Article

The International Journal of Robotics Research
2019, Vol. 38(12-13) 1329–1337

DOI: 10.1177/0278364919839137
journals.sagepub.com/home/ijr
SAGE

DART: Diversity-enhanced Autonomy in Robot Teams

Nora Ayanian

Abstract
This paper defines the research area of Diversity-enhanced Autonomy in Robot Teams (DART), a novel paradigm for the creation and design of policies for multi-robot coordination. Although current approaches to multi-robot coordination have been successful in structured, well-understood environments, they have not been successful in unstructured, uncertain environments, such as disaster response. Although robot hardware has advanced significantly in the past decade, the way we solve multi-robot problems has not. Even with significant advances in the field of multi-robot systems, the same problem-solving paradigm has remained: assumptions are made to simplify the problem, and a solution is optimized for those assumptions and deployed to the entire team. This results in brittle solutions that prove incapable if the original assumptions are invalidated. This paper introduces a new multi-robot problem-solving paradigm which uses a diverse set of control policies that work together synergistically within the same team of robots. Such an approach will make multi-robot systems more robust in unstructured and uncertain environments, such as in disaster response, environmental monitoring, and military applications, and allow multi-robot systems to extend beyond the highly structured and highly controlled environments where they are successful today.

Keywords
Multi-robot systems, multi-agent systems

1. Introduction

The field of multi-robot systems (MRS) is growing at a rapid pace. Research in MRS spans many different areas, including automated delivery (Agha-mohammadi et al., 2014; Lonsdorf, 2017; Sung et al., 2013), surveillance (Glaser, 2017), and disaster response (Jennings et al., 1997; Schurr et al., 2005). There have also been many successful demonstrations of increasing numbers of robots (Chung et al., 2016; Glaser, 2016; Hauert et al., 2011; Kushleyev et al., 2013; Preiss et al., 2017; Rubenstein et al., 2012). MRS have also been successfully deployed in the field including in warehousing (D'Andrea and Wurman, 2008), manufacturing (Hagerty, 2015), and entertainment (Barret, 2016). Whereas these outcomes show the promise of MRS, the environments in which MRS have been successful are highly controlled, and some are highly instrumented, enabling precise tuning of controllers and nearly perfect knowledge of environmental conditions.

Many environments where MRS could be beneficial are not highly controlled or equipped with the extensive infrastructure often necessary to coordinate large teams of robots with state-of-the-art algorithms. For example, containing wildfires, searching through collapsed buildings, patrolling borders, monitoring infrastructure, and containing oil spills all occur in highly dynamic and unique environments (no two collapsed buildings are the same), with high uncertainty and little control over other non-robot agents in the environment. One of the most desirable benefits of MRS is their *robustness*, wherein robots can compensate for loss of capabilities by relying on other robots in the team. However, the uncertainty of many real-world environments renders current state-of-the-art algorithms and controllers, even those designed for robustness, ineffectual. Although robot hardware has advanced significantly in the past decade, the way we solve multi-robot problems has not. Many control policies are so specialized and optimized for specific capabilities and conditions that they do not empower robots to cope with uncertainty. Incorporating diversity, in the form of an ensemble of control policies that

University of Southern California, USA

Corresponding author:
Nora Ayanian, University of Southern California, Viterbi School of Engineering, Department of Computer Science, Los Angeles, CA 90089, USA.
Email: ayanian@usc.edu

work synergistically across the team, can help to realize the true benefits of robustness in the face of uncertainty for teams of robots. Single-robot systems can also benefit from ensembles of control policies, from which they can draw when faced with control failure.

2. Motivation

In disaster response alone, the potential impact of autonomous MRS is substantial: 60,000 people die each year in natural disasters, mostly in developing countries (Kenny, 2009). This makes robots an ideal tool for disaster response. In fact, DJI announced that *one* properly equipped drone can find a missing person more than *five times faster* than traditional search methods (DJI, 2016). However, most robots used in search and rescue today are teleoperated (Liu and Nejat, 2013), requiring trained operators that may not be nearby. Autonomous robots equipped for disaster response that can *automatically* synthesize control policies without the need for an expert operator can reduce response time and save many more lives, especially when a trained operator may be hours away.

The potential applications of autonomous MRS go well beyond disaster response, including military, agriculture, transportation, manufacturing, and fulfillment applications. However, current solutions for MRS have not successfully transitioned from controlled environments such as laboratories or warehouse facilities to the inherently high uncertainty in these complex environments. Without infrastructure that provides communication and localization, and without knowledge of or control over the environment, current state-of-the-art methods fail.

