devised such that when it is operating in each individual robot, the total sum behavior is as required. Such controllers are built anew for each task. But as future robot applications grow in complexity, such controllers would need to take into account allocating and scheduling the execution of multiple tasks, taking place concurrently or in sequence. For instance, urban search and rescue (Murphy et al. 2008) applications require elements of both coverage and foraging and introduce additional novel tasks. Similarly, soccer (for example, in RoboCup) requires complex decision making, resource allocation, and task scheduling.

A key insight gained in the AAMAS field in the last 15 years is that, in fact, multiagent tasks can be decomposed—conceptually, as well as technically—into two components. The first, called taskwork, includes domain-dependent individual capabilities. The second, called teamwork in teams, and socialwork in general, includes the capabilities for collaboration (in teams), or maintaining other social relations. This socialwork component includes social choice mechanisms, for instance, protocols for allocating tasks to different team members (for example, by bidding), or protocols for reaching joint decisions (for example, by voting). The combination of taskwork and socialwork creates a working multiagent system for a given domain.

This insight has manifested itself in several different ways in robotics research. I will briefly discuss some of these areas of cross-fertilization between agents and robotics research and then dive in detail into one specific area (teamwork).

Market-Based Task Allocation

In terms of impact on robotics, the use of market-based methods for allocating tasks to robots enjoys widespread popularity. It is now being adopted and investigated by roboticists outside of the AAMAS community, certainly a positive sign. Starting with Dias and Stentz's work (2000) on the use of market mechanisms for coordinating robots in exploration and mapping tasks, there has been much work in this area, addressing challenges that are raised when working with robots (see, for example, Lin and Zheng [2005]; Gerkey and Mataric [2002]; Zlot and Stentz [2006]; Lagoudakis et al. [2004]; Vig and Adams [2006]; Köse et al. [2003]; Michael et al. [2008]; Tang and Parker [2007]; Lagoudakis et al. [2005]; Bererton, Gordon, and Thrun [2003]). Dias and colleagues (Dias et al. 2006) provide a comprehensive survey, and Xu and colleagues (Xu, Scerri, and Lewis 2006) provide a comparison with other methods.

The key idea in multirobot market-based task allocation is that robots bid on tasks for execution using virtual currency, which corresponds to their fitness for the tasks. A market mechanism is used to decide on the winner(s), and the tasks are allocated accordingly. In this way, robots bid rationally based on their individual fitness to the tasks, while the group as a whole benefits.

Applying markets in multirobot systems is not a straightforward process of simple implementation. Just as AAMAS research affects robotics in this application, so can robotics affect AAMAS research into markets. For example, several challenges are raised when considering the nature of the virtual currency to be used by robots.

A naïve proposal for virtual currency would base it on the spatiotemporal distances involved, and energy considerations, that is, each robot would bid on a task to be executed based on how quickly it can get to it, and at what cost to its energy reserves. Such currency, however, is not transferable, and not arbitrarily set. A robot, may report false estimates of its fitness to the task, because of its own self-interests. But it can never actually carry out a task at an arbitrary cost. For example, the duration of travel from point *A* to point *B* cannot be made arbitrarily short. Thus for instance second-price auctions cannot be simply used in such settings, as the winning robot cannot arbitrarily pay the lower cost of the second price, however much it may want to.

Moreover, even when robots are not self-interested, and bid their true valuations, failures in communications and/or uncertainty in perception may lead to incorrect bids. Thus there's a second challenge of market methods that are robust to failures, rather than malicious manipulation (Procaccia, Rosenschein, and Kaminka 2007).

Indeed, the issue of uncertainty raises a third challenge to standard multiagent auctions. Even under the best of conditions, robots would often not be able to know their actual cost for carrying out the task, due to uncertainty in their actuation, and the dynamic nature of their environments. For instance, even under pristine laboratory conditions, a robot has significant variance in how much time it is taking to travel a fixed path, with no obstacles (Traub, Kaminka, and Agmon 2011). When some elements of a dynamic environment are included (for example, the robot may be intermittently blocked by a passing pedestrian), this variance grows still. Thus the estimation of the cost for a given task—and calculation of a corresponding bid—is inherently uncertain. Some recent work is beginning to address this (Spaan, Gonçalves, and Sequeira 2010), but the challenge remains open.

Reaching Joint Decisions in Teamwork

More generally, AAMAS researchers have long discovered that teamwork involves more than task allocation. It also involves agreement on a common goal, agreement on a plan to reach the common goal, assisting teammates as necessary, and so on.

Teamwork has been investigated within the multiagent systems community for many years. Grosz, Sidner, and Kraus (Grosz and Sidner 1990; Grosz and Kraus 1996), and Cohen and Levesque (Cohen and Levesque 1991; Levesque, Cohen, and Nunes 1990) have published a series of articles on teamwork, developing logical models (SharedPlans, Joint Intentions Framework, respectively) to model and prescribe teamwork. Among other issues, these models describe the conditions under which an agent must inform its teammates of its own private beliefs, thus effectively maintaining synchronization in the team as to specific propositions. The SharedPlans teamwork model also specifies conditions for proactive assistance to teammates, mutual support, and so on.

The key benefit of this approach is that much of such teamwork can be algorithmitized. It can be described by a set of behavioral rules, which, if followed, would cause the agent to act appropriately in the context of a team, regardless of the task it was assigned, or the application domain. Unfortunately, in general, I think it is safe to say that roboticists took little notice of these theoretical frameworks, as groundbreaking as they were.

However, several autonomous agent researchers picked up on these logical frameworks and began investigations of how the frameworks might be applied in practice. Motivated by a seemingly endless stream of coordination failures in a distributed industrial system, Jennings (Jennings 1995) built on the joint intentions framework to propose the joint responsibility model, to automate coordination messages between agents within a distributed system, thus reducing the number of coordination failures. A short while later, Tambe (1997) extended the techniques involved, allowing his system (called STEAM) to consider the cost of communications in its decisions, and recover from fails. STEAM was evaluated empirically in a high-fidelity virtual environment in which synthetic helicopter pilots used the system to automate their coordination decisions.

