

Ibots: Learning Real Team Solutions

Cristina Versino and Luca Maria Gambardella

IDSIA, Corso Elvezia 36, CH-6900 Lugano, Switzerland
{cristina,luca}@idsia.ch, <http://www.idsia.ch>

Abstract

This paper presents “Ibots” (Integrating roBOTS), a computer experiment in team robotics designed on an artificial mission. Our aim is to understand how to use reinforcement learning to program automatically a team of robots with a shared mission. Moreover, we are interested in *learning real team solutions*. These are programs whose form strongly depends on the number of robots composing the team, on their individual skills and limitations, and on any other mission boundary condition which makes it worth to prefer “at a team level” certain solutions to others. The Ibots mission is specified implicitly by means of a single reinforcement signal which measures the team performance as a whole. This form of payoff leads to real team solutions. Benefits and drawbacks of using team reinforcement as opposed to individual robot reinforcement are discussed.

Introduction

The use of reinforcement learning to produce self-programming robots¹ is not new in the context of *single robot* missions, where a reinforcement signal directly qualifies the behavior of the single robot in charge of carrying out a given task (Barto, Sutton, & Watkins 1989; Kaelbling 1993).

A robot credit assignment problem. The picture changes radically as *many robots* are acting at the same time, with little or perhaps no knowledge at all about teammate activities. In this scenario, if the reinforcement signal reflects the whole team performance, each single robot is faced with the problem of deciding to what extent its own behavior has contributed to the overall team’s good or bad score: this is the *robot credit assignment* problem. Because of the robot credit assignment problem, each robot has a *noisy* perception of the mission it is asked to accomplish. A single robot

can behave identically many times (during different trials of the mission), and nevertheless, it may receive completely different payoffs. This occurs because it is not the only actor, and the reinforcement signal just *partially* depends on its actions.

Bypassing the robot credit assignment problem.

Instead of addressing the robot credit assignment problem directly, one can bypass it by reformulating the team learning problem.

A *first* way is to enable *broadcast communication* between teammates. If a robot is *aware* of other robots’ perceptions and actions, then it is in a position to make sense out of a global team payoff. It has been observed that explicit communication makes the team equivalent to a “big robot” (Mataric 1995), whose perception and action are the union of perceptions and actions of all team members, the team payoff being a measure of the big robot performance. The utility of communication to achieve robot awareness has been proved experimentally on real robots both in a simplified hazardous waste cleanup mission (Parker 1995a), and in a box-pushing task (Parker 1994; Mataric, Nilsson, & Simsarian 1995). However, communication is not always possible technically, and it tends to become a bottleneck as the team size increases (Kube & Zhang 1994). Even when communication is possible, the big robot <perception, action> space may be simply too large to be efficiently explored. Thus, the explicit communication way is well-suited to small teams only.

A *second* way of avoiding the robot credit assignment problem is to measure *each robot individual performance* instead of team performance. In (Mataric 1994a; 1996) this idea is applied to training a group of real robots in a foraging task. The mission is to collect pucks disseminated in the workspace and to bring them to a home area. Each robot learns through individual payoff a private policy, which specifies how it should behave under various conditions. For exam-

¹Throughout the paper, we use the term “robot” as a synonym for “agent”.

ple, the robot is rewarded whenever it grasps a puck or if it drops a puck at home. In this framework, a single robot is by no means interested in the performance of its teammates, because it addresses the mission in an individualistic sense. We see *two drawbacks* in this approach. *First: its underlying assumption is that team performance indirectly increases because individual performance increases.* However, if the robots do not learn the task at a *similar pace*, it cannot be guaranteed that each robot will reliably learn and participate to the mission. As an example, suppose that in the foraging task one robot in the team manages (maybe just by chance) to learn the individually optimal policy after a few trials. This “super-robot” will collect most of the pucks by itself, thus diminishing the learning opportunities of its teammates because pucks are a *limited, shared resource*. Ideally, to improve this state of affairs, the super-robot should behave in a sub-optimal way (by forgetting part of its optimal policy) so as to let the other robots take part in the task and learn. However, there is no reason for the super-robot to recede from its optimal policy, because it is designed to be the best possible individual. Let us observe that the super-robot phenomenon is not as unrealistic as it appears at first glance. It may arise, for instance, in *incremental learning experiments* where the team size is progressively increased: elder, experienced robots would tend to carry out the mission by themselves, thus limiting the possibility of novice robots to participate and learn. In (Mataric 1994b), it is argued that *social rules* can be learnt by the robots to minimize resource competition and to direct their behavior away from individual greediness and towards global efficiency. Though the idea is interesting, the experimental results reported in (Mataric 1994b) are preliminary, and, to the best of our knowledge, no follow-up work has been produced along this line. The *second* drawback we find in the individualistic approach is summarized in the answer to the following question: even if the super-robot phenomenon does not arise, where are the policies learnt by the robots expected to converge? They will converge to the optimal policy for a robot carrying out the mission by itself (this is confirmed by the experiments reported in (Mataric 1994a)). The robots will behave as “clones” of *a robot designed to work alone*.

Real team solutions. We feel this violates the spirit of team learning, which should be aimed at producing *efficient real team solutions*. *These are policies whose shape is strongly influenced by the number of robots composing the team, by each robot’s skills and weaknesses, and by any other mission boundary condition*

which could be of some relevance for preferring “at a team level” some solutions to others.

Facing the robot credit assignment problem. *To obtain truly team solutions, one should use team pay-offs at the price of dealing with the ambiguity posed by the robot credit assignment problem.* Multi-agent learning experiments based on a team payoff are illustrated in recent works (Murciano & Millán 1996; Crites & Barto 1996; Schmidhuber 1996).

In (Murciano & Millán 1996) a simulated team of agents learns signaling behaviors to efficiently solve an object-gathering task in an unknown and changing environment. The reinforcement signal is based on the total time needed by the team to gather all the objects. Experiments are carried out under several conditions: with teams of different sizes, with different numbers of objects, and with different object distributions in the environment (uniformly scattered or clustered); moreover, during learning the object distribution is occasionally modified to test the agents ability of tracking environmental changes. Under these heterogeneous experimental conditions, each agent learns the appropriateness of exhibiting a given signaling behavior. For example: an agent perceiving an object, is faced with the problem of deciding whether to activate its “object-signaling” behavior to attract other agents towards its position. The agent learns that exhibiting this signaling behavior is, in general, rewarding for the team when the objects are distributed in clusters: in fact, in this case, the detection of one object gives a high chance of finding other objects in the neighborhood, and these could be collected by the called agents. On the contrary, the same signaling behavior is inappropriate if the objects are uniformly distributed in the environment and the agent should refrain from calling other teammates whenever it finds an object. Finally, certain mission boundary conditions require that the agents acquire different signaling policies: in this case, they specialize into *signallers* and *harvesters*. The authors show by statistical analysis of the results that the team discovers by trial-and-error a near-to-optimal signaling policy given specific mission conditions.