Although the field of MRS has advanced significantly, the same problem-solving paradigm has remained. First, the problem is defined. Next, complexity is reduced by making several assumptions to simplify the problem, such as terrain and communication range. Finally, an optimal solution to that specific problem is designed and applied to all the robots in the team. This paradigm, pictured in Figure 1(left), limits the capability of MRS to cope with real-world environments. The solutions are brittle, as the assumptions made are easily invalidated and the optimized controller is not designed for real environments. In the best case, the controller is able to overcome these challenges, but it is not the best solution to the problem, defeating the purpose of optimization. In the worst case, the controller cannot cope, potentially causing mission failure, loss of high-value assets, and casualties; after all, if the environment violates the assumptions and the same ill-equipped controller is applied to all robots, it is possible that all of them will fail.

2.1. Leveraging diversity

Instead of applying the same controller to all robots, a new approach leveraging diversity in policies within the robot team can allow MRS to better cope with uncertain environments. Using an ensemble of diverse control policies to accomplish a coordinated task within a single team of robots can enable the team to adjust to different conditions. For example, with two unmanned aerial vehicles (UAVs) on a large security task, a natural result of using an ensemble of controllers is for one UAV to position itself high, to view the entire area, while the other UAV takes a closer look at areas of interest. Diversity allows the robots to perceive and respond to failure as they encounter it in the environment.

2.1.1. Diversity in human workgroups. Diversity is well established as a way to improve the performance of human workgroups: studies have shown repeatedly that diverse groups outperform homogeneous groups (Hoffman and Maier, 1961; Hoffman, 1978; Jackson, 1992; Nemeth, 1986). Whereas diverse groups do have a higher likelihood of conflict (Ancona and Caldwell, 1992; O'Reilly and Flatt, 1989; Steiner, 1972), that conflict can be productive. In studies where conflict due to diverse skill sets was purposely introduced into the workgroup, it was shown to consistently lead to higher-quality solutions (de Wit et al., 2012; Jehn, 1995; van de Vliert and de Dreu, 1994).

2.1.2. Diversity in insect and animal behavior. Heterogeneity has also been studied extensively in insect and animal behavior. Jandt et al. (2014) studied personality at multiple levels with regard to behavioral syndromes and insect societies, discussing fitness consequences of intra-colony behavioral variation. Specifically, under varying environmental conditions, maintaining a mixture of individuals with different behavioral types may be more effective than individuals switching between behavioral types, which might be costly and inefficient. Slower, more accurate individuals can bring large quantities of food back to the colony when good abundance is constant, whereas faster "sloppier" individuals might be more efficient at exploiting resources in more frequently changing environments (Chittka et al., 2009). Burns and Dyer (2008) found that ant colonies that maintain a mixture of different foraging types within a group allows the colonies to respond more quickly to environmental fluctuation. In certain species, groups with a mixture of aggression types tend to have higher fitness than groups with only one type (Modlmeier and Foitzik, 2011; Pruitt and Riechert, 2011). On the other hand, maintaining a mixture of *inflexible* behavioral types can incur costs to the colony, such as overly aggressive types being aggressive to their own nestmates (Crosland, 1990).

These results in insect and animal behavior studies point strongly to behaviorally heterogeneous teams with the ability to adapt to the environment and task having higher fitness in uncertain and dynamic environments, which has inspired many multi-robot approaches. However, there is a need to further study the use of diversity as a tool for MRS,

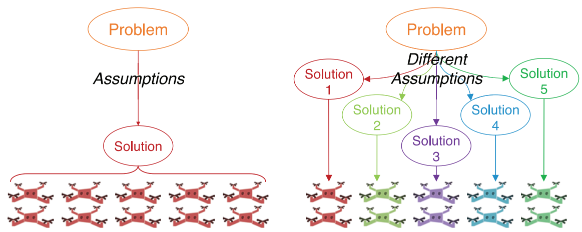

Fig. 1. (Left) The current MRS problem-solving paradigm is linear, making the same assumptions and deploying the same solution to all robots in the team. (Right) The proposed novel paradigm takes advantage of diversity in controllers to handle various scenarios. Varying assumptions are made, and a complementary ensemble of control policies is deployed across the team.

especially in tightly coordinated tasks, as well as single-robot systems operating in uncertain environments.

2.2. A new paradigm

Thus, *the current problem-solving paradigm in MRS may not reflect an effective approach to working in teams*. In the current paradigm, one set of assumptions is made, a single control policy is developed, and it is uniformly deployed to all robots in the team as in Figure 1(left). Instead, an ensemble of control policies may be able to leverage the strengths of the different control policies under different conditions much as in human workgroups: varying control policies should be developed using different sets of assumptions and/or different styles of interaction, and the best approaches combined synergistically within the team, as in Figure 1(right). In this way, MRS may leverage diversity much as human workgroups do, to improve robustness in uncertain environments. Single-robot systems may also benefit from ensembles of control policies, as they would allow the system to swap policies in the case of control failure.

