Evolving division of labor in varying environments

Paweł Lichocki
EDIC, I&C, EPFL LIS, STI, EPFL

***Abstract*—The field of multi-agent systems is concerned with societies of autonomous agents that interact to efficiently achieve their goals. In this work, we will evolve *in silico* teams of agents that are capable of displaying division of labor. We hope to broaden the understanding of the evolutionary dynamics of fixed and adaptive mechanisms of division of labor. The main innovation of our work stems from the fact that we consider varying environments, which have not yet been studied broadly in the context of division of labor. The purpose of our research is twofold. From engineering perspective, we wish to improve techniques of team optimization and intend to explore the area of bio-inspired task allocation algorithms. From biological perspective, we hope to shed some light on evolutionary history and mechanical explanations of division of labor in social insects.**

***Index Terms*—division of labor, task allocation, multi-agent systems, evolution, team optimization, adaptation, distributed control, social insects**

I. Introduction

Research in the field of *multi-agent systems* spans both natural and engineering sciences (Figure 1). Engineers uncover the complexities arising from the interactions between multiple agents, seeking for distributed control algorithms that would result in the desired joint behavior of the entire system [1]. Biologists study insect societies, seeking for evolutionary roots and mechanistic explanations for behavioral traits observed in colonies of ants, bees and termites [2]. Finally, there is a growing interest on both sides to understand the optimization and evolution of complex systems in dynamically changing scenarios [3].

Proposal submitted to committee: June 25th, 2010; Candidacy exam date: July 2nd, 2010; Candidacy exam committee: Martin Hasler, Dario Floreano, Laurent Keller, Alcherio Martinoli.

This research plan has been approved:

Date: ______________________

Doctoral candidate: ______________________
(P. Lichocki) (name and signature)

Thesis director: ______________________
(D. Floreano) (name and signature)

Thesis co-director: ______________________
(L. Keller) (name and signature)

Doct. prog. director: ______________________
(R. Urbanke) (signature)

Consider a *team* composed of multiple *agents*[1] to be distributed among several types of tasks. In particular, agents exhibit high level of autonomy, meaning there is no central controller coordinating their decisions. We say a team displays *division of labor* [4] if workers self-distribute to different types of tasks. We define the *target distribution* [5], as the optimal agents distribution to different types of tasks under given environmental conditions. In the division of labor problem agents are expected to distribute themselves, using some *task allocation mechanism*, in such a manner that the team reaches the target distribution. We distinguish two types of task allocation mechanisms: *fixed* [6], [7] (i.e. constant during the team's lifetime) and *adaptive* [8], [9] (i.e. agents are adopting to the environment during the team's lifetime). In case of *varying environments* [3] the target distribution can change over time [5]. We recognize that the target distribution may change with two different types of frequencies: during the *lifetime* of the team [2] (i.e. during the simulation of the multi-agent system) or on the *evolutionary* time-scale (i.e. during the optimization process, every several generations [3]).

Our goal is to understand the evolutionary dynamics of multi-agent systems capable of displaying division of labor in varying environments. Our main scientific goal is to study the *team composition* [1], [10], [11] and the *agent behavioral complexity* [1], [2] (i.e. whether agents use fixed or adaptive task allocation mechanisms) with respect to different division of labor scenarios. There are two particular questions we wish to address. How diverse does a team need to be in order to efficiently display a given target distribution? Do evolution selects for adaptive mechanisms of division of labor in varying environments? From engineering perspective, addressing those issues could lead to new scalable design methods and optimization techniques suited for systems composed of large number of agents. From biological perspective, we hope to shed light on evolutionary explanations of fixed and adaptive task allocation mechanisms proposed for social insects [4], [6], [8]. The contributions of our work stems from the fact, that we consider varying environments, which has not yet been explored broadly in the context of division of labor. We suspect that varying environments may enhance both the selections for diverse teams and complex agents. However, the exact interplay between these two types of selection forces remains unknown.

[1]In the biological context, we talk about a *colony* composed of multiple *workers*.

Fig. 1. The field of multi-agent systems is concerned with societies of autonomous agents that interact to efficiently achieve their goals. Research spans both natural and engineering sciences, with some examples: (a) Ants nest in the forest of eastern North America ©Alex Wild, www.myrmecos.net (b) Termite soldiers rushing forward to guard a breach in their nest ©Alex Wild, www.myrmecos.net (c) Micro-robots gathering pucks ©lis.epfl.ch (d) Team of robots exploring the area ©Jiuguang Wang, Wikimedia Commons (e) Cellular automata implementing Conway's game of life ©Justin Windle, blog.soulwire.co.uk (f) Computational cluster balancing the load ©lacal.epfl.ch

Additionally, we wish to explore the area of task allocation algorithms for multiple autonomous agents, that take inspiration from biological models [4], [12], [6], [13], [8], [14], [15]. It has been shown that bio-inspired approaches give a promise of light-weighted and robust control mechanisms for coordinating the team behavior that could be competitive with classical approaches [16], [17]. In fact, there are many domains where those biological concepts have been forged into engineering solutions [18], with examples in manufacturing or scheduling systems [19], [20], [21], multi-robots [22], [23] and computation balancing [24], [25]. However, the work done so far has focused on specific applications. We believe a unifying approach should be developed, as it could integrate research effort from different fields. Furthermore, we believe that considering changing target distribution is a major step forward, as it allows to study more realistic scenarios [5], [26], [27].