One of the unique features of the AAMAS conference is that it is a rare forum in which both researchers of virtual humans (virtual agents) and roboticists can meet to exchange ideas. This presents a tremendous opportunity for both sides to affect each other. The demonstration of automated teamwork in virtual environments brought teamwork models close enough to robotics to get some attention from that community, especially when STEAM was shown to be applicable to the domain of virtual RoboCup soccer (Tambe et al. 1999). The benefits of automated explicit teamwork are concrete contributions that roboticists, in principle, should be very happy to adopt. Currently, many roboticists do not differentiate teams that collaborate toward shared goals, from loosely coordinating groups of essentially autistic robots. Relatively few roboticists (for example, Parker [1998]) have recognized that robot group interactions are an independent and specific object of study. Many others, instead, focus on investigating various mechanisms through which a group of robots would appear to act as a team, or in a coordinated manner, despite the robots' not having any explicit notion of their teamwork (or even of other robots).

Recognition of the benefits of automated teamwork in virtual environments has thus migrated into robotics, resulting in the Bar-Ilan teamwork engine (BITE) architecture (Kaminka and Frenkel 2005, 2007; Kaminka et al. 2007), the Machinetta framework (Scerri et al. 2003), and the CogniTAO commercial high-level control SDK.[3] All are related to STEAM, but contain various novel features. In particular, both BITE and CogniTAO target robotics specifically by providing specific features important in mobile robots: a computationally active world model, in which sensor readings are preprocessed to address uncertainty; support for maintenance goals; and flexible interaction protocols, to match the conditions of the task.

To illustrate the contribution of teamwork—as understood in state-of-the-art AAMAS—to robotics, I will describe my groups' utilization of teamwork software as part of an technology-transfer project, intended to implement a canonical multirobot task—formation maintenance—familiar to many roboticists. Given the space constraints, I settle here for a relatively high-level description; details are in Traub (2011).

In formation maintenance, robots must move in unison along a given path, while maintaining a given geometric shape. Various formation maintenance methods have been investigated (for example, Balch and Arkin [1998]; Desai [2002]; Fredslund and Mataric [2002]; Balch and Hybinette [2000]; Desai, Ostrowski, and Kumar [2001]; Carpin and Parker [2002]; Inalhan, Busse, and How [2000]; Tabuada, Pappas, and Lima [2005]; Kaminka, Schechter-Glick, and Sadov [2008]; Elmaliach and Kaminka [2008]). All of these schemes are distributed; all require each robot to run a local control process, which executes the controller that fits the role of the robot. For instance, a left-following robot in an equilateral triangle formation would keep the leader in a fixed distance (matching the distance kept by the right-following robot), such that the leader robot is at bearing 30 degrees to the right. A right-following robot would do the same, but its controller would maintain the leader at a bearing of 330 degrees (that is, 30 degrees to the left). Figures 1–4 show a number of formations, the basis for the work in Kaminka and Frenkel (2005; 2007); Elmaliach and Kaminka (2008); Kaminka, Schechter-Glick, and Sadov (2008); and Traub (2011).

The various control schemes differ in the type of operating conditions they assume, as well as in the type of performance they provide. For instance, some control schemes (called SBC for separation-bearing control) require each follower robot to be able to identify the distance and angle to a leader robot in the formation (Fredslund and Mataric 2002), based on sensor readings. In contrast, communication-based formation maintenance can be used to eliminate the need for perception by relying on dead-reckoning and communications from the leader robots (Elmaliach and Kaminka 2008). Others still use more robust schemes that allow robots to switch which robots are to be followed (Desai, Ostrowski, and Kumar 2001).

The goal of the project was to create a robust controller by tying these different control schemes together and switching between them as necessary. This creates a formation-maintenance scheme that is robust to intermittent perception and communications failures, as long as they do not coincide. The key is to switch between the different schemes, based on availability of the perception and communication processes.

Now suppose we adopt a standard robotics approach to this problem. This would entail writing a switching controller that switches between the different modes. Each such switching controller would operate on a different robot, and thus we immediately face a challenge: we need to make sure that when one robot switches, the others do as well (since mixing up formation maintenance schemes is not, in general, likely to work well). This means that we need to add code that manages communications between robots, so that when one robot finds it necessary to switch, it automatically lets the other ones know, and awaits confirmation of their switching, too. Of course, the conditions under which a robot may want to switch are not necessarily those that another robot senses, and so we also need code for them to negotiate and agree as to which control scheme the team should use. Now we just need to get all of this working for more than two robots, and more than two schemes, and across potential communication errors. And all of this still not taking into account issues such as changing roles in the formations, and so on—just a simple matter of programming, as the expression goes.

Agent researchers have long recognized that the challenges above are general. Teamwork architectures offer a general solution to cases where agents must decide on (1) when to communicate (and to some degree, what to communicate about), (2) how to come to a joint agreement (in this case,

Figure 1.Triangular AIBO Formation.

From Elmaliach and Kaminka (2008) and Kaminka and Frenkel (2005).

Figure 2. Diamond AIBO Formation.

From Kaminka, Schechter-Glick, and Sadov (2008).

which mode mode should be used by the robots), and (3) how to allocate (and reallocate) tasks to different robots. The need for such decisions comes up again and again.

Thus my research group used a teamwork architecture to manage the joint switching of controllers and the allocation of roles and tasks. We utilized the CogniTAO (CogniTeam, Ltd. 2009) commercial teamwork architecture to integrate together robust SBC (Kaminka, Schechter-Glick, and Sadov 2008) and communication-based formation-maintenance controllers (Elmaliach and Kaminka 2008). The details of the integration are well beyond the scope of this article, but the lessons are not.

Traub 2011 has carried out an analysis of the benefits of using a teamwork architecture by using a standard software engineering model (CoCoMo) (Boehm 1981) to measure its impact in automating the coordination processes described above, contrasting it with conservative and optimistic estimates of the size of the project given a standard robotics approach. The results show 50 percent to 68 percent savings in programming effort within the project, which of course translate into significant savings in both development time and number of programmers. These numbers are compatible with earlier reported results (Tambe 1997) (in fact, they are more conservative).