3. Current state of the art

Diversity of robots with different physical embodiments or capabilities has been studied previously (Huang et al., 2006; Parker, 1998; Pimenta et al., 2008; Prorok et al., 2017), but has not led to significant improvements in robustness of MRS. However, controllers for those teams are developed using a similar paradigm, making the same assumptions across the entire team. There has been relatively little exploration into diversity in control policies within a single team of robots. Most research in this area is a result of studying ants that take different roles in foraging and house hunting (Berman et al., 2007; Dorigo et al., 2006; Sugawara et al., 2004) or collective transport (Kumar et al., 2013), and applied to similar problems in robotics. Unfortunately, in trying to model ant algorithms closely, these works do not take advantage of robot capabilities, including communication, sensing, and computation, which could expand the solution space and lead to better, more appropriate solutions for a larger set of problems.

Behavioral or control diversity in teams of robots has also been explored. In Tang and Parker's ASyMTRe architecture based on schemas, robots take different roles depending on environmental conditions (Parker and Tang, 2006; Tang and Parker, 2007), but the robots are all programmed to react the same. This leaves them vulnerable to unforeseen changes in capabilities or the environment, and does not enable robots to individually adjust their approaches.

A majority of work exploring control diversity in robots exists in behavior-based systems, most notably Balch's work in learning behavioral specialization for robot teams (Balch, 1997, 2000). Goldberg and Matarić (1997) evaluated multi-robot controllers based on the amount of interference and describe caste arbitration, where all robots have the same capabilities, but have different conditions for activating behaviors. Schneider-Fontan and Matarić (1998) concluded that adapting group behavior is a balance between minimizing interference and maximizing synergy, and interference is the key stumbling block in the way of efficient group interactions.

More recently, evolutionary robotics and agent-based systems have been appearing as a method for encouraging behavioral diversity and plasticity (individuals changing roles over time). Mouret and Doncieux (2011) reviewed and benchmarked published approaches to behavioral diversity, and showed that fostering behavioral diversity substantially improves the evolutionary process in the investigated experiments, regardless of task. Pugh et al. (2016) reviewed quality diversity algorithms, which have resulted in a new class of algorithms that return an archive of diverse, high-quality behaviors in a single run. Vassiliades and Christodoulou (2016) designed behaviorally plastic agents (capable of switching between different behaviors in response to environmental changes), by taking inspiration from neuroscience, using artificial neural networks, neuromodulation, and synaptic gating. Umedachi et al. (2015) attempted to understand the underlying mechanism of the behavioral diversity of animals, then used the findings to build truly adaptive robots. However, all of these approaches focused on training agents to act independently in the environment, and thus are not directly applicable to multi-robot problems where task completion relies on tight coordination, such as box-pushing, shape formation, wildfire containment, cooperative transport, etc. Furthermore, agents are trained in the environments where they will be used, which, in natural disasters in particular, may not be possible.

There is also a significant body of work on large-scale simulation of crowds, that generate realistic-appearing simulated crowds (Bera et al., 2016; Guy et al., 2010; Lee et al., 2007; Pelechano et al., 2007). Although these approaches generate diverse behaviors, diversity in crowd simulation plays a different role than the one we seek.

Whereas we seek diversity to enhance performance, diversity in crowd simulation is cosmetic.

3.1. Diversity improving performance

Diversity has already been shown to improve performance in some scenarios. In multi-agent systems, Marcolino et al. (2013) addressed the problem of selecting the best possible team to accomplish a goal given limited resources. Varied agents form a team and vote on the best course of action in the computer game Go. They compared teams with the strongest individual members and teams with diverse members, and found that a team of diverse agents can outperform a uniform team of the strongest agents when individual agents outperform the overall strongest agent in certain states. It is important to keep in mind, however, that their diverse agents combine to make a single decision, whereas in MRS, each agent encounters different scenarios and must make their own decision. Nonetheless, in scenarios where the team of robots encounters similar challenges, they may be able to share useful information about which actions were successful and unsuccessful with certain parameters. If this can propagate throughout the team, the entire team may be more informed and thus perform better.

In single-robot systems, diversity can be provided by switching controllers. For example, Zefran and Burdick (1998) designed stable control schemes for systems with changing dynamics; in their case, a different controller is activated in each dynamic regime of the system. However, to successfully design such a system, one must fully understand the dynamics of the system and how to stabilize the system in each regime. For MRS in uncertain environments, it may not be possible to fully understand the dynamics of each robot under all possible environmental conditions, thus it may be difficult to know how to stabilize the system in advance. Equipping the team with varied control policies or a set of parameters enables performance observation of each control policy or parameter set *in parallel*, which can then quickly inform individuals of how to stabilize themselves.

Lyu et al. (2016) explored k-survivability in MRS. The k-survivability of n paths is the probability that “at least k out of n robots following those paths through a stochastic threat environment reach goals.” The main idea is that if the best and safest path is known, then it is still not robust for all robots to take that path, because all the robots on that path can fall to a single trap. k-survivability demonstrates that in uncertain environments, diversity can be leveraged to improve the performance, and in this case the survivability, of a team of robots.