The big challenge is to automate the process of creating systems composed of multiple agents, capable of displaying given target distribution. Panait and Luke [1] provide a broad survey on the general subject of multi-agent optimization[2]. They show different algorithms used to automatically derive teams of agents jointly solving a given problem. We discuss some of them in section II, focusing on tools which we aim to use in this work, namely *evolutionary computation* and *team optimization*. Furhermore, Panait and Luke expose open issues of optimal team composition and multi-agent systems scalability, which are important motivations for our work.

Panait and Luke explore in details the subject of multi-agent optimization, but they do not mention the division of labor problem explicitly. Waibel et al. [2] fill in this gap from a biological perspective by evolving $in\ silico$ colonies of artificial ants, capable of displaying division of labor. The authors propose a framework to model and study the emergent results of the joint behavior of multiple agents, which we intent to use and extend. Within the framework, they compare the performance of colonies (i.e. how close the displayed division of labor is to the target distribution) using three simple task allocation mechanisms. Interestingly, it is investigated which cases allow an efficient colony response, when external perturbations are applied during the lifetime of the colony. We describe this work in section III.

In section IV we present in more details the notion of varying environments. We describe the work of Kashtan et al. [3] who use standard genetic algorithms to evolve networks that are able to make computations. Alternatively to Waibel et al. [2], they vary the environments every several generations and not during the lifetime of the examined system. They show that this might significantly contribute to the speed of natural evolution and suggest a way to accelerate optimization algorithms. The authors consider many types of evolving networks, but they do not concern themselves with multi-agent systems directly. We conclude this paper with a discussion about foreseen challenges and a thesis plan.

[2]Panait and Luke use a general term multi-agent *learning*. However, biologist typically understand learning as a behavioral development within the life-time of the individual (i.e. durable modification of behavior in response to information acquired) [28]. In order to avoid confusions we consequently use the term *optimization* instead.

II. Multi-agent optimization

We intend to evolve $in\ silico$ cooperative agents that display division of labor at the team level. The approach we take falls into the category of team optimization, presented in details by Panait and Luke [1]. Team optimization involves a single process that discovers a set of behaviors for the entire team. Typically, this is achieved by stochastic search methods, with a notable example of evolutionary computation.

A. Evolutionary computation

For instance, using multiple agents imposes an explosion of the state space (if agent A or B can be in any of 100 states, than agents A and B together can be in 100×100 states). This makes it difficult to apply popular *reinforcement learning* techniques, which seek how to map states to actions, so as to maximize a numerical reward signal [1], [29]. Rather than that, team optimization relies on stochastic search methods, that directly learn entire behaviors without referring to separate states. Often, this is achieved by means of evolutionary computation, which abstract Darwinian principles of natural selection to iteratively refine a population of candidate solutions (called individuals). This class of method uses a *fitness*-oriented procedure, meaning that individuals are evaluated and the ones with higher performance get selected. The selected individuals are slightly and randomly changed (mutated) and/or crossed-over with other selected individuals (breed) in order to create the population of the next iteration step (generation). This process is repeated and the population gets refined, until the time is exhausted or individual with sufficient performance is discovered.

In particular, Panait and Luke mention collaborative evolution (or ‘one-population’ coevolution), where a single population of agents is evolved using standard genetic algorithms, but sets of agents are evaluated together by randomly teaming them up. It has been reported that this specific heterogeneous approach may yield better results than homogenous team optimization [1], [30]. Furthermore, the appearance of ‘sub-species’ (i.e. *specialization*) has been observed when using this method [1], [31]. We believe it could be especially useful when evolving multi-agent systems displaying division of labor, as typically some level of agents’ specialization is required.

B. Team composition

Usually, evolutionary computation better than reinforcement learning handles huge search spaces [1]. However, due to using multiple agents the size of the search space still remains a key issue. To some extent, we can influence the size of the search space, as it is inherently related to the team composition. Agents can be *homogeneous*, meaning that all are assigned identical behavioral controller (Figure 2.a). Alternatively, agents can be *heterogenous*, meaning that each is assigned a unique behavioral controller (Figure 2.b). Homogeneity reduces the search space, as the optimization process needs to discover just one prototypical behavior. On the other hand, heterogeneity potentially allows for more diversity in a team, which might help when agent specialization is needed (which is often the case in division of labor scenarios [26], [32]). There is also a middle-ground approach, called *hybrid* team optimization, where teams are divided into squads (castes) and there is one prototypic behavior evolved for every squad (Figure 2.c). It is believed that hybrid methods might yield benefits of both homogenous and heterogenous team optimization [1].