In (Crites & Barto 1996) a team of Q-learning agents is engaged in the challenging real-world problem of elevator dispatching. Each agent is responsible for controlling one elevator car. Two different control architectures are tested. In the *parallel architecture*, the agents share a single neural network which models a common policy: this allows the agents to learn from each others experiences but forces them to use identical policies. In the fully *decentralized architecture*, the agents learn their own networks, allowing them to

specialize their control policies. The team receives a global payoff which is the sum of squared wait times of passengers. Despite this noisy reinforcement signal and the inherent stochastic nature of the task, results obtained in simulation on both architectures surpass the best known heuristic elevator control algorithm. Moreover, the authors expect an additional advantage of reinforcement learning over heuristic controllers in buildings with heterogeneous arrival rates at each floor, because Q-learning agents may adapt to each floor traffic profile.

In (Schmidhuber 1996) a general method for incremental self-improvement and multi-agent learning in unrestricted environments is presented. In one of the implementations, a recurrent neural net is applied to a non-Markovian maze task. Each connection in the net is viewed as an agent: the connection's weight represents the agent's policy. The net learns to guide an animat to a goal by using a team payoff whose value turns from 0 to 1 only once the animat hits the goal. Looking at the net's connections as if they were a team of agents is an unconventional point of view. But why not doing so? After all, a neural net is a good example of a set of partially independent agents (the net's weights, or the neurons themselves) which learn to act well "as a team". Following this view, the structural credit assignment problem in connectionist reinforcement learning can be regarded as being equivalent to the robot credit assignment problem described above.

Finally, this paper presents "Ibots" (Integrating roBOTS), an experiment in collective robotics designed to understand how to use reinforcement learning to program automatically a team of robots with a shared mission. As in (Murciano & Millán 1996; Crites & Barto 1996; Schmidhuber 1996), Ibots learn through a reinforcement signal which measures the team performance as a whole. In this way, Ibots manage to learn real team solutions.

General issues

Before describing the Ibots mission, we summarize the general issues addressed by the experiment, as well as the underlying assumptions and design choices.

Team size. In the Ibots experiment, the same mission is handled with teams of *different size*. By this, our aim is to assess whether learnt policies change as the team size changes.

Team composition. Ibots can be *homogeneous* or *heterogeneous*. When they are homogeneous, they all have same sensors and same actuators. When they are heterogeneous, they have different sensors and/or

different actuators. Clearly, when Ibots are heterogeneous, they have potentially different skills and weaknesses due to physical characteristics. The question is whether they can learn to specialize their behavior so as to emphasize skills and minimize the impact of weaknesses. The same question is addressed in (Parker 1995b).

Mission's boundary conditions. Ibots are confronted with *different boundary conditions* to the mission. We are interested in checking whether the strategy learnt by the team to accomplish the mission changes as the boundary conditions change. (Murciano & Millán 1996) also address this issue by dealing with different object distributions in their object-gathering task.

Control programs. Ibots run and learn control programs which define their behavior in the mission. The Ibots' control programs may be *public* or *private*. When they are public, all Ibots share the same program; when they are private, each Ibot works with a different, personal program.

Robots working with private programs instead of public programs is more general. The public program option relies on the assumption that there exists a shared solution which works for every robot in the team. This is not case when the robots are heterogeneous or when the mission requires that the robots specialize their behaviors into different typologies. On the other hand, if the robots are homogeneous and the mission does not require specialization, it definitely makes sense to consider the public option, because it is faster to learn 1 shared program than n distinct programs, n being the number of robots in the team. Learning a shared program is faster for two reasons: first, the search space for the learning algorithm gets reduced, and, second, the robot credit assignment problem completely vanishes. Sometimes, even when the mission requires the robots to specialize their behaviors, it is possible to define the team learning problem in such a way that the public program option is still valid. As an example, we guess that if the agents in (Murciano & Millán 1996) were designed to learn directly the distribution of object-signaling agents over the whole team (instead of the individual agent's tendency to exhibit the object-signaling behavior), it would be easier for the team to learn the correct proportion between signaling and non-signaling agents. Finally, we stress that learning 1 shared program in a team of n robots is not the same as learning 1 program with a single robot and then cloning it n times. *A shared program is a real team solution, while cloned programs are not.* The

Ibots experiment illustrates this point clearly.

In our opinion, the possibility of learning a single shared program instead of several private programs has been overlooked by most collective robotics works. This is curious if one considers the recent trend of designing robot teams inspired to Swarms (Deneubourg *et al.* 1991). Robot Swarms have shown that an interesting group behavior may emerge from the interaction between robots running a public control program. Odd enough, when learning techniques are used to derive automatically a team behavior, people usually neglect the possibility of learning a public program.

The only work we are aware of where both public and private policies have been considered, is the recent paper on elevator dispatching by (Crites & Barto 1996). Their “parallel” and “decentralized” architectures are equivalent to our “public” and “private” policies, respectively. In this application, however, it remains unclear which are the benefits of using the decentralized architecture in a problem where the elevator cars to be controlled are homogeneous.

Communication. Ibots do *not communicate*. By this we precisely mean that, when Ibots work with private control programs, they do not tell each other which these programs are. In this restricted sense, there is no communication. By this choice we depart from works in collective robotics where robots are aware of teammates activities through explicit communication (Parker 1994; 1995a; 1995b; Mataric, Nilsson, & Simsarian 1995).

Sharing limited resources. In collective robotics, robots usually share physical resources (a workspace, objects to gather, etc). In our experiment, Ibots share a limited, non-physical resource. However, there is no a priori, externally imposed rule which decides how this resource should be shared. We thought that Ibots should tend to be more or less resource *greedy* depending on their control programs. Thus, while Ibots learn their control programs, they implicitly learn to share the resource in a way which is convenient for the team. In particular, when Ibots run *public programs*, they exhibit the same degree of greediness, and, on the average, each Ibot takes the same amount of the total resource. This is not the case when Ibots run *private programs*. In this situation, each Ibot has a different propensity in consuming the resource, and this fact has a great impact on the team performance in the mission.

Learning through reinforcement. Ibots learn to accomplish the mission through *trial-and-error*. At each trial, each Ibot receives a *reinforcement signal*

which measures the *team performance* as a whole. Ibots are confronted with their performance as a group because we want them to learn a team strategy. This view is shared with (Murciano & Millán 1996; Crites & Barto 1996; Schmidhuber 1996), but it is different from that pursued in (Mataric 1994a; 1996), as discussed in the introductory section.

In the next section, we describe the Ibots experiment in details.

Ibots

The *mission* for the Ibots is to guess the integral I ($0 \leq I \leq 1$) of an arbitrary gray *Region* drawn on the white ground of a squared arena (Fig. 1(a),(d)).

How are robots turned into Ibots? For the sake of clarity, let us first consider the case of a team composed of one Ibot.

One Ibot

The Ibot’s dowry is a *control program Prog* which lets it explore the arena while sampling the ground color. By activating *Prog*, the Ibot performs a *trial* run (Fig. 1 (b),(d)).