Teamwork in robots is a success story for AAMAS research, with measurable effects as demonstrated above. But if anyone is looking for a more bottom-line kind of evaluation, they need not look further than Kiva System's use of AI and multiagent techniques in their materials handling systems (Wurman, D'Andrea, and Mountz 2008). Kiva System's products use robots to automate order fulfillment in large warehouses. Its commercial success relies on AAMAS techniques, integrated beautifully with robotics.

Many challenges remain open in robotic teamwork, beyond those discussed above for market-based coordination. For example, when we look at teamwork in soccer, we see that humans coordinate not only based on communications, but also based on observations of each other, as well as the shared environment. But with rare exceptions (see, for instance, Huber and Durfee [1995, 1996]; Gmytrasiewicz and Durfee [2000]; Kaminka and Tambe [2000]; Agogino and Tumer [2008]), there is very little work on observation-based teamwork; and certainly no integrated architecture in which observations are used in an automated fashion. There are also very few studies that quantitatively and comprehensively demonstrate the type of software engineering improvements that are gained from using teamwork architectures. Because of this, it can be difficult to affect engineering and research outside of AAMAS and outside of the academic world (Kiva Systems' being the exception rather than the rule). Getting complete methodologies (such as Padgham and Winikoff's seminal work (2002, 2004), with serious case studies, should be a priority for the field.

A Call to Arms

My argument in this article is that AAMAS has a lot to offer robotics and also a lot to benefit from robotics, that as a field of science goes, working with robots is a useful endeavor. To show this, I have reported from the trenches of ongoing work in two specific ares of contribution, where past contributions and ongoing work are showing significant promise, both in robotics as well as in AAMAS. But there's quite a bit more; I've described the tip of the iceberg, hoping to convince you, the reader, to look at the iceberg underneath. There's a lot going on. Following are a few sample tastes.

AAMAS papers have been producing theoretical and empirical results in understanding swarms, self-organizing systems, and robot ants, all of which represent different approaches and different type of systems than the proscribed, orderly, teamwork implemented in teamwork architectures. Some of these works focus on swarms of robots, where computationally simple agents can sense each other only in their local, limited-range surroundings (for example, Turgut et al. [2008]; Gökçe and Sahin [2009]; Yu, Werfel, and Nagpal [2010]). Others focus on analyzing self-organization and theoretical properties (for example, Yu and Nagpal [2008]; Yamins and Nagpal [2008]; Mermoud, Brugger, and Martinoli [2009]; Shiloni, Agmon, and Kaminka [2009]).

While commercial vehicles with some autonomous driving capabilities are becoming a serious issue for car companies, AAMAS researchers have gone beyond the individual navigation and obstacle avoidance, investigating ways to automate intersections and traffic lights, to make the complete flows more efficient (see, for instance, Bazzan [2004]; Dresner and Stone [2008]; Bazzan [2009]; Fajardo et al. [2012]).

The use of robots in surveillance and security is not new to AAMAS roboticists (see, for instance, Rybski et al. [2000, 2002]). But in the last few years, AAMAS researchers have begun to approach the issue with game-theoretical and adversarial reasoning notions in mind, which allowed them to provide guarantees as to the optimality of the algorithms. Recent work in robots for security has focused on patrolling (Elmaliach, Shiloni, and Kaminka 2008; Agmon et al. 2008; Jensen et al. 2011; Basilico, Gatti, and Amigoni 2009; Agmon, Urieli, and Stone 2011). Much remains to be done, for instance, in integrating with reasoning for deciding on checkpoints and static security devices and in addressing motion and perception uncertainties (Agmon et al. 2009; Agmon 2010).

The use of learning as part of an integrated, complete robot has fascinated agenticists and roboticists for many years. There exists vast research on the use of learning in robots, and certainly I will not be able to do any justice to the literature in this single paragraph. Instead, let me highlight recent surveys in key areas of interest: learning from demonstration (Argall et al. 2009), multirobot / multiagent learning (Stone and Veloso 2000; Yang and Gu 2004; Hoen et al. 2006; Panait and Luke 2005; Busoniu, Babuska, and De Schutter 2008), and control (Nguyen-Tuong and Peters 2011). The opportunities and challenges for learning offered by working with robots are essentially endless, and given the high cost of manually tweaking code for robots, the motivation for learning is made crisp and clear.

Figure 3. Column Shrimps-III Formation.

From Traub (2011).

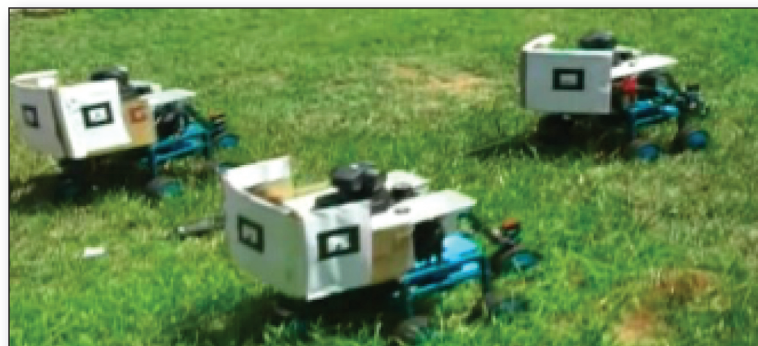

Figure 4. Triangular Shrimps-III Formation.

From Traub (2011).

Indeed, this article is also intended to be a call to arms, to invest in robot-based research. The drop in robot prices and consequent rise of the machines make robot-based artificial intelligence research in general, and AAMAS research in particular, both compelling and practical. One no longer needs to have an in-house mechanical and electronics shop to successfully conduct research involving robots. Stable platforms are now cheaply available, and their commercial maintenance makes maintaining a robot lab a feasible effort. As for venues for meeting other like-minded researchers, the appropriately named ARMS (Autonomous Robots and Multirobot Systems) workshop works hand in hand with the AAMAS conference to promote robotics research within the AAMAS community. This is a good place to start, even if your paper is yet not quite up to AAMAS's rigorous standards. Similar workshops take place on occasion with other AI conferences, including AAAI.