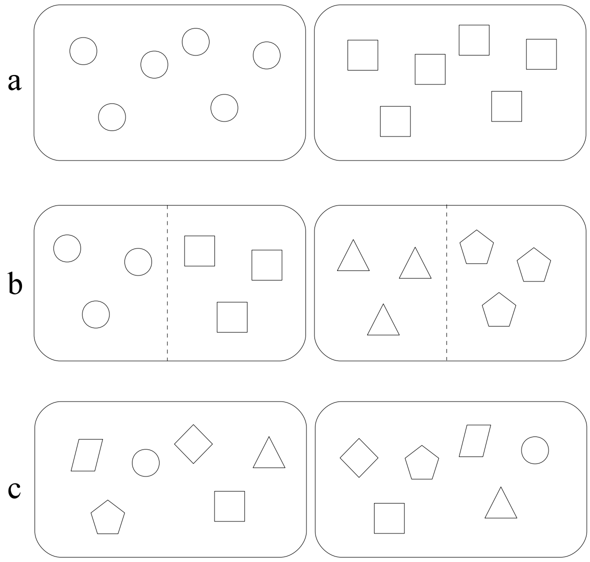

Fig. 2. Examples of three categories of team composition: (a) two homogeneous teams composed of six agents with identical behavioral controllers, (b) two hybrid teams composed of two squads of three agents, behavioral controllers of agents are identical within the squad, but different between the squads (c) two heterogeneous teams composed of six agents with different behavioral controllers.

Generally, it is suggested that choosing among homogenous and heterogeneous approaches depends on whether specialists are required or not [1], [33]. However, even if agents are homogenous they may act heterogeneously, due to different conditions. For example two agents with the same behavioral controller may perform different actions, depending on their actual sensory state [1]. This suggests that in cases of more complex agents, team diversity might no longer be an important issue. However, it has been also shown that increasing team diversity can significantly boost the optimization process, despite apparent domain homogeneity [1], [34]. To take advantage over this fact, there has been proposed a hybrid method that tries to automatically discover the optimum number of groups and their composition [1], [35]. However in the end, it remains unclear what team composition is optimal and under which conditions [1]. We wish to address this question in the context of multi-agent systems displaying division of labor.

C. Other open issues

In the context of multiple agents, one of the goals is to understand the dynamics of the optimization procedure, which is not trivial due to interactions between the agents. Panait and Luke recognize that varying environments present a particular challenge in finding optimal team behaviors [1]. They advise to intensify the research by considering dynamically changing scenarios. Our study of division of labour in varying environments aims to partially address this. However, while Panait and Luke are interested in the difficulties arising from co-adaptations between concurrent agents, we are focused on the evolutionary dynamics of the entire team.

Finally, Panait and Luke state that one overlooked property of multi-agent optimization methods is their scalability [1]. The authors express the opinion that for large systems it is impossible to optimize a joint behavior of a fully heterogeneous team. They believe that in such cases, one would always be forced to resort to the hybrid team optimization. Implicitly, they put forward a question about the optimal team composition which could facilitate the evolutionary process (for a particular problem at hand). We think that the division of labor problem constitutes a suitable framework to address this issue, because teams can be of any arbitrary size (from just few up to many thousands, or even more). Thus, it enables us to study the efficiency of the optimization process with respect to the size of the multi-agent system.

III. Biological perspective

We find it important to review biological literature on division of labor, as we are interested in bio-inspired task allocation mechanisms. Here, we focus on the paper by Waibel et al. [2] to present some biological explanations of division of labor observed in insect societies [15]. Waibel et al. study *response threshold models* [4], [6], [8], [12], [32], which assume that workers in a colony vary in the response threshold at which they begin to perform the corresponding task. Workers with a low threshold are more likely to perform that task than workers with a high threshold. Consequently, intra-colony variation in individual thresholds results in division of labour among the workers. In particular, Bonabeau et al. [6] proposed a model in which the response thresholds of an agent are determined genetically and are fixed during its lifetime. Theraulaz et al. [8] extended this model by introducing adaptive thresholds that change during the lifetime of the agent, according to the outcome of its actions.

Waibel et al. [2] evolve $in\ silico$ colonies of artificial ants capable of displaying division of labor. They consider three different response thresholds mappings between genotype and behavior. Each worker has five genes ($g_1 - g_5$) which it uses to allocate itself to one of the five tasks. In the 'deterministic mapping' a worker chooses the task with the highest genetically encoded value $\max(g_i)$. In the 'probabilistic mapping' a worker choses the task i with the probability $P_i = g_i / \sum_{j=1}^{5} g_j$. In the 'dynamic mapping' a worker choses a task that maximizes the ratio g_i/a_i, where a_i is the proportion of the colony performing task i.