A trial starts from a random location in the arena. It is a sequence of N_{max} *elementary movements* separated by stops. Whenever the Ibot stops, it samples the ground color. A color reading equal to gray gives evidence for the sample to be “inside *Region*”, while a white reading is interpreted as “outside *Region*”. At the end of the trial, the Ibot returns the number N_{in} of samples which happened to be inside *Region*. $\hat{I} = N_{in}/N_{max}$ is its estimate of I at this trial. $E = |I - \hat{I}|$ is the error in the estimate.

How are elementary movements generated? The control program *Prog* depends stochastically on two parameters ($\alpha_{prog}, \delta_{prog}$) which are fixed during a trial. α_{prog} is used to generate a rotation instruction for the Ibot, while δ_{prog} induces a translation. The semantics of α_{prog} and δ_{prog} is as follows. First, a number α is drawn from a uniform distribution in $[-\alpha_{prog}, +\alpha_{prog}]$. This is interpreted as: “Rotate α degrees.”. Second, a number δ is drawn from a uniform distribution in $[0, \delta_{prog}]$. The interpretation for δ is: “Translate δ units in your current heading direction, calling the *bumping rule* if necessary.”. The bumping rule is called when the Ibot meets the arena border before having covered the whole distance δ . In this case, the Ibot rotates 180° and covers the remaining distance *minus* 1: bumping against the arena border consumes 1 unit of translation to prevent the Ibot from bumping forever without consuming δ , a rare case which can occur when the Ibot lies in a corner. The bumping rule




Figure 1: (From left to right) (a) The squared arena and the “half-full” *Region*: $I = 0.49$. The arena side is 500 units long. (b) Trajectory of one Ibot running a solitary trial with $N_{max} = 100$ and $Prog = (30^\circ, 200)$. Crosses indicate sampled points. The result is $N_{in} = 52$, $\hat{I} = 0.52$. (c) Trajectories of two Ibots running a shared trial with $N_{max} = 100$. They started from a scattered configuration and ran private control programs: $Prog^1 = (180^\circ, 50)$ (gray trace), $Prog^2 = (20^\circ, 150)$ (black trace). The result is $N_{in} = 55$, $N_{in}^2 = 30$, $\hat{I} = 0.85$. Ibot 1 took $N_{sam}^1 = 61$ samples, Ibot 2 took $N_{sam}^2 = 39$. (d) Trajectory of one Ibot running a trial on the “chessboard” *Region*: $I = 0.49$. $N_{max} = 100$ and $Prog = (10^\circ, 40)$. The result is $N_{in} = 45$, $\hat{I} = 0.45$.

can be called recursively. We decided the bumping rule should just reverse the Ibot’s travelling direction to minimally affect its natural motion angle which already depends on the parameter α_{prog} .

In our simulation, a rotation instruction is executed by the Ibot in one time unit whatever the rotation angle α ; while the execution time of a translation instruction is directly proportional to the distance δ .

Finally, a last remark on α_{prog} and δ_{prog} : both parameters take values in finite ranges, $0 \leq \alpha_{prog} \leq \alpha_{max}$ and $0 \leq \delta_{prog} \leq \delta_{max}$.

A Team of Ibots

Each Ibot i ($i = 1, \dots, N_{ibots}$) being equipped with a control program $Prog^i = (\alpha_{prog}^i, \delta_{prog}^i)$, there are several ways of generalizing from the single Ibot case to the team case (Fig. 1 (c)). We have considered both the cases where the $Prog^i$ s may be *public* or *private*. Public means that all instances of $Prog^i$ s are constrained to be the same $Prog$, while private means that each $Prog^i$ may be different.

The Ibots activate these $Prog^i$ s in parallel to run a shared trial. At the beginning of the trial, the Ibot locations are chosen at random in the arena, and these can be either *clustered* or *scattered*. In the clustered configuration, all Ibots have the same initial position and orientation; in the scattered configuration, they have different positions and orientations. During a trial, the Ibots are granted a total of N_{max} elementary movements to collect N_{max} samples of the ground color *overall: samples are a limited, shared resource*. As a consequence, whenever an Ibot stops to take a sample, it is allowed to do so only if less than N_{max} samples have been taken by the team so far. Otherwise, the Ibot gives up the sampling, and the team trial

is terminated. Notice that if the Ibots work with private $Prog^i$ s, each Ibot i will collect a different number of samples N_{sam}^i : Ibots with small δ_{prog}^i s (travelling for shorter distances) will take on the average more samples than Ibots with larger δ_{prog}^i s. At the end of the trial, each Ibot i returns the number N_{in}^i of samples it counted inside *Region*. These contributions are summed in $N_{in} = \sum_i N_{in}^i$, leading to the team integral estimate $\hat{I} = N_{in}/N_{max}$, the error being $E = |I - \hat{I}|$.

Finally, Ibots are *immaterial*, they do not collide when their trajectories intersect. This may result disappointing to real roboticists who must take into account the physical interaction between robots. But, in the context of this paper, this fact is irrelevant. We are not trying to simulate real robots, and, though a physical interaction rule could be learnt as well, we preferred to emphasize the Ibots interaction due to the shared number of samples.

Programmed Ibots vs. Learning Ibots

How can a team of Ibots learn to provide good estimates \hat{I} of the integral I ? Referring to the proposed framework, *the goal of learning is to find control programs $Prog^i$ s which lead to estimates \hat{I} close to I* . As the behavior of each control program depends on its parameters α_{prog}^i and δ_{prog}^i , *the target of learning is to discover “good” pairs $(\alpha_{prog}^i, \delta_{prog}^i)$* . Notice that we know a *general* solution, namely:

$$\forall i : \begin{cases} \alpha_{prog}^i &= 180^\circ \\ \delta_{prog}^i &= \text{“length of arena diagonal”} \end{cases} \quad (1)$$

With this choice of parameters, an Ibot may reach *any* position in the arena starting from *any* other position and orientation in just *one* elementary move-

ment. Which means, knowing the Ibot's current position and orientation, its next position remains highly unpredictable. An external observer would describe the points sampled by this Ibot as being *uniformly distributed* in the arena. This brings us to the hypothesis of the Monte Carlo method (Hammersley & Handscomb 1964) for integration. This states that, by drawing N_{max} points from a uniform distribution in the arena, the error E in the estimate of the integral is probabilistically bounded by N_{max} :

$$P \left\{ E \leq \frac{1}{\sqrt{N_{max}}} \right\} \geq 0.9999 \quad (2)$$

For example, by drawing 100 points, we are almost guaranteed that the error in the estimate will not exceed 0.1, whatever the *Region's* integral and shape (remember that $0 \leq I \leq 1$). This result quantifies the *admissible error* for the Ibot mission, given N_{max} . We call the control programs defined by Eq. 1 the "programmer solution", because it reflects, in our opinion, the way a programmer would address this robot programming task: by looking for a *general* solution, which will work for any *Region*, whatever the number of Ibots in the team, independently of their starting configuration in the arena. Though appealing, we are not interested in this *a priori* solution. Rather, we are looking for *real team solutions* established through experience. These should depend on the specific *Region*, on the number of Ibots, on their initial configuration, and on their specific skills when these latter are no longer homogeneous.