Comments on earlier drafts of this article have repeatedly raised the issue of the use of physical simulations and virtual environments in robotics research. Many roboticists utilize simulations as part of the development effort to ease the transition from theory and ideas to an actual working system. Paradoxically, however, roboticists often frown at research ideas that remain proven in simulation alone. In contrast, agenticists do not always see the value in taking the extra effort (which can be significant, even with good simulations) to bring about the transition of the code from simulation to a real robotics platform. Moreover, since many agenticists focus on virtual environments as the target environment (working on virtual agents), the confusion increases.

There are a number of reasons roboticists assign significantly higher value to experiments performed on real platforms. First, even with simulated noise, simulated environments are too clean, too sterile, to really test the limits of a system. For instance, the transition from a simulated sensor to a real sensor is significant and affects the results of experiments and, thus, the validity of the conclusions. Second, simulations limit the scope of human-robot interactions that can be effectively evaluated. Thus much of the work in this exciting area cannot be done in simulation alone. Finally, simulations hide the computational effort: they can run slower than real time, thus obscuring the responsiveness of the systems. The bottom-line claim is that evaluating using robots does not relax the experiment assumptions, while simulations might.

However, simulations are incredibly important not only in shortening the development cycle but also in enabling experiment designs that would otherwise be too expensive to be feasible. For instance, simulations are often used in experiments where the number of robots is scaled up (typically after validating the simulation results on a smaller number of real robots). Simulations are also used in running experiments with a wide variety of physical environments, which in reality would be too expensive to construct. And yes, there are agenticists utilizing the fact that virtual agents are a domain of study in themselves, and thus affect both virtual agents as well as robotics. A good example of this is recent work on the reciprocal velocity obstacles (RVO) family of algorithms, which has been evaluated in both types of environments (van den Berg, Lin, and Manocha 2008; Guy, Lin, and Manocha 2010).

There is a great opportunity for AI researchers to begin exploring essentially philosophical ideas in real-world bodies. The potential impact that we can have on the scientific community, and on the exploding commercial world of robotics, is huge. The alternative—letting roboticists rediscover and reinvent decades of AAMAS work—is unforgivable. Go hug a robot today!

Notes

1. The use of the phrase in this context is due to Manuela Veloso.

2. See CogniTeam, Ltd., www.cogniteam.com.

3. See the CogniTAO (Think as One) software development kit and run-time kernel (CogniTeam, Ltd. 2009), www.cogniteam.com/cognitao.

References

Agmon, N. 2010. On Events in Multirobot Patrol in Adversarial Environments. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems* (AAMAS-2010). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Agmon, N.; Hazon, N.; and Kaminka, G. A. 2008. The Giving Tree: Constructing Trees for Efficient Offline and Online Multirobot Coverage. *Annals of Math and Artificial Intelligence* 52(2–4): 143–168.

Agmon, N.; Kraus, S.; and Kaminka, G. A. 2008. Multi-Robot Perimeter Patrol in Adversarial Settings. In *Proceedings of the IEEE International Conference on Robotics and Automation* (ICRA-08). Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Agmon, N.; Kraus, S.; Kaminka, G. A.; and Sadov, V. 2009. Adversarial Uncertainty in Multirobot Patrol. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence* (IJCAI-09). Menlo Park, CA: AAAI Press.

Agmon, N.; Sadov, V.; Kaminka, G. A.; and Kraus, S. 2008. The Impact of Adversarial Knowledge on Adversarial Planning in Perimeter Patrol. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-08), volume 1, 55–62. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Agmon, N.; Urieli, D.; and Stone, P. 2011. Multiagent Patrol Generalized to Complex Environmental Conditions. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.

Agogino, A., and Tumer, K. 2008. Analyzing and Visualizing Multiagent Rewards in Dynamic and Stochastic Environments. *Journal of Autonomous Agents and Multiagent Systems* 17(2): 320–338.

Argall, B.; Chernova, S.; Veloso, M.; and Browning, B. 2009. A Survey of Robot Learning from Demonstration. *Robotics and Autonomous Systems* 57(5): 469–483.

Balch, T., and Arkin, R. 1998. Behavior-Based Formation Control for Multirobot Teams. *IEEE Transactions on Robotics and Automation* 14(6): 926–939.

Balch, T., and Hybinette, M. 2000. Social Potentials for Scalable Multirobot Formations. In *Proceedings of the IEEE*

International Conference on Robotics and Automation (ICRA-00). Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-Follower Strategies for Robotic Patrolling in Environments with Arbitrary Topologies. In *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-09), 57–64. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Batalin, M., and Sukhatme, G. 2002. Spreading Out: A Local Approach to Multirobot Coverage. In *Proceedings of the 6th International Symposium on Distributed Autonomous Robotic Systems,* 373–382. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Bazzan, A. L. C. 2009. Opportunities for Multiagent Systems and Multiagent Reinforcement Learning in Traffic Control. *Journal of Autonomous Agents and Multiagent Systems* 18(3): 342–375.

Bazzan, A. L. C. 2004. A Distributed Approach for Coordination of Traffic Signal Agents. *Journal of Autonomous Agents and Multiagent Systems* 10(1): 131–164.

Beetz, M. 2001. Structured Reactive Controllers. *Journal of Autonomous Agents and Multiagent Systems*. 4(1–2): 25–55.

Beetz, M.; Arbuckle, T.; Bennewitz, M.; Burgard, W.; Cremers, A.; Fox, D.; Grosskreutz, H.; Hähnel, D.; and Schulz, D. 2001. Integrated Plan-Based Control of Autonomous Service Robots in Human Environments. *IEEE Intelligent Systems* 16(5): 56–65.

Bererton, C. A.; Gordon, G. J.; and Thrun, S. 2003. Auction Mechanism Design for Multirobot Coordination. In *Advances in Neural Information Processing Systems 16*. Upper Saddle River, NJ: Prentice Hall.

Boehn, B. W. 1981. *Software Engineering Economics*. Upper Saddle River, NJ: Prentice Hall.