In swarm robotics, Li et al. (2004) studied the effect of diversity and specialization in self-organizing, distributed, artificial systems, correlating the degree of specialization with the swarm’s overall performance. As not all diversity will lead to better performance, they defined specialization as diversity that is evoked for better performance. They studied the stick-pulling problem, where robots search an arena pulling sticks out of the ground; each stick requires two (or more, in the generalized case) robots to remove the stick from the ground, and each robot has a gripping time parameter, which can be homogeneous or heterogeneous across the team (Ijspeert et al., 2001; Martinoli et al., 2004; Martinoli and Mondada, 1997). When the number of robots was smaller than the number of sticks, this problem could lead to a deadlock situation where all robots wait for assistance for extended periods of time. In this scenario, heterogeneous swarms, in which agents had in parallel learned their own gripping time parameters, outperformed both swarms with hardwired homogeneous and learned homogeneous gripping time parameters, owing to the specialization of the agents.

4. Some open problems in DART

Much as human workgroups, as well as insects and animals, benefit from diversity in the composition of the group, such variation of behavior would be beneficial for teams of robots operating in uncertain and unstructured environments. Thus, it is natural to consider enhancing multi-robot autonomy with diverse control policies designed to work synergistically together. There exist many open problems in DART; some of the challenging open problems that must be addressed by the community are described here.

4.1. Learning for MRS

In order to learn from humans, or to learn directly from simulations, new machine learning tools must be developed for many-agent systems. Multi-agent learning is an area that is not yet well represented in the literature, save for several works (Foerster et al., 2016; Lowe et al., 2017; Matignon et al., 2012; Peshkin et al., 2000), most of which cannot handle more than a few agents. Other works focus on tasks that can be learned and completed alone (Matarić, 1994; Recchia et al., 2013). Tight coordination between a large team of agents, for example in wildfire containment, currently presents a significant computational challenge for existing multi-agent learning tools. Those that are suitable for tightly coordinated tasks for a few (two or three) agents are intractable for tasks that require tight coordination among large numbers of agents (Amato et al., 2015, 2014; Buffet et al., 2007). *Given tightly coordinated multi-robot tasks, such as pattern formation, border patrol, or wildfire containment, automatically learn sets of control policies for individual robots that enable a team of robots, equipped with one or more of those control policies, to complete those tasks.*

4.1.1. Multi-robot learning from humans. Humans provide a pool of diverse resources that can be tapped to develop diverse controllers that work well together. However, owing to differences in human and robot capabilities (communication, locomotion, sensing, etc.), it is difficult to learn controllers by observing human in-person interaction. By limiting interaction to an interface (such as a mobile phone, tablet, or laptop), communication, locomotion, and sensing can be restricted to robot-like capabilities (Tavakoli et al., 2016). A major benefit of human-inspired controllers is the ability to communicate with and easily motivate study participants, as opposed to animal-inspired controllers. However, learning from human cooperation requires multi-agent learning tools for many agents. Although there is recent interest in learning from demonstration for MRS (Chernova and Veloso, 2010; Freelan et al., 2015; Martins and Demiris, 2010), these works either require a significant amount of domain knowledge, leading to potential bias in the creation of design policies, or do not apply to tightly coordinated tasks. New approaches to learning are necessary in order to learn truly novel behaviors from demonstration. *Given data generated by humans completing a coordinated task, find a set of control policies that, when deployed on a team of agents, produce qualitatively and quantitatively similar results.* In order to produce functional behaviors, this requires automatic solutions to the correspondence problem, which would eliminate designer bias and pave the way for novel behaviors; the agents to understand both the state of the human as well as the goal of the behavior; and measures of behavioral similarity between humans and agents.

4.1.2. Automatic correspondence. In learning from humans or other animals, it is necessary to solve the mapping between demonstrators and imitators automatically. Using primitive behaviors that are hand-coded, as has traditionally been done, limits the behavior space of the team of robots to those determined by the system designer, which can lead to a biased set of behaviors. *Given demonstrative data generated by a human or animal completing a coordinated task, find automatically, without hand-coding, a set of robot behaviors that produce qualitatively similar results to the demonstrator.*

4.2. Automatic abstractions for complex observations

With large numbers of agents interacting in a space, and a very rich set of possible observations or features (potentially millions or more), it is necessary to use abstractions for the observation or feature space of each agent. For example, variations of regular or polar grids can be used with occupancies or sampling can be utilized. Abstractions can simplify the decision-making process, but their utility depends on encoding relevant information, which can be task-specific. Thus, developing useful abstractions, whether automatic or hand-coded, is imperative; if done automatically, one can avoid the use and potential bias of expert knowledge. *Given a large number of agents interacting in a space, find efficient representations for the agents' observable state that enable decision making in individual agents. With respect to learning from demonstration, find efficient representations of the observable state of the demonstrator that enable computer-controlled agents to mimic the behaviors of the demonstrator.* Efficiency can be measured by how much information is encoded in the representation, and how quickly it can be accessed and shared if necessary.