To see why this "adaptation to circumstances" makes sense, consider the extreme case when *Region* coincides with the whole arena ($I = 1$). In this situation, it would be *unnatural* to see the Ibots running the programmer's solution, as the integral is perfectly guessable by the more economic "staying in place" program:

$$\forall i : \begin{cases} \alpha_{prog}^i & = 0^\circ \\ \delta_{prog}^i & = 0 \end{cases}$$

How can real team solutions be derived? In general, exhaustive search in the multi-dimensional program space of the Ibots is unfeasible. Suppose the Ibot's *continuous* program space is discretized so that each Ibot selects its control program only among a set of N_{progs} different control programs. Then, in a team of size N_{ibots} , the number of possible distinct program selections for the team is (Tucker 1995):

$$\binom{N_{ibots} + N_{progs} - 1}{N_{ibots}} = \frac{(N_{ibots} + N_{progs} - 1)!}{N_{ibots}!(N_{progs} - 1)!}$$

Moreover, each program selection should be tested several times to be sure about its quality, because "bad"

control programs have a stochastic performance. As an alternative to exhaustive search, true team solutions can be derived through learning: by measuring the Ibots team performance as a whole.

Learning to Be Good Ibots

"Good" *Prog's* are programs which lead *repeatedly to admissible* estimates of I . Admissible estimates are defined with respect to N_{max} by Eq. 2. Moreover, estimates must be repeatable because we are interested in programs with a *stable* performance.

The Ibots learn good control programs *Prog's* by reinforcement through a sequence of trials. Each time instant t corresponds to a team trial. Let $Prog^i(t) = (\alpha_{prog}^i(t), \delta_{prog}^i(t))$ be Ibot i control program at time t . The following 4 steps are repeated forever.

1. Each Ibot " i " independently generates a new, tentative control program $New^i(t) = (\alpha_{new}^i(t), \delta_{new}^i(t))$ by slightly modifying $Prog^i(t)$.

Given:

$$\begin{aligned} \alpha_{temp}^i &= \alpha_{prog}^i(t) + \alpha_{rand}^i(t) \cdot \alpha_{step} \\ \delta_{temp}^i &= \delta_{prog}^i(t) + \delta_{rand}^i(t) \cdot \delta_{step} \end{aligned}$$

it is:

$$\begin{aligned} \alpha_{new}^i(t) &= \begin{cases} 0^\circ & \text{if } \alpha_{temp}^i < 0^\circ \\ \alpha_{max} & \text{if } \alpha_{temp}^i > \alpha_{max} \\ \alpha_{temp}^i & \text{otherwise} \end{cases} \\ \delta_{new}^i(t) &= \begin{cases} 0 & \text{if } \delta_{temp}^i < 0 \\ \delta_{max} & \text{if } \delta_{temp}^i > \delta_{max} \\ \delta_{temp}^i & \text{otherwise} \end{cases} \end{aligned}$$

where:

- $\alpha_{rand}^i(t)$ and $\delta_{rand}^i(t)$ are uniform random numbers in $[-\rho(t), +\rho(t)]$ (see below for the definition of $\rho(t)$);
 - α_{step} and δ_{step} are constants.
2. The Ibots collectively carry out a trial with the newly generated programs $New^i(t)s$.

At the beginning of the trial, the Ibots are positioned at a random configuration (which may be clustered or scattered). At the end of the trial, each Ibot returns $N_{in}^i(t)$. The team integral estimate is $\hat{I}(t) = N_{in}(t)/N_{max}$, with $N_{in} = \sum_i N_{in}^i$. The error $E(t)$ in the estimate is:

$$E(t) = \frac{|I - \hat{I}(t)|}{\max(I, 1 - I)}$$

3. *The team reinforcement signal $R(t)$ is computed and communicated to each Ibot.*

$R(t)$ is defined as the difference between the team errors in two successive trials:

$$R(t) = E(t - 1) - E(t) \quad (3)$$

$E(t - 1)$ can be thought of as a naïve predictor of $E(t)$.

4. *Each Ibot “ i ” independently computes a modification for its control program and updates it.*

Being:

$$\begin{aligned} \Delta\alpha_{new}^i(t) &= \alpha_{new}^i(t) - \alpha_{new}^i(t - 1) \\ \Delta\delta_{new}^i(t) &= \delta_{new}^i(t) - \delta_{new}^i(t - 1) \end{aligned}$$

the program modification is:

$$\Delta\alpha_{prog}^i(t) = \epsilon \cdot R(t) \cdot \Delta\alpha_{new}^i(t) \quad (4)$$

$$\Delta\delta_{prog}^i(t) = \epsilon \cdot R(t) \cdot \Delta\delta_{new}^i(t) \quad (5)$$

where ϵ , a real parameter, is the learning rate. Given:

$$\begin{aligned} \alpha_{temp}^i &= \alpha_{prog}^i(t) + \Delta\alpha_{prog}^i(t) \\ \delta_{temp}^i &= \delta_{prog}^i(t) + \Delta\delta_{prog}^i(t) \end{aligned}$$

the updated control programs is:

$$\begin{aligned} \alpha_{prog}^i(t + 1) &= \begin{cases} 0^\circ & \text{if } \alpha_{temp}^i < 0^\circ \\ \alpha_{max} & \text{if } \alpha_{temp}^i > \alpha_{max} \\ \alpha_{temp}^i & \text{otherwise} \end{cases} \\ \delta_{prog}^i(t + 1) &= \begin{cases} 0 & \text{if } \delta_{temp}^i < 0 \\ \delta_{max} & \text{if } \delta_{temp}^i > \delta_{max} \\ \delta_{temp}^i & \text{otherwise} \end{cases} \end{aligned}$$

The description of the algorithm is completed by the following remarks and definitions.

- When the Ibots work with a public control program, only one Ibot is allowed to generate the tentative program $New^i(t)$ at step 1; then, $New^i(t)$ is communicated to the other team members.
- About the error measure $E(t)$: by dividing $|I - \hat{I}(t)|$ by $\max(I, 1 - I)$, $E(t)$ varies between 0 and 1, no matter what the value of I . As a consequence, the reinforcement signal also varies in a fixed interval, namely $[-1, +1]$.

- $E(-1) = E(0)$. At trial 0, we assume that the expected error $E(-1)$ is equal to the measured error $E(0)$. This implies $R(0) = 0$, so no change is made to $Prog(0)$.
- $\rho(t) = \max(|R(t - 1)|, \rho_c)$. The amount of variation in the new control programs is proportional to the reinforcement signal received at previous trial. This is to enhance the tendency of escaping from programs with unpredictable performance, and, viceversa, to favor the convergence towards programs with stable performance. ρ_c is a small positive constant which maintains a minimal level of exploration in program space when $R(t - 1) = 0$.

