DOI: http://dx.doi.org/10.7551/978-0-262-32621-6-ch136

Real-time Evolution of iAnt Robot Foraging Strategies

Joshua P. Hecker[1] and Melanie E. Moses[1,2,3]

[1]Department of Computer Science, University of New Mexico, Albuquerque, NM 87131
[2]Department of Biology, University of New Mexico, Albuquerque, NM 87131
[3]External Faculty, Santa Fe Institute, Santa Fe, NM 87501
{jhecker,melaniem}@cs.unm.edu

Introduction

Central-place foraging is a canonical task in swarm robotics. For this task, robots are programmed to search an area for resources and aggregate these resources at a central location. Foraging can be instantiated in a number of real-world applications, such as hazardous waste clean-up, search and rescue, and in-situ resource utilization.

We extend our prior work by using coevolutionary methods to evolve iAnt robot swarm foraging strategies in real robots and in real time. Each robot maintains a private agent-based simulation, which models the physical environment according to the robot's own approximation of the real-world resource distribution. The robot uses a private genetic algorithm (GA) to evolve parameters for a central-place foraging algorithm (CPFA) which maximize the foraging success of the simulated agents; it then employs these evolved parameters as a foraging strategy in the real world.

In addition to evolving individual foraging strategies, the entire swarm executes a distributed GA to evolve a population of resource distribution approximations that is shared across all robots. In this way, coevolution evolves both the foraging strategy and the resource distribution approximation for each robot. Over time, this allows the swarm to adapt its behavior to previously unknown environments.

Background

A key challenge in evolutionary robotics is narrowing the reality gap between simulated agents and real robots. Bongard and Lipson (2004) address this by iteratively adapting robots and robot simulators in real time to resolve brain-body discrepancies and thus generate robust machines. O'Dowd et al. (2011) use distributed coevolution to coevolve their foraging strategy in conjunction with their simulated world model in order to adapt to different simple foraging tasks.

Our approach differs from this previous work in that we use interdependent foraging strategies, memory, and communication to accomplish more complex foraging tasks. In this paper, we adapt foraging behaviors to maximize foraging success on the specific resource distribution that the robots encounter in each foraging experiment.

Methods

The CPFA implements robot foraging behaviors as a series of states connected by directed edges with transition probabilities (Fig. 1(a)). Each robot begins its search at a central nest site and sets a search location. Robots traveling to a random location with no prior information search using an uninformed correlated random walk. Robots traveling to a previously found resource location via memory or communication search using an informed random walk that is initially undirected and localized, then becomes more directed and straighter over time. When a robot locates a resource, it first collects the resource, and then records a count of resources in the neighborhood of the found resource. Robots use this count to decide whether to exploit information through memory or communication. Robots who have not found a resource will probabilistically return to the nest.

In prior work, we used a GA to evolve a population of CPFA parameters that maximized the foraging efficiency of simulated robot swarms evaluated in an agent-based model (Hecker et al., 2013). These parameters control the sensitivity threshold for triggering CPFA behaviors, the likelihood of transitioning from one behavior to another, and the length of time each behavior should last. The fitness function of the GA is the collective foraging success (number of resources collected in fixed time). Each parameter set is evaluated with an identical copy running on each robot. We demonstrated that such "group selection" evolved successful foraging strategies for each resource distribution, and that robot swarms were most efficient when using the specialist strategy adapted for a given distribution. However, the approach required *a priori* knowledge of the resource distribution, and it did not evolve in real time.

In this work, we simulate individual robots, each running a private version of the GA in real time while foraging given a particular unspecified resource distribution (Fig. 1(b), point 1). The parameter set evolved by the private GA is then transferred to the robot and evaluated for fitness (Fig. 1(b), point 2). Each robot defines its fitness as the number of actual resources it individually collects during the experiment (Fig. 1(b), point 3). Robots periodically communicate