4.3. Measures of behavioral similarity

It is necessary to measure quantitatively how similar two behaviors are, whether comparing between two computer-generated behaviors or a demonstrator–imitator pair. *Given two behavioral policies, develop metrics that measure their similarity.*

4.4. Measures of diversity

Taking inspiration from the study of behavioral diversity in social insect colonies, there is a need to understand the impact of behavioral diversity on MRS in tightly coordinated tasks, and, if learning from demonstration, to measure the diversity in the demonstrators' behaviors and the resulting imitators' behaviors. This is distinct from measuring behavioral similarity, as similarity compares a pair of behaviors whereas diversity pertains to the team as a whole. To that end, measures of diversity must be developed that apply to MRS, such as the Hierarchic Social Entropy of Balch (2000) or the diversity metric of Prorok et al. (2017). *Given an ensemble of agent behaviors or control policies, find a metric that quantifies the diversity of the team in a way that meaningfully differentiates between agents.* As in the social science literature, diversity of a robot team can be measured in many different ways. Thus, the word meaningful here can be subjective: there are many possible measures, and perhaps different measures would be useful for different tasks, for example, depending on whether a team with small differences in parameters that do not lead to significant differences in behavior should be considered diverse.

4.5. Value of diversity

Diversity alone will not solve the robustness issue; it should bring value to the team as they work on the task. Some tasks, for example where the optimal policy can be computed and all robots can execute the optimal policy, may not require diversity. In the presence of uncertainty, however, diversity can bring significant value to the team. *Given a team task, predict the value that diversity will bring to that task.* This value can depend on the level of uncertainty, the size and nature of the task, and potential restrictions on robot capabilities (whether inherent or induced by the environment). Prediction is necessary in order to avoid wasting significant time and resources on developing diversity within the team.

4.6. Architectures for diversity

Diversity of behavior within the team may also lead to different architectures for cooperation. For example, it may lead to leader–followers-type behavior, a hierarchy of subgroups within the team, a flat organization with each agent making independent decisions, and many other possibilities. These architectures will likely be influenced by the nature of the task as well as the capabilities of the robots themselves (e.g., computation and communication). *Given a team of robots on a team task, determine the best architecture for cooperation among the robots.*

4.7. Integrating small group theory

Considering potential architectures for diversity leads naturally to the problem of creating an effective mix of behaviors within the team. Outside of robotics, human workgroups have been studied extensively. Many theories have been developed on effective groups (see, for example, Poole and Hollingshead (2005) for a collection of works). These theories can be leveraged in order to better coordinate teams of robots on complex tasks. Small group theory, however, must be adapted in order to compensate for differences in human and robot capabilities. Collaboration with social scientists who study human workgroups can lead to novel methods of cooperation for MRS, but only if we understand the role diversity plays in team success. Quantify such a property, then use it to construct effective teams for completing team tasks that perform better than existing approaches. *Given a team of robots on a team task, determine the best combination of diverse behavioral types to include in the team.*

4.8. Adjusting policies online

To successfully utilize a diverse set of controllers, the team of robots must collectively reason about the role that each team member plays and automatically adjust their own roles to achieve an appropriately diverse team with an effective skill set. To do so, they must have the ability to measure the success of individual agents on a coordinated task, adjusting based on their own and others' shortcomings and successes. Evaluating an individual's performance within the team may not be straightforward. For example, the value of a defender in a soccer game cannot be quantified by the number of goals the defender's team scores. Thus, *given a robot completing part of a team task, develop a metric to evaluate its performance within that team.* Once we can quantify an individual's performance, we can consider adjusting control policies within the team accordingly. *Given an ensemble of control policies within a team of robots, develop methods to adjust individual agents' behaviors, either parametrically or otherwise, based on the agent's observations of itself or other agents.*

5. Discussion

This paper has proposed DART, Diversity-enhanced Autonomy for Robot Teams, a new research thrust that represents a paradigm shift in problem-solving for MRS. The current problem-solving paradigm is linear: control policies are optimized for a specific set of assumptions and applied to the entire team. We have proposed a paradigm wherein an ensemble of control policies is developed, with multiple sets of assumptions or interaction strategies, and exists synergistically within a team of robots. Such diversity in control policies may better prepare the team of robots for unstructured and uncertain environments, much like diversity in the knowledge base in human workgroups leads to higher-quality solutions. Adoption of this new paradigm may lead to expanded success of MRS in the field, especially in challenging dynamic environments. The DART philosophy can also be applied to single-robot systems, by enabling individual robots to switch or adjust controllers in response to failure conditions. In this way, single-robot systems may also be more successful in uncertain and unstructured scenarios.