A. Division of labor

Importantly, Waibel et al. provide a useful framework to implicitly model target distributions [2]. They design a scenario in which workers can engage in five different tasks, for which they do not specify an explicit target distribution.

Rather than that, they introduce a payoff measure for each task as a function of proportion of workers engaged in the corresponding task. The performance of the entire colony is a sum of all five payoff functions, which implicitly defines the target distribution. A notable fact is that Waibel et al. use payoff functions that are sublinear. They define the payoff for a task i as $1 - e^{-ix_i}$, where x_i is the proportion of workers performing task i. Importantly, with sublinear payoff functions, the gain of allocating more workers to a given task drops, as the number of worker allocated rise. Intuitively, this approach models the overcrowding of workers, which might occur in real-life scenarios and applications [12], [22], [36]. We intend to use and extend this approach to model division of labor with various target distributions (e.g. uniform, unimodal, bi-modal, gaussian).

B. Colony relatedness

Waibel et al. reveal that the type of mapping between genotype and individual behavior (i.e. the type of task allocation mechanism) strongly influences colony productivity [2]. They also recognize the important role of *colony relatedness*, which corresponds to the notion of team composition. In details, Waibel et al. show that the deterministic mapping performs well only in the case of a very low colony relatedness. On contrary, for treatments with the probabilistic and dynamic mappings, it is the higher relatedness that translates into higher performances. The reason is that in deterministic case, the only source of different worker's behaviors (i.e. displaying division of labor) is the genetic variation. Whereas, both probabilistic and dynamic mappings increase the phenotypic variance, which suffices to reach the vicinity of the target distribution, even in case of highly related workers. Furthermore, high relatedness allows a more efficient selection of genotypes, than in situation where workers are unrelated. Concluding, the comparison of the three mappings illustrates the importance of maintaining sufficient diversity in teams displaying division of labor (i.e. the interplay between team composition and agent behavioral complexity).

C. Perturbations

Finally, Waibel et al. introduce perturbations during the lifetime of the colony (*not* on the evolutionary time-scale). Repetitively in the simulations, they choose one task at random and remove all workers performing it. Then, they add new workers to recreate the colony in order to observe how the workers re-allocate themselves to the five tasks. Waibel et al. conclude that only the dynamical task allocation scheme yields top performance (i.e. close to global optimum), regardless of the experimental conditions [2]. The explanation is that the dynamic mapping allows the replacement of workers doing a particular task by new workers also performing the same task. On contrary, both deterministic and probabilistic mappings do not allow the preferential replacement of workers doing the same task, hence introducing the perturbations is detrimental to the colony performance in those cases. In conclusion, only the dynamical mappings give a promise of efficient task allocation in environments varying during the lifetime of the colony.

IV. Varying environments

Waibel et al. explain the impact of perturbing the colony during its lifetime into the performance of dividing the labor. One other line of investigation is to change the target distribution less frequently, namely on the evolutionary time-scale. Kashtan et al. [3] show that such varying environments may speed-up the evolution of networks capable of doing computations. If a similar phenomena occurs in the evolution of multi-agent systems, one could significantly improve the process of team optimization.

Kashtan et al. evolve $in\ silico$ four types of networks capable of doing computations: logic circuits, feed-forward logic circuits, feed-forward neural networks and continues function circuits. They vary environments on the evolutionary time-scale, i.e. they switch between two, or more, different fitness functions every 20 generations. In particular, Kasthan et al. focus on modularly varying environments, where each new goal shares some of the subproblems with the previous goal. For example, consider two subgoals described by functions $f(x, y)$ and $h(w, z)$. They can be combined by a function g into a modular goal $G = g(f(x, y), h(w, z))$. Or they can be combined by other function g' into another modular goal $G' = g'(f(x, y), h(w, z))$. These two goals G and G' constitute a pair of modularly varying goals, as they share common subgoals (i.e. functions f and h) and only the aggregation function (i.e. functions g and g') changes.

A. Evolutionary speed-up and modularity

It turns out that modular variations significantly speed up the optimization process. The reason is that each time a goal changes, a gradient for the new goal is generated in the search space [3]. This prevents evolution from being stuck with locally optimal solutions, which often happens in the case of the static fitness function. Interestingly, if two environments vary modularly, the local maxima or plateaus in one goal correspond to areas with positive fitness gradient for the second goal. In result, the population rapidly reaches an area in the search space where solutions maximizing both fitness functions exist. Then, if the goal switches, the networks can rapidly re-evolve to cope with the new environment. Importantly, Kasthan et al. notice that modularly varying environments have much stronger impact that a pure randomized approach (i.e. switching between main goal and randomly generated goals) [3]. They hypothesize that it might be because modularity in the environment imposes modularity in the solutions (as shown in their previous work [37]). This enforces the evolution of modules corresponding to shared subgoals, which is believed to be the main source of observed evolutionary speed-ups. In the light of this finding, we hypothesize that team composition might be a major factor dictating weather the speed-up occurs in the evolution of systems displaying division of labor. We think that a very diverse team might correspond to a solution with no modularity, whereas a team composed of few big squads, might correspond to a solution with few big modules. We wish to verify that claim during our work.