Busoniu, L.; Babuska, R.; and De Schutter, B. 2008. A Comprehensive Survey of Multiagent Reinforcement Learning. *IEEE Transactions on Systems, Man, and Cybernetics,* Part C: Applications and Reviews 38(2): 156–172.

Butler, Z.; Rizzi, A.; and Hollis, R. L. 2000. Complete Distributed Coverage of Rectilinear Environments. Paper presented at the 4th International Workshop on the Algorithmic Foundations of Robotics. Hanover, NH, 16–18 March.

Carpin, S., and Parker, L. 2002. Cooperative Leader Following in a Distributed Multirobot System. In *Proceedings of the IEEE International Conference on Robotics and Automation*. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Cohen, P. R., and Levesque, H. J. 1991. Teamwork. Nous 35(4): 487–512.

Costelha, H., and Lima, P. U. 2008. Modelling, Analysis and Execution of Multirobot Tasks Using Petri Nets. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-08), 1187–1190. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Delle Fave, F.; Canu, S.; Iocchi, L.; Nardi, D.; and Ziparo, V. 2009. Multi-Objective Multirobot Surveillance. In *Proceedings of the 4th International Conference on Autonomous Robots and Agents,* 68–73. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Desai, J. P. 2002. A Graph Theoretic Approach for Modeling Mobile Robot Team Formations. *Journal of Robotic Systems* 19(11): 511–525.

Desai, J. P.; Ostrowski, J. P.; and Kumar, V. 2001. Modeling and Control of Formations of Nonholonomic Mobile Robots. *IEEE Transactions on Robotics and Automation* 17(6): 905–908.

Dias, M. B., and Stentz, A. T. 2000. A Free Market Architecture for Distributed Control of a Multirobot System. Paper presented at the 6th International Conference on Intelligent Autonomous Systems (IAS-6), Venice, Italy, July 25–27.

Dias, M. B.; Zlot, R.; Kalra, N.; and Stentz, A. 2006. Market-Based Multirobot Coordination: A Survey and Analysis. *Proceedings of the IEEE* 94(7): 1257–1270.

Dresner, K., and Stone, P. 2008. A Multiagent Approach to Autonomous Intersection Management. *Journal of Artificial Intelligence Research* 31: 591–656.

Earl, C., and Firby, R. J. 1997. Combined Execution and Monitoring for Control of Autonomous Agents. In *Proceedings of the First International Conference on Autonomous Agents* (Agents-97), 88–95. New York: Association for Computing Machinery.

Elmaliach, Y., and Kaminka, G. A. 2008. Robust Multirobot Formations Under Human Supervision and Control. *Journal of Physical Agents* 2(1): 31–52.

Elmaliach, Y.; Shiloni, A.; and Kaminka, G. A. 2008. A Realistic Model of Frequency-Based Multirobot Fence Patrolling. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-08), volume 1, 63–70. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Fajardo, D.; Au, T.-C.; Waller, T.; Stone, P.; and Yang, D. 2012. Automated Intersection Control: Performance of a Future Innovation Versus Current Traffic Signal Control. *Transportation Research Record* 2259: 223–32.

Ferranti, E.; Trigoni, N.; and Levene, M. 2007. Brick&mortar: An On-Line Multiagent Exploration Algorithm. In *Proceedings of the IEEE International Conference on Robotics and Automation* (ICRA). Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Fredslund, J., and Mataric, M. J. 2002. A General Algorithm for Robot Formations Using Local Sensing and Minimal Communications. *IEEE Transactions on Robotics and Automation* 18(5): 837–846.

Gat, E. 1992. Integrating Planning and Reacting in a Heterogeneous Asynchronous Architecture for Controlling Real-World Mobile Robots. In *Proceedings of the Tenth National Conference on Artificial Intelligence* (AAAI-92). Menlo Park, Calif.: AAAI Press.

Georgeff, M. P.; Pell, B.; Pollack, M. E.; Tambe, M.; and Wooldridge, M. 1998. The Belief-Desire-Intention Model of Agency. In *Proceedings of Intelligent Agents V: The 5th International Workshop on Agent Theories, Architectures, and Languages* (ATAL-98), 1–10. Berlin: Springer.

Gerkey, B. P., and Mataric, M. J. 2002. Sold!: Auction Methods for Multirobot Coordination. *IEEE Transactions on Robotics and Automation* 18(5): 758–768.

Gmytrasiewicz, P. J., and Durfee, E. H. 2000. Rational Coordination in Multiagent Environments. *Journal of Autonomous Agents and Multiagent Systems* 3(4): 319–350.

Gökçe, F., and Sahin, E. 2009. To Flock or Not to Flock: The Pros and Cons of Flocking in Long-Range "Migration" of Mobile Robot Swarms. In *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-09), 65–72. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Goldberg, D., and Mataríc, M. 2001. Design and Evaluation of Robust Behavior-Based Controllers for Distributed Multirobot Collection Tasks. In *Robot Teams: From Diversity to Polymorphism,* 315–344. Boca Raton, FL: A. K. Peters/CRC Press.

Grosz, B. J., and Kraus, S. 1996. Collaborative Plans for Complex Group Actions. *Artificial Intelligence* 86(2): 269–358.

Grosz, B. J., and Sidner, C. L. 1990. Plans for Discourse. In *Intentions in Communication*, ed. P. R. Cohen, J. Morgan, and M. Pollack. Cambridge, MA: The MIT Press. 417–445.

Guy, S. J.; Lin, M. C.; and Manocha, D. 2010. Modeling Collision Avoidance Behavior for Virtual Humans. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-10), 575–582. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Haigh, K. Z., and Veloso, M. M. 1997. High-Level Planning and Low-Level Execution: Towards a Complete Robotic Agent. In *Proceedings of the First International Conference on Autonomous Agents* (Agents-97), 363–370. New York: Association for Computing Machinery.

Haigh, K. Z., and Veloso, M. M. 1998. Planning, Execution and Learning in a Robotic Agent. In *Proceedings of the International Conference on Artificial Intelligence Planning Systems,* 120–127. Menlo Park, CA: AAAI Press.