We conclude this section by placing this experiment in the reinforcement learning panorama. We have set the stage for a *nonassociative, immediate* reward learning experiment. This experiment qualifies as nonassociative, because there is no perception-action mapping to be learned: the only input to the learners is the reinforcement signal. Moreover, since this is fed to the Ibots as soon as the trial is finished and the trial is the team’s atomic action, this is immediate reinforcement learning. Equation 3 establishes that the payoff is positive if the $New^i(t)$ s provided a better estimate of the integral than the $New^i(t - 1)$ s, negative if the estimate was worse, zero if it was equal (equally good or equally bad). Thus $R(t)$ decides the direction of change in the control programs. As an example², if $\alpha_{new}^i(t)$ is greater than $\alpha_{new}^i(t - 1)$, $\alpha_{prog}^i(t + 1)$ will increase or decrease with respect to $\alpha_{prog}^i(t)$ depending on whether $R(t)$ was positive or negative (Eq. 4). On the contrary, if $\alpha_{new}^i(t)$ is less than $\alpha_{new}^i(t - 1)$, $\alpha_{prog}^i(t + 1)$ will increase if $R(t)$ is negative, decrease if $R(t)$ is positive. In all cases the direction of change in the control programs is meant to increase the probability of those $Prog^i$ s which proved to be better in the integral estimate process. Moreover, $Prog^i$ s remain unchanged when $R = 0$. This learning method is a non-associative version of the basic, connectionist reinforcement learning algorithm proposed in (Barto, Sutton, & Brouwer 1981). Overall the learning algorithm is expected to guide the team towards admissible and stable control programs.

Experiments

On the “half-full” *Region* ($I = 0.49$) of Fig. 1(a), we have run repeated learning experiments with Ibots’ teams of increasing size ($N_{ibots} = 1, \dots, 14$), with public or private $Prog^i$ s, and starting from clustered or scattered configurations. We have also considered the

²All the following observations made on α_{new}^i , clearly hold for δ_{new}^i too, see Eq. 5.

situation where the Ibots' skills are no longer homogeneous because of differences in sensing and acting capabilities. In all experiments we have set: $N_{max} = 100$ (the corresponding admissible error being 0.1), $\alpha_{max} = 180^\circ$ and $\delta_{max} = 500$ (length of arena side), $\alpha_{step} = \alpha_{max}/10$ and $\delta_{step} = \delta_{max}/10$, $\rho_c = 0.1$, and, for each Ibot i , $Prog^i(0) = (\alpha_{prog}^i(0), \delta_{prog}^i(0)) = (0^\circ, 0)$: this "staying in place" program was chosen *to bias* the Ibots towards "economic" programs, i.e. programs requiring small translations. To examine the form of the learnt programs, a learning experiment was stopped either after having obtained 10 consecutive admissible estimates, or after a fixed number of trials, depending on which of these two events occurred first. The former stopping criterion was introduced for convenience, for not devoting too much time to experiments which did converge rapidly.

One Ibot

The single Ibot experiment is a point of reference for comparing results obtained with teams of Ibots. It requires to discover a pair $(\alpha_{prog}, \delta_{prog})$ which produces admissible and stable integral estimates. As the program space searched by the learning algorithm is bidimensional, one can explore it in a systematic way to test the quality of a significant number of programs. Thus, before starting the actual learning experiments, we have run background trials with all combinations of α_{prog} and δ_{prog} , with α_{prog} ranging in $\{0^\circ, 10^\circ, \dots, 180^\circ\}$, and δ_{prog} ranging in $\{0, 25, \dots, 500\}$. Each combination program was tested on 20 different trials to have a sample of estimates $\{\hat{I}_k, k = 1, \dots, 20\}$. Then, for each program we computed the *mean of errors* $\mu_{prog}(E)$ and the *variance of errors* $\sigma_{prog}^2(E)$:

$$\begin{aligned}\mu_{prog}(E) &= \frac{1}{20} \sum_k |I - \hat{I}_k| = \frac{1}{20} \sum_k E_k \\ \sigma_{prog}^2(E) &= \frac{1}{20} \sum_k (E_k - \mu_{prog}(E))^2\end{aligned}$$

$\mu_{prog}(E)$ and $\sigma_{prog}^2(E)$ describe the suitability of the corresponding program for the integration task: $\mu_{prog}(E)$ gives an indication on the accuracy of the estimates, while $\sigma_{prog}^2(E)$ measures their stability. Figure 2 shows the plots in program space of $\mu_{prog}(E)$ (left) and of $\sigma_{prog}^2(E)$ (right): the program space axes are indexed by δ_{prog} and α_{prog} . On the $\mu_{prog}(E)$ plot we have highlighted the contour lines of level 0.1: these lines identify the space of admissible programs. Notice that these programs are also stable (see $\sigma_{prog}^2(E)$ plot). Most of them are distant from the "programmer solution" $(180^\circ, 700)$: admissible and stable solutions start

at $(50^\circ, 200)$. Observe also that large values of α_{prog} require large values for δ_{prog} to be admissible, because small values of δ_{prog} would confine the Ibot's motion to a localized area (as an example, see the gray trace of Ibot 1 in Fig. 1 (c)). Finally, from the $\sigma_{prog}^2(E)$ plot we remark that not only admissible programs are stable. For example, all "staying in place" programs ($\delta_{prog} = 0$) have a very predictable performance. This makes the learning task more difficult as the Ibot's initial program $Prog(0) = (0^\circ, 0)$ acts as a local minimum with respect to the stability criterion.

The results of 20 different learning experiments for the single Ibot have been overlaid on the left plot of Fig. 2. The convergence point of each experiment is indicated by a square. Learning stopped in all cases before the limit of 1000 trials. All experiments ended inside the space of admissible and stable programs.

For the sake of comparison, Fig. 3 shows the plots of $\mu_{prog}(E)$ (left) and of $\sigma_{prog}^2(E)$ (right) for one Ibot running on the "chessboard" *Region* ($I = 0.49$, Fig. 1 (d)). The space of admissible programs is considerably larger than the one for the "half-full" *Region*, because the integral of the "chessboard" *Region* can be predicted even by a more localized motion. Hence, the Ibot can afford travelling shorter distances to do a good job (Fig. 1 (d)).

Teams of Ibots

The form of admissible and stable programs changes radically for teams of Ibots.

Public Control Programs. The first case study we have addressed is that of Ibots equipped with a public control program. As in the single Ibot case, the program space is bidimensional, so we first ran background trials (with no learning) for teams of increasing size, both for the clustered and the scattered configuration. The results for the largest team of 14 Ibots are shown in Fig. 4. The left plot is the contour plot of $\mu_{prog}(E)$ for the clustered configuration, while the right plot shows the behavior of $\mu_{prog}(E)$ for the scattered configuration. On the left plot, the bold line delimits the space of admissible programs; for the scattered configuration, all programs result to be admissible (all contour lines are below level 0.1).

By comparing these results with those of Fig. 2 (left), one observes how the single Ibot's solutions space "shrinks" or "expands" depending on whether the Ibots are started in the clustered or in the scattered way. Why?