A small sample of open problems in DART were discussed, but there exist many open problems in this space. To further the reach of MRS into such environments will require collaboration of roboticists with experts in machine learning, biological and social sciences, human–computer interaction, and many other fields. By explicitly defining DART, we hope to inspire the development of new tools for coping with uncertain, unstructured environments such as disaster response, precision agriculture, surveillance, and others.

Acknowledgements

Special thanks to Gaurav Sukhatme, M. Ani Hsieh, and Fei Sha for conversations and guidance in formulating this line of research.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by NSF CAREER Award IIS-1553726 and the Okawa Foundation Research Award.

ORCID iD

Nora Ayanian https://orcid.org/0000-0002-8394-6912

References

Agha-mohammadi A, Ure NK, How JP and Vian J (2014) Health aware stochastic planning for persistent package delivery missions using quadrotors. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3389–3396.

Amato C, Konidaris GD, Cruz G, Maynor CA, How JP and Kaelbling LP (2015) Planning for decentralized control of multiple robots under uncertainty. In: *IEEE International Conference on Robotics and Automation*, pp. 1241–1248.

Amato C, Konidaris GD and Kaelbling LP (2014) Planning with macro-actions in decentralized POMDPs. In: *International Conference on Autonomous Agents and Multi-agent Systems*, pp. 1273–1280.

Ancona D and Caldwell D (1992) Demography and design: Predictors of new product team performance. *Organization Science* 3: 321–341.

Balch T (1997) Learning roles: Behavioral diversity in robot teams. In: *AAAI Workshop on Multiagent Learning*.

Balch T (2000) Hierarchic social entropy: An information theoretic measure of robot group diversity. *Autonomous Robots* 8(3): 209–238.

Barret B (2016) Disney's latest attraction? 300 drones flying in formation. *Wired*, https://www.wired.com/2016/11/disneys-latest-attraction-300-drones-flying-formation/ .

Bera A, Kim S and Manocha D (2016) Online parameter learning for data-driven crowd simulation and content generation. *Computers and Graphics* 55: 68–79.

Berman S, Halasz A, Kumar V and Pratt S (2007) Bio-inspired group behaviors for the deployment of a swarm of robots to multiple destinations. In: *IEEE International Conference on Robotics and Automation*, pp. 2318–2323.

Buffet O, Dutech A and Charpillet F (2007) Shaping multi-agent systems with gradient reinforcement learning. *Autonomous Agents and Multi-Agent Systems* 15(2): 197–220.

Burns JG and Dyer AG (2008) Diversity of speed-accuracy strategies benefits social insects. *Current Biology* 18: R953–R954.

Chernova S and Veloso M (2010) Confidence-based multi-robot learning from demonstration. *International Journal of Social Robotics* 2(2): 195–215.

Chittka L, Skorupski P and Raine NE (2009) Speed–accuracy tradeoffs in animal decision making. *Trends in Ecology and Evolution* 24: 400–407.

Chung TH, Clement MR, Day MA, Jones KD, Davis D and Jones M (2016) Live-fly, large-scale field experimentation for large numbers of fixed-wing UAVs. In: *IEEE International Conference on Robotics and Automation*, pp. 1255–1262.

Crosland MWJ (1990) Variation in ant aggression and kin discrimination ability within and between colonies. *Journal of Insect Behavior* 3: 359–379.

D'Andrea R and Wurman P (2008) Future challenges of coordinating hundreds of autonomous vehicles in distribution facilities. In: *IEEE International Conference on Technology for Practical Robot Applications*, pp. 80–83.

de Wit FRC, Greer LL and Jehn KA (2012) The paradox of intragroup conflict: A meta-analysis. *Journal of Applied Psychology* 97(2): 360–390.

DJI (2016) DJI documents faster search and rescue responses with drones. DJI Newsroom, http://www.dji.com/newsroom/news/dji-documents-faster-search-and-rescue-responses-with-drones.

Dorigo M, Birattari M and Stutzle T (2006) Ant colony optimization. *IEEE Computational Intelligence Magazine* 1(4): 28–39.

Foerster J, Assael YM, de Freitas N and Whiteson S (2016) Learning to communicate with deep multi-agent reinforcement learning. In: *Advances in Neural Information Processing Systems*, pp. 2137–2145.

Freelan D, Wicke D, Sullivan K and Luke S (2015) Towards rapid multi-robot learning from demonstration at the robocup competition. In: RAC Bianchi, HL Akin, S Ramamoorthy and K Sugiura (eds.) *RoboCup 2014: Robot World Cup XVIII*. Cham: Springer International Publishing, pp. 369–382.

Glaser A (2016) Intel invented a way for a single operator to fly hundreds of drones at once. *Recode*, https://www.recode.net/2016/11/4/13517550/intel-single-operator-fly-hundreds-drones-shooting-star .

Glaser A (2017) These surveillance robots will work together to chase down suspects. *Recode*, https://www.recode.net/2017/4/18/15264908/surveillance-robots-network-cornell-suspects.