B. Other results

Interestingly, Kashtan et al. notice that the speed-up increases with the complexity of the goal (defined as T_{FG}, which is the time needed to evolve solution for a fixed goal 'from scratch') [3]. They support this observation with a formula showing that the speed-up roughly follows a power law in terms of the goal complexity: $S \sim T_{FG}^{\alpha}$, where $alpha = 0.7 \pm 0.1$. This suggests that, if a phenomena of the evolutionary speed-up actually occurs for multi-agent systems displaying division of labor, it might be significant, due to the high complexity of the team optimization problem.

Next, Kasthan et al. compare their results with other optimization methods. First, they find similar speed-ups for modularly varying goals in the case of the hill-climbing algorithm (an iterative approach that moves from a current solution to a better one in the local neighborhood in the search space and stops when no local refinement is possible [38]). This suggests that their findings are related to the shape of the fitness landscape, rather than to particular optimization method used. Finally, Kashtan et al. also test a multiobjective optimization problem, that considers two different goals at once (previously varying in time). They observe no evolutionary speed-up in this case. This means that multiple modular goals by themselves, with no variation in time, are not sufficient for the speed-up to occur. We find it interesting to continue the work of Kashtan et al., by bringing it into the domain of multi-agent systems.

V. DISCUSSION

In general, our goal is to study the evolutionary dynamics of multi-agent systems capable of displaying division of labor in varying environments. First, Panait and Luke motivated our venture into the field of multi-agent optimization by posing several open issues (i.e. optimal team composition, scalability, dynamics of the optimization procedure) [1]. In addition, they equipped us with appropriate optimization methods (i.e. team optimization, evolutionary computation, collaborative evolution). Next, Waibel et al. provided us with the tools to model division of labor problem and a biological inspiration for task allocation algorithms [2]. They also show that the perturbations introduced during the lifetime of the colony, select for dynamical mechanisms of division of labor. Finally, Kashtan et al. give evidence that varying environments less frequently (every few generations) may significantly speed-up the evolutionary process [3].

We identify several main steps in our work, each requiring important decisions to make, that we need to justify. First of all, we aim at developing a generic framework to study multi-agent systems displaying division of labor. We take both engineering and biological perspectives, which is a considerable challenge. Fortunately, engineers commonly use evolutionary computation in team optimization problems [1] and biologists are interested in evolutionary history of division of labor observed in nature [39]. In other words, they all use an evolutionary approach (although to accomplish different objectives). For this reason, we focus on one particular search and optimization technique, namely artificial evolution. There still remains a question which genetic algorithm to use. We will adopt the collaborative evolution, briefly described by Panait and Luke [1], as it seems to be well-suited for evolving systems displaying division of labor (e.g. it allows specialization to occur spontaneously). Importantly, we are still left with the problem of defining an appropriate fitness function, that will drive the optimization process [40]. Our current idea is to minimize the difference between target distribution and the actual division of labor displayed by the evolved system. However it is unknown, whether such a fitness function will provide a strong enough gradient to guide the evolutionary process.

Also, we must develop a framework to model division of labor in a multi-agent system. We have decided to follow Waibel et. al [2] and their idea of implicitly defining target distribution with payoff functions. This a very flexible tool, which will allow us to depict realistic division of labor scenarios (e.g. with overcrowding effects). Next, we must choose the task allocation mechanism, according to which workers will divide the labor. To this purpose we intent to extend the response threshold models within a formalism of *artificial neural networks* [41], [42]. We find this approach particularly interesting, as neuronal formalism can easily accommodate many variations. For example, neural networks may easily incorporate plasticity [43], [44] which would allow to model adaptive mechanisms of division of labor [8], [45]. Furthermore, the neuronal formalism will make it possible to extend our work in other interesting directions in the future (e.g. including memory [46]).

Finally, we will model the varying environments by switching between two or more target distributions. We recognize that these variations could be applied on the evolutionary time-scale [3] or during the lifespan of the system [2]. We currently do not know how frequently we should switch between different target distributions. Furthermore, we suspect that the evolution may favor different mechanisms of task allocation, depending on this frequency. This is one of the main issues that we need to investigate and resolve.