First, consider the clustered starting condition. As the Ibots begin a trial from the same position and with the same orientation, they have to disperse in order to

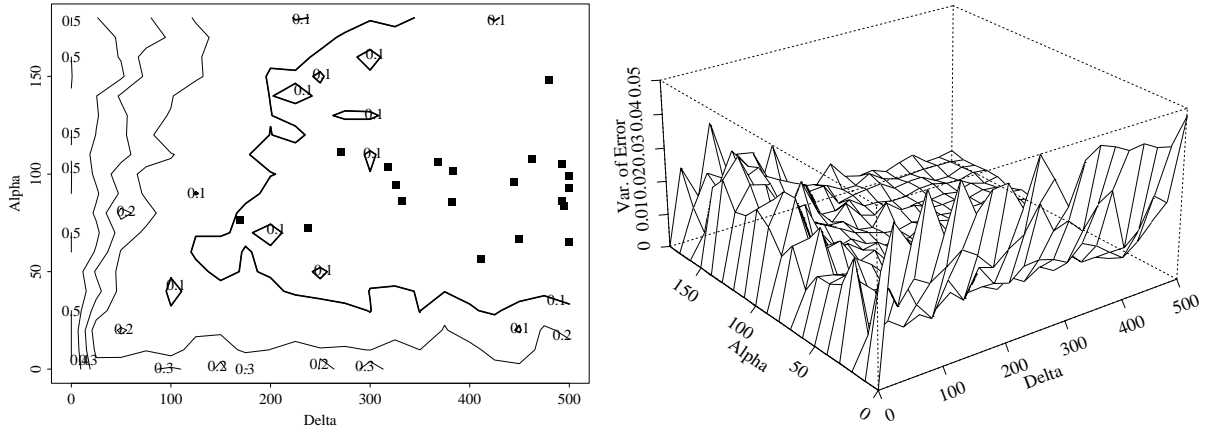


Figure 2: One Ibot on the “half-full” *Region*. (Left) Contour plot of $\mu_{prog}(E)$ and convergence points for 20 learning experiments. (Right) Plot of $\sigma^2_{prog}(E)$.




Figure 3: One Ibot on the “chessboard” *Region*. (Left) Contour plot of $\mu_{prog}(E)$ and convergence points for 20 learning experiments. (Right) Plot of $\sigma^2_{prog}(E)$.

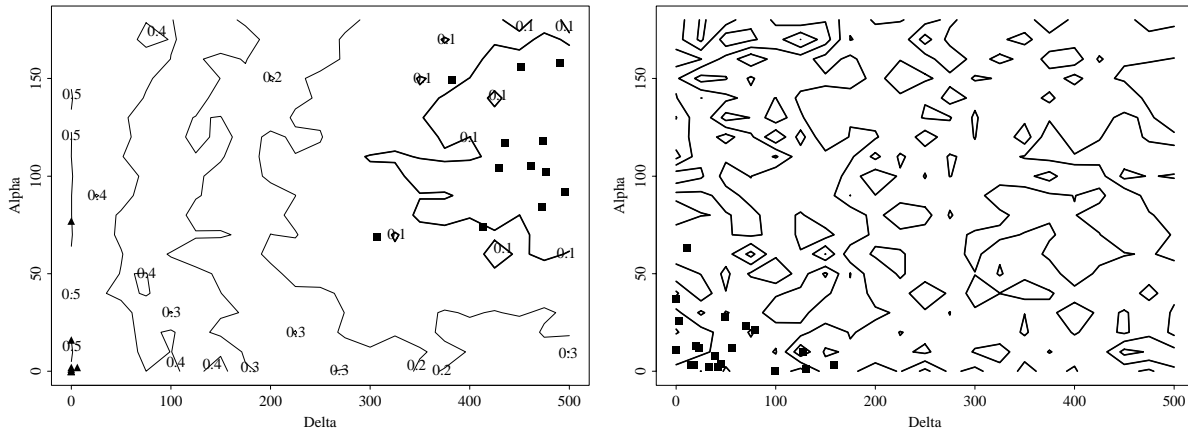


Figure 4: A team of 14 Ibots with a public control program on the “half-full” *Region*. Contour plot of $\mu_{prog}(E)$ and convergence points for 20 learning experiments for clustered configurations (left) and for scattered configurations (right).




Figure 5: A team of 14 Ibots running trials with learnt public control programs. (From left to right) (a) Samples and (b) trajectories for a team started in the clustered configuration and running the control program: $Prog = (100^\circ, 450)$. (c) Samples and (d) trajectories for a team started in the scattered configuration and running the control program: $Prog = (10^\circ, 50)$.

explore the arena. Moreover, the quantity of samples they are allowed to take as individuals decreases as the number of Ibots increases, because the samples budget is fixed for the team. Thus, each Ibot is granted fewer samples and fewer elementary movements to disperse in the arena. Given this constraint, the only way of achieving rapid dispersion is through control programs with large variability both in translation and in rotation. *In conclusion, most of the solutions which are valid for the single Ibot would not work for this team.*

Second, consider the opposite case where the Ibots start a trial from scattered positions. As they are already uniformly distributed in the arena, any kind of motion program would lead to admissible estimates. In principle, the solution which would profit the most from this favorable start would be to perform a very localized motion around each Ibot’s initial position. *We stress that this team solution would not be admissible for the single Ibot.*

Figure 4 also reports the convergence points of 20 learning experiments for both initial configuration types. Experiments which converged to admissible and stable programs within the time limit of 2000 trials are represented by squares, while non-converging experiments are represented by triangles. Not surprisingly, all the experiments for the scattered configuration converged very rapidly (right). On the contrary, for the clustered configuration, not all experiments managed to converge to admissible programs within the predefined time limit (left). This is due to the fact that the Ibots are initiated with the shared $Prog^i(0) = (0^\circ, 0)$, a bad but very stable program. As a matter of fact, the stability of this program becomes stronger as the number of Ibots grows. Therefore, it may take a considerable amount of time for the team to get away from this inconvenient initial program.

Finally, figure 5 shows trajectories and samples taken by a clustered ((a) and (b)) and a scattered team ((c) and (d)) of 14 Ibots running trials with learnt public programs.

Private Control Programs. Figure 6 shows a representative sample of learning experiments performed with a team of 14 Ibots working with private control programs and started from clustered (plots on the first row) or scattered configurations (plots on the second row). Each plot shows the programs learnt by the team within a time limit of 20000 trials; programs which did not converge have been represented by triangles (second and third plot on the first row).

Essentially, these solutions are similar to those obtained with the public control program. Within the same category of initial configuration, Ibots learn the

same typology of programs: clustered Ibots need large variability in angle and translation, while scattered Ibots don’t. This uniformity in the shape of the solutions is not surprising, because the integration task does not require any differentiation in behavior as long as the Ibots have homogeneous skills.

From the point of view of learning, the main difference between dealing with a single public program or with many private programs is the robot credit assignment problem, which clearly arises when the Ibots learn personal programs. Figure 7 illustrates the robot credit assignment problem for a team of two Ibots. The first and the second time plot present the evolution of the Ibots’ δ_{prog}^i s and α_{prog}^i s parameters, respectively; the third time plot shows the error E in the integral estimate produced by the team. Observe that, around time 30, Ibot 1 (dashed line) has already acquired an admissible program, but this does not appear at the level of team performance because Ibot 2 (solid line) is still locked to the initial “staying in place” program. Consider also that Ibot 2’s contribution to the team integral weighs more than Ibot 1’s contribution, because Ibot 2 takes more samples: “bad Ibots count more”. This is also the cause for the two non-converging experiments of Fig. 6, where a minority of Ibots translating for short distances damage the team performance. To improve from this state of affairs, Ibot 1 (Fig. 7), first backtracks from its admissible program, and then it relearns at a similar pace with Ibot 2.