Goldberg D and Matarić MJ (1997) Interference as a tool for designing and evaluating multi-robot controllers. In: *Proceedings AAAI*, Providence, RI, pp. 637–642.

Guy SJ, Chhugani J, Curtis S, Dubey P, Lin M and Manocha D (2010) Pledestrians: A least-effort approach to crowd simulation. In: *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '10)*. Goslar, Germany: Eurographics Association, pp. 119–128.

Hagerty JR (2015) Meet the new generation of robots for manufacturing. *Wall Street Journal*, https://www.wsj.com/articles/meet-the-new-generation-of-robots-for-manufacturing-1433300884.

Hauert S, Leven S, Varga M, et al. (2011) Reynolds flocking in reality with fixed-wing robots: Communication range vs. maximum turning rate. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 5015–5020.

Hoffman L and Maier N (1961) Quality and acceptance of problem solutions by members of homogeneous and heterogeneous groups. *Journal of Abnormal and Social Psychology* 62: 401–407.

Hoffman LR (1978) The group problem-solving process. In: L Berkowitz (ed.) *Group Processes*. New York: Academic Press, pp. 101–114.

Huang J, Farritor SM, Qadi A and Goddard S (2006) Localization and follow-the-leader control of a heterogeneous group of mobile robots. *IEEE/ASME Transactions on Mechatronics* 11(2): 205–215.

Ijspeert AJ, Martinoli A, Billard A and Gambardella LM (2001) Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots* 11(2): 149–171.

Jackson S (1992) Team composition in organizations. In: S Worchel, W Wood and J Simpson (eds.) *Group Process and Productivity*. London: Sage.

Jandt JM, Bengston S, Pinter-Wollman N, et al. (2014) Behavioural syndromes and social insects: Personality at multiple levels. *Biological Reviews* 89: 48–67.

Jehn KA (1995) A multimethod examination of the benefits and detriments of intragroup conflict. *Administrative Science Quarterly* 40(2): 256–282.

Jennings JS, Whelan G and Evans WF (1997) Cooperative search and rescue with a team of mobile robots. In: *International Conference on Advanced Robotics*, pp. 193–200.

Kenny C (2009) *Why Do People Die In Earthquakes? The Costs, Benefits And Institutions Of Disaster Risk Reduction In Developing Countries*. The World Bank, https://elibrary.worldbank.org/doi/abs/10.1596/1813-9450-4823 .

Kumar GP, Buffin A, Pavlic TP, Pratt SC and Berman SM (2013) A stochastic hybrid system model of collective transport in the desert ant *Aphaenogaster cockerelli*. In: *International Conference on Hybrid Systems: Computation and Control*. New York: ACM Press, pp. 119–124.

Kushleyev A, Mellinger D, Powers C and Kumar V (2013) Towards a swarm of agile micro quadrotors. *Autonomous Robots* 35(4): 287–300.

Lee KH, Choi MG, Hong Q and Lee J (2007) Group behavior from video: A data-driven approach to crowd simulation. In: *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '07)*. Aire-la-Ville, Switzerland: Eurographics Association, pp. 109–118.

Li L, Martinoli A and Abu-Mostafa YS (2004) Learning and measuring specialization in collaborative swarm systems. *Adaptive Behavior* 12(3–4): 199–212.

Liu Y and Nejat G (2013) Robotic urban search and rescue: A survey from the control perspective. *Journal of Intelligent Robotics Systems* 72(2): 147–165.

Lonsdorf K (2017) Hungry? call your neighborhood delivery robot. *NPR Morning Edition*, http://www.npr.org/sections/alltechconsidered/2017/03/23/520848983/hungry-call-your-neighborhood-delivery-robot.

Lowe R, Wu Y, Tamar A, Harb J, Abbeel P and Mordatch I (2017) Multi-agent actor-critic for mixed cooperative–competitive environments. *arXiv preprint arXiv:1706.02275*.

Lyu Y, Chen Y and Balkcom D (2016) *k*-survivability: Diversity and survival of expendable robots. *IEEE Robotics and Automation Letters* 1(2): 1164–1171.

Marcolino LS, Jiang AX and Tambe M (2013) Multi-agent team formation: Diversity beats strength? In: *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI '13)*. AAAI Press, pp. 279–285.

Martinoli A, Easton K and Agassounon W (2004) Modeling swarm robotic systems: A case study in collaborative distributed manipulation. *The International Journal of Robotics Research* 23(4–5): 415–436.

Martinoli A and Mondada F (1997) Collective and cooperative group behaviours: Biologically inspired experiments in robotics. In: O Khatib and JK Salisbury (eds.) *Experimental Robotics IV.* Berlin: Springer, pp. 1–10.

Martins MF and Demiris Y (2010) Learning multirobot joint action plans from simultaneous task execution demonstrations. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, pp. 931–938.

Matarić MJ (1994) Learning to behave socially. In: *International Conference on Simulation of Adaptive Behavior*, Vol. 617, pp. 453–462.