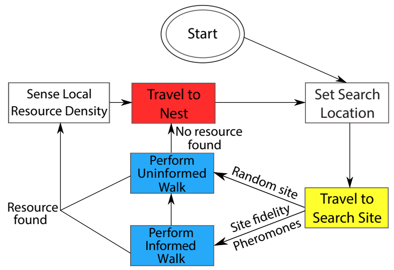

(a) CPFA state diagram

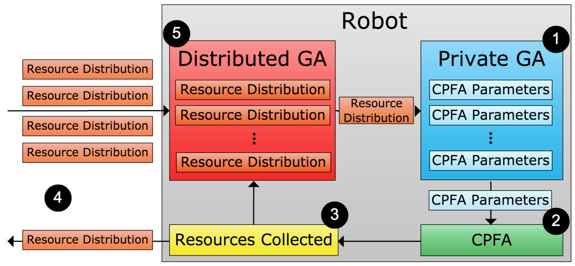

(b) Coevolution state diagram

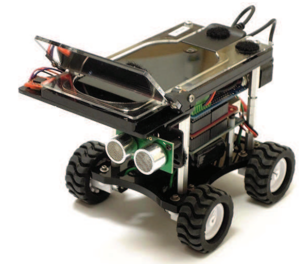

(c) iAnt robot

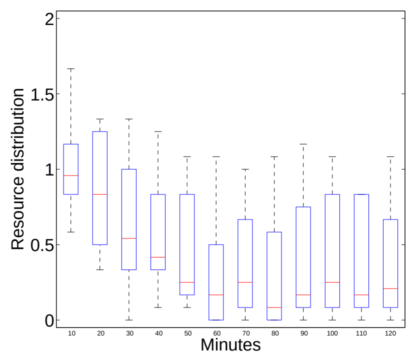

(d) Clustered

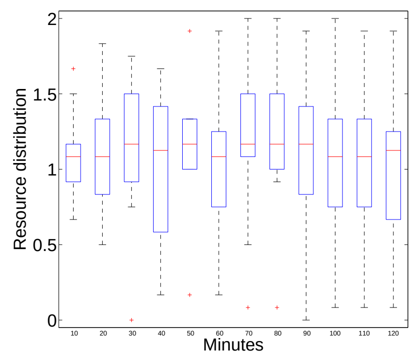

(e) Power law

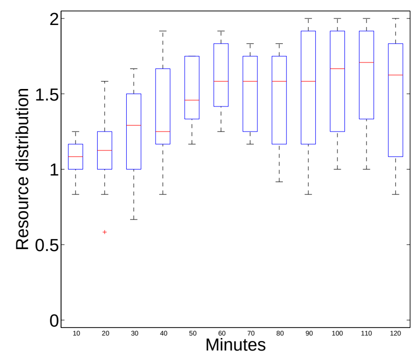

(f) Random

Figure 1: Top: Diagrams explaining implementation of (a) CPFA and (b) coevolution, as well as (c) an iAnt robot. Bottom: Resource distribution approximation values for swarms foraging on (d) clustered, (e) power law, and (f) random distributions.

their current foraging efficiency and simulation resource description to the entire swarm via wireless transmission (Fig. 1(b), point 4). Each robot evolves a distinct subpopulation of simulated resource placements (Fig. 1(b), point 5). In this way, the swarm's distributed GA maximizes the correspondence between agent-based simulations and the resource distribution by selecting for simulations which produce more efficient foraging strategies in real robots (Fig. 1(c)).

Preliminary Experiments

We conduct preliminary experiments with simulated robot swarms in an agent-based model. Swarms of 12 robot agents forage for 256 resources placed on a 125 x 125 cellular grid, simulating a 100 m^2 physical area over 120 minutes. The resources are arranged in one of three distributions: clustered (4 randomly placed clusters of 64 resources each), power law (1 large cluster of 64, 4 medium clusters of 16, 16 small clusters of 4, and 64 randomly scattered), or random (each resource placed at a random location).

Each robot's private simulation (Fig. 1(b), point 1) is randomly initialized with one of the three distributions. The CPFA parameters of each swarm in the private simulation are randomly and independently initialized; agents within a swarm use identical parameters. The robot's private GA evolves a population of 100 simulated swarms over 20 generations. When the private simulation is complete (approximately 10 minutes), the robot transfers the most efficient CPFA parameters from the simulation to its own CPFA (Fig. 1(b), point 2). Each robot then communicates to the entire swarm its current foraging efficiency, along with a value representing the current resource distribution used in its private simulation (clustered, power law, or random).

Figures 1(d)–(f) show the resource distribution approximation values for the entire swarm at 10 minute increments across 10 replicates. The swarm converges on the true distribution in each experiment (clustered = 0, power law = 1, random = 2) within approximately 60 minutes. Adapting foraging behavior to a particular resource distribution in real time is 2.7 times more efficient for clustered resources, and 1.1 times more efficient for random resources, compared to using a generalist fixed strategy evolved *a priori* for power-law-distributed resources. Experiments replicating these observations in real iAnt robots are ongoing.

References

Bongard, J. C. and Lipson, H. (2004). Once More Unto the Breach: Co-evolving a robot and its simulator. In *ALIFE9*, pp. 57–62.

Hecker, J. P., Stolleis, K., Swenson, B., Letendre, K., and Moses, M. E. (2013). Evolving Error Tolerance in Biologically-Inspired iAnt Robots. In *ECAL 2013*, pp. 1025–1032.

O'Dowd, P. J., Winfield, A. F., and Studley, M. (2011). The distributed co-evolution of an embodied simulator and controller for swarm robot behaviours. In *IROS 2011*, pp. 4995–5000.