Private Control Programs for Heterogeneous Ibots. As a last experiment, we wanted to make the learnt control programs more specialized. A way of achieving this is by differentiating the Ibots’ individual skills. Figure 8 refers to a learning experiment with a team of two heterogeneous Ibots. Ibot 1 (dashed line) translates four times as fast as Ibot 2 (solid line). In addition to that, Ibot 2 is blind: its ground color sensor reads “white” whatever its position in the arena.

The strategy discovered by the team to provide admissible and stable estimates is clear from the δ_{prog}^i s and α_{prog}^i s plots of Fig. 8. The blind Ibot minimizes its catastrophic contribution to the team integral by travelling long distances ($\delta_{prog}^2 = 2000$, having set $\delta_{max} = 2000$ for this particular experiment). Observe that the error E stabilizes to low values only when the difference between δ_{prog}^1 and δ_{prog}^2 is sufficiently large. Still, Ibot 1 can afford a parameter of $\delta_{prog}^1 = 800$ because it moves very fast.

Table 1 reports the control programs learnt by this team in 5 repeated experiments. Each Ibot program assumes different values for α_{prog} and δ_{prog} , but Ibot 1 always travels for shorter distances than Ibot 2. In

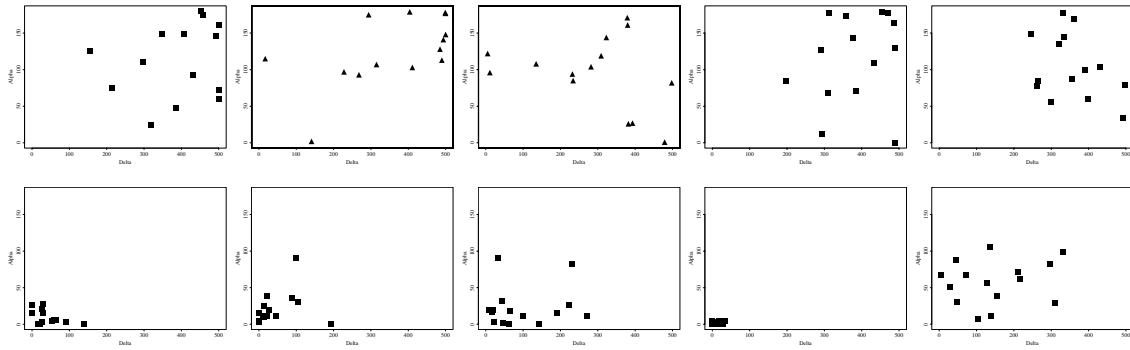


Figure 6: A team of 14 Ibots with private control programs on the “half-full” *Region*. Convergence points for 5 learning experiments for clustered configurations (first row) and for scattered configurations (second row). For each experiment, the 14 private programs learnt the team members are shown.




Figure 7: Robot credit assignment problem for a team of two Ibots: Ibot 1 (dashed line) and Ibot 2 (solid line). Time evolution of δ_{prog}^i (first plot), of α_{prog}^i (second plot), and of the error E (third plot).

δ_{prog}^1	α_{prog}^1	δ_{prog}^2	α_{prog}^2	$\mu(N_{sam}^1)$	$\sigma(N_{sam}^1)$
565	95	1528	179	85.5	1.5
876	95	2000	171	85.9	1.7
361	176	1549	2	88.9	1.3
537	165	1646	102	86.9	2.0
422	75	1061	66	84.1	0.5

Table 1: (Columns 1–4) Programs learnt by the heterogeneous team in 5 repeated experiments (one experiment per row). (Columns 5 and 6) Mean and standard deviation of the number of samples taken by Ibot 1 over 10 trials.

all experiments, the balance between δ_{prog}^1 and δ_{prog}^2 is such that Ibot 1 consistently manages to collect at least 85 samples out of the 100 available to the team.

Finally, figure 9 shows the behavior of $\mu_{prog}(E)$ in program space of Ibot 1, when Ibot 2 works with a fixed control program. On the left plot, Ibot 2 is running the program $Prog^2 = (0^\circ, 2000)$. This gives Ibot 1 the possibility of choosing its admissible program in a rather large set of programs: for example, it can easily run up to $\delta_{prog}^1 = 1000$ because its speed is four times the speed of the teammate. However, if the blind Ibot reduces its translations to $\delta_{prog}^2 = 1000$ (right plot), the fast Ibot is forced to fit its control program to a space of admissible programs which is considerably reduced.

Conclusions

The overall objective of the Ibots experiment was to understand how to use reinforcement learning to program automatically a team of robots with a common mission. In addition, we wanted to derive *real team solutions*. The key issue is: *to optimize team performance instead of individual performance*.

The “integration” mission of the Ibots is an artificial task for robots. However, the mission could also be interpreted as an exploration task, where the Ibots learn patterns of movement to reliably collect evidences about the region extension. Interestingly, it has been pointed to us (Martinoli 1996) that the Ibots resemble networks of patrolling ants engaged in the task of monitoring events occurring throughout their territory (Adler & Gordon 1992).

The learning scenario for the Ibots is applicable to other missions because it relies on very weak assumptions. There is a *team reinforcement signal* which qualifies the behavior of the group as a whole; a single Ibot has no direct way of assessing its own performance, as distinct from the performance of its teammates. There is a *limited, common resource* which constrains the

Ibots, and there is no a priori rule to decide how this resource should be shared. When working with private control programs, the Ibots are *unaware* of teammate programs; during learning, each Ibot changes its own program independently, and has no information on how teammates are changing theirs.

As a general conclusion, experiments have demonstrated how different mission conditions require completely different control programs, and that a very simple reinforcement learning procedure can find them by using a team payoff.

As far as the specific Ibots experiment is concerned, we cannot claim that the “general pattern” of the solutions discovered through learning were completely unexpected. However, as robot programmers, we have only a limited intuition for program parameters tailored to specific mission conditions (i.e. for a specific *Region*, for a given team size, or for a particular set of robot skills). Sometimes, human programmers are able to specify “a priori” programs which work for every possible mission condition (like those of Eq. 1); but, in certain mission contexts, these general solutions look unnatural. To write “ad hoc” programs for a robot team, a programmer will usually need to learn by trial-and-error himself (Kaelbling 1993). Therefore, why not consider letting the robots do this, i.e learn by trial-and-error on their own? *Second*, a program which is admissible for a single Ibot, in general is not ammissible for a team of Ibots, and viceversa. This means that we cannot simply find a solution for one Ibot and clone it n times, n being the number of team members. The form of the solution to a problem changes as the number of “problem solvers” changes. *Third*, the space of admissible programs strongly depends on the number of Ibots involved in the mission and on their initial configuration in the arena. In particular, the admissibility space “shrinks” (and learning requires more time) when the Ibots are started in the clustered configuration and the team size grows. The increasing difficulty is due to the fact that the number of samples is shared. On the contrary, the admissibility space “enlarges” (and learning requires less time) when the Ibots are started in the scattered configuration and the team size grows. The mission becomes easier in this case because, by initially distributing the Ibots at random in the arena, we are providing them with a strong bias towards the solution. *Fourth*, when the Ibots work with private control programs, the robot credit assignment problem arises clearly, resulting in longer learning times. Interestingly, the robot credit assignment problem forces the Ibots to learn admissible programs at a similar pace, to prevent “slow” learners to jeopardize the team mission. *Fifth*, the robot credit as-




Figure 8: A team of two heterogeneous Ibots: Ibot 1 (dashed line) moves 4 times faster than Ibot 2 (solid line), which is blind. Time evolution of δ_{prog}^i s (first plot), of α_{prog}^i s (second plot), and of the error E (third plot).