Matignon L, Laurent GJ and Le Fort-Piat N (2012) Independent reinforcement learners in cooperative Markov games: A survey regarding coordination problems. *The Knowledge Engineering Review* 27(1): 1–31.

Modlmeier AP and Foitzik S (2011) Productivity increases with variation in aggression among group members in Temnothorax ants. *Behavioral Ecology* 22: 1026–1032.

Mouret JB and Doncieux S (2011) Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evolutionary Computation* 20(1): 91–133.

Nemeth C (1986) Differential contributions of majority and minority influence. *Psychological Review* 93: 23–32.

O'Reilly CA and Flatt S (1989) Executive team demography: Organizational innovation and firm performance. Working paper, School of Business, University of California at Berkeley.

Parker LE (1998) Alliance: An architecture for fault tolerant multirobot cooperation. *IEEE Trans Robotics and Automation* 14(2): 220–240.

Parker LE and Tang F (2006) Building multirobot coalitions through automated task solution synthesis. *Proceedings of the IEEE* 94(7): 1289–1305.

Pelechano N, Allbeck JM and Badler NI (2007) Controlling individual agents in high-density crowd simulation. In: *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '07)*. Aire-la-Ville, Switzerland: Eurographics Association, pp. 99–108.

Peshkin L, Kim K, Meuleau N and Kaelbling LP (2000) Learning to cooperate via policy search. In: *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pp. 489–496.

Pimenta LCA, Kumar V, Mesquita RC and Pereira GAS (2008) Sensing and coverage for a network of heterogeneous robots. In: *47th IEEE Conference on Decision and Control*, pp. 3947–3952.

Poole MS and Hollingshead AB (eds.) (2005) *Theories of Small Groups: Interdisciplinary Perspectives*. Thousand Oaks, CA: Sage Publications Inc.

Preiss JA, Hoenig W, Sukhatme GS and Ayanian N (2017) Crazyswarm: A large nano-quadcopter swarm. In: *IEEE International Conference on Robotics and Automation*.

Prorok A, Hsieh MA and Kumar V (2017) The impact of diversity on optimal control policies for heterogeneous robot swarms. *IEEE Transactions on Robotics* 33(2): 346–358.

Pruitt JN and Riechert SE (2011) How within-group behavioural variation and task efficiency enhance fitness in a social group. *Proceedings of the Royal Society B: Biological Sciences* 278: 1209–1215.

Pugh K J, Soros LB and Stanley KO (2016) Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI* 3: 40.

Recchia T, Chung J and Pochiraju K (2013) Improving learning in robot teams through personality assignment. *Biologically Inspired Cognitive Architectures* 3: 51–63.

Rubenstein M, Ahler C and Nagpal R (2012) Kilobot: A low cost scalable robot system for collective behaviors. In: *IEEE International Conference on Robotics and Automation*, pp. 3293–3298.

Schneider-Fontan M and Matarić MJ (1998) Territorial multi-robot task division. *IEEE Transactions on Robotics and Automation* 15(5): 815–822.

Schurr N, Marecki J, Tambe M, Scerri P, Kasinadhuni N and Lewis J (2005) The future of disaster response: Humans working with multiagent teams using defacto. In: *AAAI Spring Symposium on AI Technologies for Homeland Security.*

Steiner ID (1972) *Group Process and Productivity.* San Diego, CA: Academic Press.

Sugawara K, Kazama T and Watanabe T (2004) Foraging behavior of interacting robots with virtual pheromone. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 3, pp. 3074–3079.

Sung C, Ayanian N and Rus D (2013) Improving the performance of multi-robot systems by task switching. In: *IEEE International Conference on Robotics and Automation*. pp. 2999–3006.

Tang F and Parker LE (2007) A complete methodology for generating multi-robot task solutions using ASyMTRe-D and market-based task allocation. In: *IEEE International Conference on Robotics and Automation*, pp. 3351–3358.

Tavakoli A, Nalbandian H and Ayanian N (2016) Crowdsourced coordination through online games (Late Breaking Report). In: *ACM/IEEE International Conference on Human–Robot Interaction*, Christchurch, New Zealand.

Umedachi T, Ito K and Ishiguro A (2015) Soft-bodied amoeba-inspired robot that switches between qualitatively different behaviors with decentralized stiffness control. *Adaptive Behavior* 23(2): 97–108.

van de Vliert E and de Dreu CKW (1994) Optimizing performance by conflict stimulation. *International Journal of Conflict Management* 5(3): 211–222.

Vassiliades V and Christodoulou C (2016) Behavioral plasticity through the modulation of switch neurons. *Neural Networks* 74: 35–51.

Zefran M and Burdick JW (1998) Design of switching controllers for systems with changing dynamics. In: *Proceedings of the 37th IEEE Conference on Decision and Control*, Vol. 2, pp. 2113–2118.