Hazon, N., and Kaminka, G. 2008. On Redundancy, Efficiency, and Robustness in Coverage for Multiple Robots. *Robotics and Autonomous Systems* 56(12): 1102–1114.

Hoen, P. J.; Tuyls, K.; Panait, L.; Luke, S.; and Poutré, J. A. L. 2006. An Overview of Cooperative and Competitive Multiagent Learning. In *First International Workshop on Learning and Adaption in Multiagent Systems,* volume 3898 of Lecture Notes in Computer Science, ed. K. Tuyls, P. Hoen, K. Verbeeck, and S. Sen, 1–46. Berlin: Springer.

Huber, M. J., and Durfee, E. H. 1995. Deciding When to Commit to Action During Observation-Based Coordination. In *Proceedings of the First International Conference on Multiagent Systems* (ICMAS-95), 163–170. Menlo Park, CA: AAAI Press.

Huber, M. J., and Durfee, E. H. 1996. An Initial Assessment of Plan-Recognition-Based Coordination for Multiagent Teams. In *Proceedings of the Second International Conference on Multiagent Systems* (ICMAS-96). Menlo Park, CA: AAAI Press.

Inalhan, G.; Busse, F.; and How, J. 2000. Precise Formation Flying Control of Multiple Spacecraft Using Carrier-Phase Differential GPS. Paper presented at the AAS/AIAA Space Flight Mechanics Meeting, Clearwater, FL, 23–26 January.

Ingrand, F. F.; Georgeff, M. P.; and Rao, A. S. 1992. *An Architecture for Real-Time Reasoning and System Control.* IEEE Expert 7(6): 34–44.

Jager, M., and Nebel, B. 2002. Dynamic Decentralized Area Partitioning for Cooperating Cleaning Robots. In *Proceedings of the IEEE International Conference on Robotics and Automation* (ICRA), 3577–3582. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Jennings, N. R. 1995. Controlling Cooperative Problem Solving in Industrial Multiagent Systems Using Joint Intentions. *Artificial Intelligence* 75(2): 195–240.

Jensen, E.; Franklin, M.; Lahr, S.; and Gini, M. 2011. Sustainable Multirobot Patrol of an Open Polyline. In *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA-11), 4792–4797. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Kaminka, G. A., and Frenkel, I. 2005. Flexible Teamwork in Behavior-Based Robots. In *Proceedings of the Twentieth National Conference on Artificial Intelligence* (AAAI-05). Menlo Park, CA: AAAI Press.

Kaminka, G. A., and Frenkel, I. 2007. Integration of Coordination Mechanisms in the BITE Multirobot Architecture. In *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA-07). Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Kaminka, G. A., and Tambe, M. 2000. Robust Multiagent Teams Via Socially-Attentive Monitoring. *Journal of Artificial Intelligence* Research 12: 105–147.

Kaminka, G. A.; Erusalimchik, D.; and Kraus, S. 2010. Adaptive Multirobot Coordination: A Game-Theoretic Perspective. In *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA-10). Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Kaminka, G. A.; Schechter-Glick, R.; and Sadov, V. 2008. Using Sensor Morphology for Multirobot Formations. *IEEE Transactions on Robotics* 24(2): 271–282.

Kaminka, G. A.; Yakir, A.; Erusalimchik, D.; and Cohen-Nov, N. 2007. Towards Collaborative Task and Team Maintenance. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-07). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Köse, H.; Meriçli, C.; Kaplan, K.; and Akin, H. 2003. All Bids for One and One Does for All: Market-Driven Multiagent Collaboration in Robot Soccer Domain. In *Computer and Information Sciences* (ISCIS 2003), volume 2869 of Lecture Notes in Computer Science, ed. A. Yazici and C. Sener, 529–536. Berlin: Springer Berlin.

Lagoudakis, M. G.; Berhault, M.; Koenig, S.; Keskinocak, P.; and Kleywegt, A. J. 2004. Simple Auctions with Performance Guarantees for Multirobot Task Allocation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS-04), volume 1, 698–705. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Lagoudakis, M. G.; Markakis, E.; Kempe, D.; Keskinocak, P.; Kleywegt, A. J.; Koenig, S.; Tovey, C. A.; Meyerson, A.; and Jain, S. 2005. Auction-Based Multirobot Routing. In *Robotics: Science and Systems,* 343–350. Cambridge, MA: The MIT Press.

Lee, J.; Huber, M. J.; Kenny, P. G.; and Durfee, E. H. 1994. UM-PRS: An Implementation of the Procedural Reasoning System for Multirobot Applications. In *Proceedings of the Conference on Intelligent Robotics in Field, Factory, Service, and Space* (CIRFFSS-94), 842–849. Washington, DC: National Aeronautics and Space Administration.

Levesque, H. J.; Cohen, P. R.; and Nunes, J. H. T. 1990. On Acting Together. In *Proceedings of the Eigth National Conference on Artificial Intelligence* (AAAI-90). Menlo-Park, CA: AAAI Press.

Lin, L., and Zheng, Z. 2005. Combinatorial Bids Based Multirobot Task Allocation Method. In *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA-05), 1145–1150. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Lötzsch, M.; Risler, M.; and Jüngel, M. 2006. XABSL—A Pragmatic Approach to Behavior Engineering. In *Proceedings of IEEE/RSJ International Conference of Intelligent Robots and Systems* (IROS), 5124–5129. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Marino, A.; Parker, L. E.; Antonelli, G.; and Caccavale, F. 2009. Behavioral Control for Multirobot Perimeter Patrol: A Finite State Automata Approach. In *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA-09). Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Mermoud, G.; Brugger, J.; and Martinoli, A. 2009. Towards Multilevel Modeling of Self-Assembling Intelligent Micro-Systems. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems* AAMAS (1), 89–96. Richland, SC:

International Foundation for Autonomous Agents and Multiagent Systems.

Michael, N.; Zavlanos, M. M.; Kumar, V.; and Pappas, G. J. 2008. Distributed Multirobot Task Assignment and Formation Control. In *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA-08), 128–133. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Murphy, R. R.; Tadokoro, S.; Nardi, D.; Jacoff, A.; Fiorini, P.; Choset, H.; and Erkmen, A. M. 2008. Search and Rescue Robotics. In *Springer Handbook of Robotics,* ed. B. Siciliano and K. Oussama, 1151–1173. Berlin: Springer Verlag.

Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Nguyen-Tuong, D., and Peters, J. 2011. Model Learning for Robot Control: A Survey. *Cognitive Processing* 12(4): 319–340.

Ostergaard, E.; Sukhatme, G.; and Mataríc, M. 2001. Emergent Bucket Brigading. In *Proceedings of the Fifth International Conference on Autonomous Agents* (Agents-01), 29–30. New York: Association for Computing Machinery.

Padgham, L., and Winikoff, M. 2002. Prometheus: A Methodology for Developing Intelligent Agents. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-02), 37–38. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Padgham, L., and Winikoff, M. 2004. *Developing Intelligent Agent Systems: A Practical Guide*. Wiley Series in Agent Technology. New York: Wiley.

Panait, L., and Luke, S. 2005. Cooperative Multiagent Learning: The State of the Art. *Autonomous Agents and Multiagent Systems* 11(3): 387–434.

Parker, L. E. 1998. ALLIANCE: An Architecture for Fault Tolerant Multirobot Cooperation. *IEEE Transactions on Robotics and Automation* 14(2): 220–240.

Pell, B.; Bernard, D. E.; Chien, S. A.; Gat, E.; Muscettola, N.; Nayak, P. P.; Wagner, M. D.; and Williams, B. C. 1997. An Autonomous Spacecarft Agent Prototype. In *Proceedings of the First International Conference on Autonomous Agents* (Agents-97), 253–261. New York: Association for Computing Machinery.

Pollack, M. E. 1990. Plans as Complex Mental Attitudes. In *Intentions in Communication,* ed. P. R. Cohen, J. Morgan, and M. Pollack. Cambridge, MA: The MIT Press.

Pollack, M. E., and Horty, J. F. 1999. There's More to Life Than Making Plans: Plan Management In Dynamic, Multiagent Environments. *AI Magazine* 20(4): 71–84.

Procaccia, A. D.; Rosenschein, J. S.; and Kaminka, G. A. 2007. On the Robustness of Preference Aggregation in Noisy Environments. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-07). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Py, F.; Rajan, K.; and McGann, C. 2010. A Systematic Agent Framework for Situated Autonomous Systems. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-10). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Rao, A. S., and Georgeff, M. P. 1995. BDI Agents: From Theory to Practice. In *Proceedings of the First International Conference on Multiagent Systems* (ICMAS-95), 312–319. Menlo Park, CA: AAAI Press.

Rekleitis, I.; Dudek, G.; and Milios, E. 2001. Multi-Robot Collaboration for Robust Exploration. *Annals of Mathematics and Artificial Intelligence* 31(1–4): 7–40.

Rekleitis, I.; Lee-Shue, V.; New, A. P.; and Choset, H. 2004. Limited Communication, Multirobot Team Based Coverage. In *IEEE International Conference on Robotics and Automation,* 3462–3468. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Risler, M., and von Stryk, O. 2008. Formal Behavior Specification of Multirobot Systems Using Hierarchical State Machines in XABSL. Paper presented at the AAMAS-08 Workshop on Formal Models and Methods for Multi-Robot Systems. Estoril, Portugal, May 13.

Rosenfeld, A.; Kaminka, G. A.; Kraus, S.; and Shehory, O. 2008. A Study of Mechanisms for Improving Robotic Group Performance. *Artificial Intelligence* 172(6–7): 633–655.

Rybski, P. E.; Gini, M. L.; Hougen, D. F.; Stoeter, S.; and Papanikolopoulos, N. 2002. A Distributed Surveillance Task Using Miniature Robots. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-02), 1393–1394. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Rybski, P.; Larson, A.; Lindahl, M.; and Gini, M. 1998. Performance Evaluation of Multiple Robots in a Search and Retrieval Task. In *Proceedings of the Artificial Intelligence and Manufacturing Research Planning Workshop,* 153–160. Menlo Park, CA: AAAI Press.

Rybski, P. E.; Stoeter, S.; Erickson, M. D.; Gini, M. L.; Hougen, D. F.; and Papanikolopoulos, N. 2000. A Team of Robotic Agents for Surveillance. In *Proceedings of the Fourth International Conference on Autonomous Agents* (Agents-00), 9–16. New York: Association for Computing Machinery.

Sardiña, S., and Padgham, L. 2011. A BDI Agent Programming Language with Failure Handling, Declarative Goals, and Planning. *Journal of Autonomous Agents and Multiagent Systems* 23(1): 18–70.

Scerri, P.; Johnson, L.; Pynadath, D.; Rosenbloom, P.; Si, M.; Schurr, N.; and Tambe, M. 2003. A Prototype Infrastructure for Distributed Robot-Agent-Person Teams. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-03). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Schneider-Fontan, M., and Mataríc, M. 1996. A Study of Territoriality: The Role of Critical Mass in Adaptive Task Division. In *From Animals to Animats IV*, ed. P. Maes, M. Mataríc, J.-A. Meyer, J. Pollack, and S. Wilson, 553–561. Cambridge, MA: The MIT Press.

Shiloni, A.; Agmon, N.; and Kaminka, G. A. 2009. Of Robot Ants and Elephants. In *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-09). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Simmons, R. G.; Goodwin, R.; Haigh, K. Z.; Koenig, S.; O'Sullivan, J.; and Veloso, M. M. 1997a. Xavier: Experience with a Layered Robot Architecture. *ACM Sigart Bulletin* 8(1–4): 22–33.

Simmons, R. G.; Goodwin, R.; Haigh, K. Z.; Koenig, S.; and O'Sullivan, J. 1997b. A Layered Architecture for Office Delivery Robots. In *Proceedings of the First International Conference on Autonomous Agents* (Agents-97), 245–252. New York: Association for Computing Machinery.