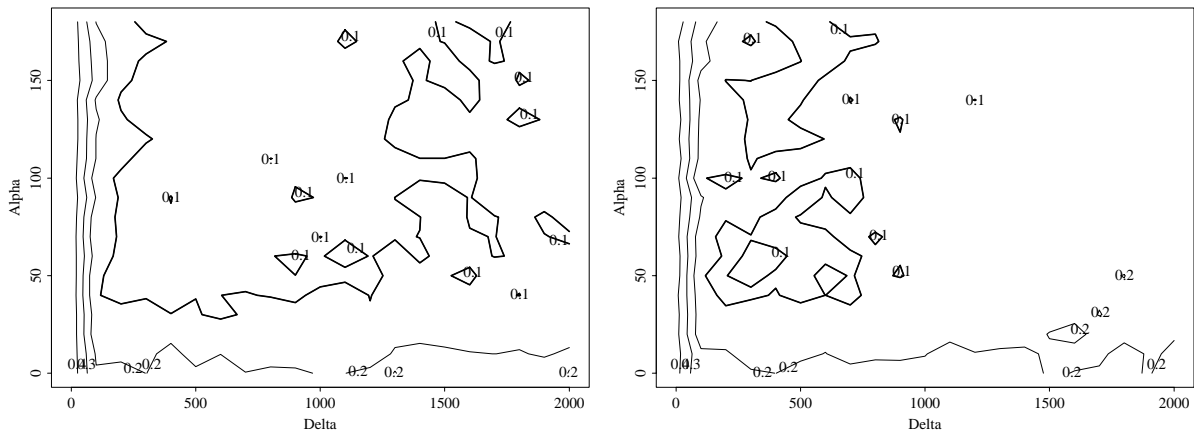


Figure 9: Contour plots of $\mu_{prog}(E)$ in program space of Ibot 1 when Ibot's 2 program is: (Left) $Prog^2 = (0^\circ, 2000)$ or (right) $Prog^2 = (0^\circ, 1000)$.

signment problem completely vanishes when the Ibots learn a shared policy, and the learnt policy is still a real team solution. The possibility of learning a single public program instead of several private programs should be not overlooked in missions where specialization of the robots behavior is not required, because the time necessary for the team to learn a public program is much shorter. Finally, *sixth*, the Ibots with their heterogeneous acting and sensing capabilities manage to specialize their control programs so as to take advantage of their skills and to minimize the impact of their weaknesses.

Acknowledgements

Cristina Versino is supported by project No. 2129-042413.94/1 of the Fonds National de la Recherche Scientifique, Berne, Suisse. We thank Jean-Bernard Billeter, Alcherio Martinoli, José del R. Millán, and Francesco Mondada for interesting discussions on collective robotics. Thanks to Marco Dorigo, Rafał Sałustowicz, Jürgen Schmidhuber, Marco Wiering, Jieyu Zhao for their comments on earlier versions of this paper.

References

Adler, F., and Gordon, D. 1992. Information collection and spread by networks of patrolling ants. *The American Naturalist* 140(3):373–400.

Barto, A.; Sutton, R.; and Brouwer, P. 1981. Associative search network: A reinforcement learning associative memory. *Biological Cybernetics* 40:201–211.

Barto, A.; Sutton, R.; and Watkins, C. 1989. Learning and sequential decision making. Technical Report COINS-89-95, Dept. of Computer and Information Science, University of Massachusetts, Amherst.

Crites, R., and Barto, A. 1996. Improving elevator performance using reinforcement learning. In Touretzky, D.; Mozer, M.; and Hasselmo, M., eds., *Advances in Neural Information Processing Systems 8*, 1017–1023. Cambridge MA: MIT Press.

Deneubourg, J. C.; Goss, S.; Franks, N.; Sendova-Franks, A.; Detrain, C.; and Chrétien, L. 1991. The dynamics of collective sorting: Robot-like ants and ant-like robots. In J.A., M., and Wilson, S., eds., *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, 356–363. MIT Press.

Hammersley, J., and Handscomb, D. 1964. *Monte Carlo Methods*. London: Methen & Co.

Kaelbling, L. 1993. *Learning in Embedded Systems*. Cambridge, MA: MIT Press.

Kube, C., and Zhang, H. 1994. Collective robotics: From social insects to robots. *Adaptive Behavior* 2(2):189–218.

Martinoli, A. 1996. Personal Communication.

Mataric, M.; Nilsson, M.; and Simsarian, K. 1995. Cooperative multi-robot box-pushing. In *IROS95, IEEE/RSJ International Conference on Intelligent Robots and Systems, Pittsburgh, PA, August*.

Mataric, M. 1994a. Interaction and intelligent behavior. Technical Report AI-TR-1495, MIT Artificial Intelligence Lab, Boston.

Mataric, M. 1994b. Learning to behave socially. In J.A., M., and Wilson, S., eds., *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, 453–462. Cambridge, MA: MIT Press.

Mataric, M. 1995. Personal Communication.

Mataric, M. 1996. Learning in multi-robot systems. In Weiss, G., and Sen, S., eds., *Adaptation and Learning in Multi-Agent Systems*, volume 1042. Springer-Verlag, Lecture Notes in Artificial Intelligence. 152–163.

Murciano, A., and Millán, J. d. R. 1996. Learning signaling behaviors and specialization in cooperative agents. *To appear in: Adaptive Behavior* 5(1).

Parker, L. 1994. ALLIANCE: An architecture for fault tolerant, cooperative control of heterogeneous mobile robots. In *IROS94, IEEE/RSJ International Conference on Intelligent Robots and Systems, Munich, Germany, September*, 776–783.

Parker, L. 1995a. The effect of action recognition and robot awareness in cooperative robotic teams. In *IROS95, IEEE/RSJ International Conference on Intelligent Robots and Systems, Pittsburgh, PA, August*, volume 1, 212–219.

Parker, L. 1995b. L-ALLIANCE: A mechanism for adaptive action selection in heterogeneous multi-robot teams. Technical Report ORNL/TM-13000, Oak Ridge National Laboratory, Tennessee.

Schmidhuber, J. 1996. A general method for incremental self-improvement and multi-agent learning in unrestricted environments. In Yao, X., ed., *Evolutionary Computation: Theory and Applications*. Scientific Publ. Co., Singapore.

Tucker, A. 1995. *Applied Combinatorics*. New York: John Wiley & Sons.