We plan that the design, implementation and testing phase of all the tools needed will last for no longer than a year. With the framework in place, we will conduct three main numerical experiments. The first one aims to answer what is the optimal team composition, with respect to the team size, number of tasks and different types of constant target distribution (e.g. uniform, uni-modal, gaussian). In the second, analogous experiment we will introduce the idea of varying environments, to see whether they enforce the selection for more diverse teams. In the last experiment, we will study the evolution of adaptive mechanisms of division of labor to understand the interplay of team diversity and agent behavioral complexity. The entire experimental phase, along with analysis of the results, is intended to last for approximately two and a half years.

Additionally, it would be interesting to see whether some of the numerical experiments could be transferred into a robotic platform [47] or if the results could be applied to real-life problems related to division of labor (listed in the introduction). These extensions have a lower priority (than the

Fig. 3. The approximated schedule of the thesis.

numerical experiments) and are planned as a series of mini projects. They should be assigned to students in a continuous manner during entire time devoted for this research. Finally, we estimate that the whole project, along with thesis writing, will take us no more than three years (Figure 3).

In brief, we will evolve $in\ silico$ multi-agents systems displaying division of labor in varying environments. We hope to broaden the understanding of the evolutionary dynamics of fixed and adaptive mechanisms of division of labor. We focus on the interplay of the two key concepts: team composition and agent behavioral complexity. Furthermore, in our work we consider varying environments, which has not yet been studied systematically in the context of division of labor. The purpose of our research is twofold. From engineering perspective, we wish to improve techniques of team optimization and intend to explore the area of bio-inspired task allocation algorithms. From biological perspective, we hope to shed some light on evolutionary history and mechanical explanations of division of labor in social insects.

Finally, we recognize that although the multi-agent systems field is relatively young, it has already developed interesting connections to other research areas, with notable examples of: multi-robot task allocation [48], [49], [36], cellular automata [50], distributed decision making [51], [52], [53], [54], message passing algorithms [55], scheduling [56], [57] and game theory [58], [59]. This plethora of different domains may give birth to many currently not foreseen contributions of our work.

REFERENCES

[1] L. Panait and S. Luke, "Cooperative multi-agent learning: The state of the art," *Autonomous Agents and Multi-Agent Systems*, vol. 11, no. 3, pp. 387–434, 2005.

[2] M. Waibel, D. Floreano, S. Magnenat, and L. Keller, "Division of labour and colony efficiency in social insects: effects of interactions between genetic architecture, colony kin structure and rate of perturbations," *Proceedings of the Royal Society B*, vol. 273, no. 1595, pp. 1815–1823, 2006.

[3] N. Kashtan, E. Noor, and U. Alon, "Varying environments can speed up evolution," *Proceedings of the National Academy of Sciences*, vol. 104, no. 34, pp. 13 711–13 716, 2007.

[4] S. Beshers and J. Fewell, "Models of Division of Labor in Social Insects," *Annual Review of Entomology*, vol. 46, no. 1, pp. 413–440, 2001.

[5] S. Berman, A. Halasz, M. Hsieh, and V. Kumar, "Optimized stochastic policies for task allocation in swarms of robots," *IEEE Trans. on Robotics*, pp. 927–937, 2009.

[6] E. Bonabeau, G. Theraulaz, and J. Deneubourg, "Quantitative study of the fixed threshold model for the regulation of division of labour in insect societies," *Proceedings of the Royal Society B: Biological Sciences*, vol. 263, no. 1376, pp. 1565–1569, 1996.

[7] ——, "Fixed response thresholds and the regulation of division of labor in insect societies," *Bulletin of Mathematical Biology*, vol. 60, no. 4, pp. 753–807, 1998.

[8] G. Theraulaz, E. Bonabeau, and J. Deneubourg, "Response threshold reinforcements and division of labour in insect societies," *Proceedings of the Royal Society B: Biological Sciences*, vol. 265, no. 1393, pp. 327–332, 1998.

[9] D. Merkle and M. Middendorf, "dynamic polytheism and competition for tasks in threshold reinfrocement models of social insects," *Adaptive Behavior*, vol. 12, no. 3-4, pp. 251–262, 2004.

[10] M. Waibel, L. Keller, and D. Floreano, "Genetic team composition and level of selection in the evolution of cooperation," *IEEE Transactions on Evolutionary Computation*, vol. 13, no. 3, pp. 648–660, 2009.

[11] D. Tarapore, D. Floreano, and L. Keller, "Task-dependent influence of genetic architecture and mating frequency on division of labour in social insect societies," *Behavioral Ecology and Sociobiology*, vol. 64, no. 4, pp. 675–684, 2010.

[12] G. Robinson, "Regulation of division of labor in insect societies," *Annual Review of Entomology*, vol. 37, no. 1, pp. 637–665, 1992.

[13] D. Gordon, "The organization of work in social insect colonies," *Nature*, vol. 380, no. 6570, pp. 121–124, 1996.