Smith, S. L.; Schwager, M.; and Rus, D. 2011. Persistent Monitoring of Changing Environments Using a Robot with Limited Range Sensing. In *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA-11), 5448–5455. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Spaan, M. T. J.; Gonçalves, N.; and Sequeira, J. 2010. Multirobot Coordination by Auctioning POMDPs. In *Proceedings of the IEEE International Conference on Robotics and Automation,* 1446–1451. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Stone, P., and Veloso, M. 2000. Multiagent systems: A Survey from a Machine Learning Perspective. *Autonomous Robots* 8(3): 345–83.

Tabuada, P.; Pappas, G. J.; and Lima, P. U.

2005. Motion Feasibility of Multiagent Formations. *IEEE Transactions on Robotics and Automation* 21(3): 387–392.

Tambe, M. 1997. Towards Flexible Teamwork. *Journal of Artificial Intelligence Research* 7: 83–124.

Tambe, M.; Adibi, J.; Al-Onaizan, Y.; Erdem, A.; Kaminka, G. A.; Marsella, S. C.; and Muslea, I. 1999. Building Agent Teams Using an Explicit Teamwork Model and Learning. *Artificial Intelligence* 111(1): 215–239.

Tang, F., and Parker, L. E. 2007. A Complete Methodology for Generating Multirobot Task Solutions Using Asymtre-D and Market-Based Task Allocation. In *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA-07), 3351–3358. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Thrun, S.; Beetz, M.; Bennewitz, M.; Burgard, W.; Cremers, A. B.; Dellaert, F.; Fox, D.; Hähnel, D.; Rosenberg, C. R.; Roy, N.; Schulte, J.; and Schulz, D. 2000. Probabilistic Algorithms and the Interactive Museum Tour-Guide Robot Minerva. *International Journal of Robotics Research* 19(11): 972–999.

Thrun, S.; Burgard, W.; and Fox, D. 2005. *Probabilistic Robotics*. Cambridge, MA: The MIT Press.

Tousignant, S.; Wyk, E. V.; and Gini, M. 2011. An Overview of Xrobots: A Hierarchical State Machine Based Language. Paper presented at the ICRA-2011 Workshop on Software Development and Integration in Robotics, Shanghai, China May 9–13.

Traub, M. 2011. Topics in Multirobot Teamwork. Master's thesis, Computer Science Department, Bar-Ilan University, Ramat Gan, Israel.

Traub, M.; Kaminka, G. A.; and Agmon, N. 2011. Who Goes There? Using Social Regret to Select a Robot to Reach a Goal. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-11). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Turgut, A. E.; Çelikkanat, H.; Gokõke, F.; and Sahin, E. 2008. Self-Organized Flocking with a Mobile Robot Swarm. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-08), 39–46. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

van den Berg, J.; Lin, M. C.; and Manocha, D. 2008. Reciprocal Velocity Obstacles for Real-Time Multiagent Navigation. In *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA-08), 1928–1935. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Vig, L., and Adams, J. A. 2006. Market-Based Multirobot Coalition Formation. In *Distributed Autonomous Robotic Systems 7,* ed. M. Gini, and R. Voyles, 227–236. Berlin: Springer.

Wagner, I. A., and Bruckstein, A. M. 1997. Cooperative cleaners—A Study in Ant Robotics. In *Communications, Computation, Control, and Signal Processing: A Tribute to Thomas Kailath,* 289–308. Dordrecht, Netherlands: Kluwer Academic Publishers.

Williams, K., and Burdick, J. 2006. Multi-Robot Boundary Coverage with Plan Revision. In *Proceedings of the IEEE International Conference on Robotics and Automation* (ICRA). Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Wurman, P. R.; D'Andrea, R.; and Mountz, M. 2008. Coordinating Hundreds of Cooperative, Autonomous Vehicles in Warehouses. *AI Magazine* 29(1): 9–20.

Xu, Y.; Scerri, P.; Sycara, K. P.; and Lewis, M. 2006. Comparing Market and Token-Based Coordination. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-06), 1113–1115. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Yamins, D., and Nagpal, R. 2008. Automated Global-to-Local Programming in 1-D Spatial Multiagent Systems. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-08), 615–622. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Yang, E., and Gu, D. 2004. Multiagent Reinforcement Learning for Multirobot Systems: A Survey. Technical Report CSM-404, University of Essex, Colchester, UK.

Yu, C.-H., and Nagpal, R. 2008. Sensing-Based Shape Formation on Modular Multirobot Systems: A Theoretical Study. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-08), 71–78. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Yu, C.-H.; Werfel, J.; and Nagpal, R. 2010. Collective Decision-Making in Multiagent Systems by Implicit Leadership. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS-10), 1189–1196. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Zheng, X.; Jain, S.; Koenig, S.; and Kempe, D. 2005. Multi-Robot Forest Coverage. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS). Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Ziparo, V. A.; Iocchi, L.; Lima, P. U.; Nardi, D.; and Palamara, P. F. 2010. Petri Net Plans. *Autonomous Agents and Multiagent Systems* 23(3): 344–383.

Zlot, R., and Stentz, A. 2006. Market-Based Multirobot Coordination for Complex Tasks. *The International Journal of Robotics Research* 25(1): 73–101.

Zuluaga, M., and Vaughan, R. 2005. Reducing Spatial Interference in Robot Teams by Local-Investment Aggression. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems.* Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Gal A. Kaminka is an associate professor at the Computer Science Department and the Brain Sciences Research Center at Bar-Ilan University (Israel). His research expertise includes multiagent and multirobot systems, teamwork and coordination, behavior and plan recognition, and modeling social behavior. He received his Ph.D. from the University of Southern California and spent two years as a postdoctorate fellow at Carnegie Mellon University. Today, Kaminka leads the MAVERICK research group at Bar-Ilan, supervising more than a dozen M.Sc. and Ph.D. students. He was awarded an IBM faculty award and top places at international robotics competitions. He served as the program cochair of the 2005 European workshop on Multiagent Systems (EUMAS), the program chair of the 2008 Israeli Conference on Robotics, and the program cochair of the 2010 International Joint Conference on Autonomous Agents and Multiagent Systems. He is currently serving on the executive board of IFAAMAS (International Foundation of Autonomous Agents and Multiagent Systems) and the RoboCup executive board.