[14] A. Hirsh and D. Gordon, "Distributed problem solving in social insects," *Annals of Mathematics and Artificial Intelligence*, vol. 31, no. 1, pp. 199–221, 2001.

[15] E. Robinson, O. Feinerman, and N. Franks, "Flexible task allocation and the organization of work in ants," *Proceedings of the Royal Society B: Biological Sciences*, vol. 276, no. 1677, pp. 4373–4380, 2009.

[16] W. Agassounon and A. Martinoli, "Efficiency and robustness of threshold-based distributed allocation algorithms in multi-agent systems," in *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 3*. ACM New York, NY, USA, 2002, pp. 1090–1097.

[17] N. Kalra and A. Martinoli, "Comparative Study of Market-Based and Threshold-Based Task Allocation," *Distributed Autonomous Robotic Systems 7*, pp. 91–101, 2006.

[18] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm intelligence: from natural to artificial systems*. Oxford University Press, USA, 1999.

[19] E. Bonabeau, A. Sobkowski, G. Theraulaz, and J. Deneubourg, "Adaptive task allocation inspired by a model of division of labor in social insects," *Bio Computation and Emergent Computing*, pp. 36–45, 1997.

[20] M. Campos, E. Bonabeau, G. Theraulaz, and J. Deneubourg, "Dynamic

scheduling and division of labor in social insects," *Adaptive Behavior*, vol. 8, no. 2, pp. 83–92, 2000.
[21] W. Xiang and H. Lee, "Ant colony intelligence in multi-agent dynamic manufacturing scheduling," *Engineering Applications of Artificial Intelligence*, vol. 21, no. 1, pp. 73–85, 2008.
[22] M. Krieger, J. Billeter, and L. Keller, "Ant-like task allocation and recruitment in cooperative robots," *Nature*, vol. 406, pp. 992–995, 2000.
[23] T. Labella, M. Dorigo, and J. Deneubourg, "Division of labor in a group of robots inspired by ants' foraging behavior," *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, vol. 1, no. 1, p. 25, 2006.
[24] J. Cao, "Self-organizing agents for grid load balancing," in *Proceedings of the Fifth IEEE/ACM International Workshop on Grid Computing*, 2004, pp. 388–395.
[25] T. Freeman, K. Keahey, I. Foster, A. Rana, B. Sotomoayor, and F. Wuerthwein, "Division of labor: Tools for growing and scaling grids," *Service-Oriented Computing–ICSOC 2006*, pp. 40–51, 2006.
[26] C. Jones and M. Matarić, "Adaptive division of labor in large-scale minimalist multi-robot systems," in *In IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2003, pp. 1969–1974.
[27] M. Matarić, G. Sukhatme, and E. Østergaard, "Multi-robot task allocation in uncertain environments," *Autonomous Robots*, vol. 14, no. 2, pp. 255–263, 2003.
[28] J. Alcock, "Animal behavior: An evolutionary approach," *Journal of Mammalogy*, vol. 84, no. 1, pp. 324–325, 2003.
[29] R. Sutton and A. Barto, *Reinforcement learning: An introduction*. The MIT press, 1998.
[30] M. Quinn, "A comparison of approaches to the evolution of homogeneous multi-robot teams," in *Proceedings of the 2001 Congress on Evolutionary Computation (CEC2001)*, 2001, pp. 128–135.
[31] T. Miconi, "A collective genetic algorithm," in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, 2001, pp. 876–883.
[32] R. Jeanson, J. Fewell, R. Gorelick, and S. Bertram, "Emergence of increased division of labor as a function of group size," *Behavioral Ecology and Sociobiology*, vol. 62, no. 2, pp. 289–298, 2007.
[33] M. Potter, L. Meeden, and A. Schultz, "Heterogeneity in the coevolved behaviors of mobile robots: The emergence of specialists," in *International Joint Conference in Artificial Intelligence*, vol. 17, no. 1, 2001, pp. 1337–1343.
[34] T. Haynes and S. Sen, "Co-adaptation in a team," *International Journal of Computational Intelligence and Organizations*, vol. 1, no. 4, pp. 1–20, 1996.
[35] A. Hara and T. Nagao, "Emergence of cooperative behavior using ADG; Automatically Defined Groups," in *Proceedings of the 1999 Genetic and Evolutionary Computation Conference (GECCO-99)*, 1999, pp. 1038–1046.
[36] T. Dahl, M. Matarić, and G. Sukhatme, "Multi-robot task allocation through vacancy chain scheduling," *Robotics and Autonomous Systems*, vol. 57, no. 6-7, pp. 674–687, 2009.
[37] N. Kashtan and U. Alon, "Spontaneous evolution of modularity and network motifs," *Proceedings of the National Academy of Sciences*, vol. 102, no. 39, pp. 13 773–13 778, 2005.
[38] Z. Michalewicz and D. Fogel, *How to solve it: modern heuristics*. Springer-Verlag New York Inc, 2004.
[39] C. Smith, A. Toth, A. Suarez, and G. Robinson, "Genetic and genomic analyses of the division of labour in insect societies," *Nature Reviews Genetics*, vol. 9, no. 10, pp. 735–748, 2008.
[40] A. Nelson, G. Barlow, and L. Doitsidis, "Fitness functions in evolutionary robotics: A survey and analysis," *Robotics and Autonomous Systems*, vol. 57, no. 4, pp. 345–370, 2009.
[41] S. Lek and J. Guégan, "Artificial neural networks as a tool in ecological modelling, an introduction," *Ecological Modelling*, vol. 120, no. 2-3, pp. 65–73, 1999.
[42] P. Lichocki, D. Tarapore, D. Floreano, and L. Keller, "Neural networks as mechanisms to regulate division of labour," in preparation.
[43] Floreano, D. and Dürr, P. and Mattiussi, C., "Neuroevolution: from architectures to learning," *Evolutionary Intelligence*, vol. 1, no. 1, pp. 47–62, 2008.
[44] D. Floreano and L. Keller, "Evolution of Adaptive Behaviour in Robots by Means of Darwinian," *PLoS Biology*, vol. 8, no. 1, p. e1000292, 2010.
[45] Y. Niv, D. Joel, I. Meilijson, and E. Ruppin, "Evolution of reinforcement learning in uncertain environments: A simple explanation for complex foraging behaviors," *Adaptive Behavior*, vol. 10, no. 1, pp. 5–24, 2002.
[46] J. Fernandez-Leon, G. Acosta, and M. Mayosky, "Behavioral control through evolutionary neurocontrollers for autonomous mobile robot navigation," *Robotics and Autonomous Systems*, vol. 57, no. 4, pp. 411–419, 2009.
[47] F. Mondada, M. Bonani, X. Raemy, J. Pugh, C. Cianci, A. Klaptocz, S. Magnenat, J. Zufferey, D. Floreano, and A. Martinoli, "The e-puck, a robot designed for education in engineering," in *Proceedings of the 9th conference on autonomous robot systems and competitions*, vol. 1, no. 1, 2009, pp. 59–65.
[48] B. Gerkey and M. Matarić, "Multi-robot task allocation: Analyzing the complexity and optimality of key architectures," in *IEEE International Conference on Robotics and Automation*, vol. 3, 2003, pp. 3862–3868.
[49] J. McLurkin and D. Yamins, "Dynamic task assignment in robot swarms," in *Proceedings of Robotics: Science and Systems*. Citeseer, 2005, pp. 2007–214 838.
[50] S. Gruner, "Mobile agent systems and cellular automata," *Autonomous Agents and Multi-Agent Systems*, vol. 20, no. 2, pp. 198–233, 2010.
[51] S. Dolev and N. Tzachar, "Empire of colonies: Self-stabilizing and self-organizing distributed algorithm," *Theoretical Computer Science*, vol. 410, no. 6-7, pp. 514–532, 2009.
[52] C. Parker and H. Zhang, "Cooperative Decision-Making in Decentralized Multiple-Robot Systems: the Best-of-N Problem," *IEEE/ASME Transactions on Mechatronics*, vol. 14, no. 2, pp. 240–251, 2009.
[53] G. Mathews, H. Durrant-Whyte, and M. Prokopenko, "Decentralised decision making in heterogeneous teams using anonymous optimisation," *Robotics and Autonomous Systems*, vol. 57, no. 3, pp. 310–320, 2009.
[54] S. Garnier, J. Gautrais, M. Asadpour, C. Jost, and G. Theraulaz, "Self-Organized Aggregation Triggers Collective Decision Making in a Group of Cockroach-Like Robots," *Adaptive Behavior*, vol. 17, no. 2, pp. 109–133, 2009.
[55] J. Chalopin, E. Godard, Y. Métivier, and R. Ossamy, "Mobile agent algorithms versus message passing algorithms," *Principles of Distributed Systems*, pp. 187–201, 2006.
[56] L. Canon, E. Jeannot, and K. Carlo, "Evaluation and Optimization of the Robustness of DAG Schedules in Heterogeneous Environments," *IEEE Transactions on Parallel and Distributed Systems*, vol. 21, no. 4, 2010.
[57] N. Immorlica, L. Li, V. Mirrokni, and A. Schulz, "Coordination mechanisms for selfish scheduling," *Theoretical Computer Science*, vol. 410, no. 17, pp. 1589–1598, 2009.
[58] Y. Shoham, R. Powers, and T. Grenager, "If multi-agent learning is the answer, what is the question?" *Artificial Intelligence*, vol. 171, no. 7, pp. 365–377, 2007.
[59] P. Tang, Y. Shoham, and F. Lin, "Designing competitions between teams of individuals," *Artificial Intelligence*, pp. 749–766, 2010.