

Romero’s Pilgrimage to Santa Fe: A Tale of Robot Evolution

Christof Teuscher Eduardo Sanchez Moshe Sipper

Logic Systems Laboratory, Swiss Federal Institute of Technology,
CH-1015 Lausanne, Switzerland, E-Mail: {name.surname}@epfl.ch, URL: <http://lslwww.epfl.ch>

ABSTRACT

Many a researcher in evolutionary robotics would love to have at his or her disposal a real population of bopping critters. We describe the design of a cheap, modular, and adaptable robot and the use of our robot in a real-world version of the Santa Fe trail artificial ant navigation task.

1 A real robot population for adaptive robotics

To date, experiments with a population of dozens of robots have been mostly (if not wholly) relegated to the land of simulation. The ultimate goal of our work is to construct a large population of bona fide robots, and to use it to perform experiments in the real world. Using several real robots in an evolutionary scenario can markedly speed up the process with respect to the classic approach, wherein one has but one robot, used to test sequentially all individuals of a given generation. Moreover, a real population presents multitudinous opportunities for investigating questions pertaining to interacting robots (e.g., coevolution, cooperation, and competition). But, a population of robots might be too costly.

We found that LEGOTM components are particularly suitable for building cheap and modular robots. However, the LEGOTM RCX (Robot Command System) was too expensive for our purposes. We built our own controller—the EvoMaster—which interfaces with LEGO sensors and motors, offers an open architecture for our own specific extensions in the future, is small and cheap, and exactly meets our demands.

2 Reality versus Simulation

Most roboticists use simulations to test and validate their models. The advantage of simulation over training or evolving robots in a real environment is that simulation is much faster. However, a simulated environment is often strongly simplified compared to a real environment. Designing a simulative environment model of a real robot and its environment is often very hard (Miglino et al. 1995): sensors deliver uncertain responses,

apparently identical physical sensors are often significantly different, and not all physical laws governing the interaction between the robot and its environment are taken into account. The problem may even become much harder if the environment underlies dynamic changes.

3 Romero’s pilgrimage to Santa Fe

Using the EvoMaster controller we have built a simple robot dubbed Romero: it is cheap (costing under \$100), uses modular LEGO components, and can be easily adapted to new tasks and novel environments. To test Romero in a real-world environment, we chose to evolve a solution to the Santa Fe trail artificial ant navigation task (Koza 1992). Several researchers have used this problem as a benchmark for evolutionary experiments. Applying genetic algorithms, (Collins and Jefferson 1992) and (Jefferson et al. 1992) evolved finite state machines that control the ant’s behavior, while (Koza 1992) approached the problem in a more direct way using genetic programming. These works all involve simulated environments; indeed, to the best of our knowledge, no one has studied the Santa Fe trail ant navigation task in a real-world setup.



Figure 1 The experimental setup.

The LSL Santa Fe trail (Figure 1) is a real-world version for use with actual robots. We added a special track enclosing the environment, which is used to guide the robot back to the starting point; furthermore, the 32×32 grid is not toroidal, as with previous simulated setups. The proportion between gaps and trail has been respected. As we are currently in possession of two robots, we used two trails printed on A0 paper, i.e., of size 90cm x 120cm (Figure 1).

A number of differences between our real-world setup and previous simulated environments should be noted.

The Romero robot has—apart from the two ground-facing light sensors—no other navigation or localization system. Moreover, as opposed to the simulated worlds, the grid markings (i.e., the cells) have been removed; thus, Romero cannot move from cell to cell—as these latter do not exist. In the simulations, the food pellets were removed when the ant passed over them, whereas in our case this is not possible (the pellets form part of the printed environment). One can thus view our setup as being more akin to a pheromone trail than to a food trail.

When working in the real world, total predictability is no longer possible—nature is noisy. One way to tackle this problem is to try to “clean up” the (real) environment to some extent, rendering it more predictable. For example, in several works the exact position of the robot is kept track of at all times (Thrun 1998, Roy and Thrun 1999, Hong and Lee 1997). Position calculation can, however, be quite costly, and it is often inaccurate. We opted for a simpler approach: Romero has no idea where it is, or—to be more precise—it does not possess its absolute position in the environment.

4 The evolutionary experiment

Our first approach uses genetic programming with asynchronous local tournament selection to evolve a solution. Romero’s genome is defined as a sequence of 30 “codons” (evolutionary building blocks), e.g., --NNF--NFL--NNN--FFR-- where F (Forward), L (Left), R (Right), and N (NOP) are elemental movements. By grouping together (up to) three elemental movements, the search space is markedly reduced (in comparison to an encoding which allows for single-action codons). The genome is treated as a program that is read and executed in a linear fashion. Because the robot is not able to count the number of food pellets amassed, we defined the fitness as the time spent over food “cells”.

The local tournaments take place asynchronously (in a manner reminiscent of steady-state genetic algorithms): each robot, upon terminating an evaluation run (either due to a timeout or when the genome program has been executed), independently sends its calculated fitness along with its genome to its two neighbors—which transmit their own fitness values and genomes in turn. Crossover then takes place between the genome of the higher-fitness neighbor and the robot’s own genome. The genome thus formed replaces the robot’s old genome.

Our tale of robot evolution has not (yet!) reached its happy end. The robots were subjected to a few days’ worth of experimenting on the real-world LSL Santa Fe trail. The two most successful robots evolved, amass about 75% of the food “pellets.” Evolution in the real world is time-consuming (evaluating one 50-robot generation takes approximately 40 minutes). Obviously, we are far from being the first to stumble upon this mind-

boggling conclusion. Indeed, we had quite expected this—which is why we wish to have Romero go forth and multiply.

5 Conclusions

This paper related our first steps in the direction of an all-real robot population for artificial-life researchers. We tested our robot on a real-world version of the Santa Fe trail, obtaining good (though not perfect) solutions, despite the increased difficulty with respect to the simulated version. There are several avenues of future research that suggest themselves at this point. One experiment which we wish to carry out in the near future involves the cooperation of two robots: both are to traverse the Santa Fe trail in line formation. Other planned experiments include evolvable hardware (i.e., FPGA circuits), self-organization, coevolution, cooperative behavior, competitive behavior, and predator-prey interactions.

References

- Collins, R. J. and D. R. Jefferson (1992). AntFarm: Towards simulated evolution. In: *Artificial Life II* (C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen, Eds.). Vol. X of SFI Studies in the Sciences of Complexity. Addison-Wesley. Redwood City, CA. pp. 579–601.
- Hong, S.-G. and J.-J. Lee (1997). A local motion planner for car-like robots in a cluttered environment. *Artificial Life and Robotics* 1, 39–42.
- Jefferson, D., R. Collins, C. Cooper, M. Dyer, M. Flowers, R. Korf, C. Taylor and A. Wang (1992). Evolution as a theme in Artificial Life: The Genesys/Tracker system. In: *Artificial Life II* (C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen, Eds.). Vol. X of SFI Studies in the Sciences of Complexity. Addison-Wesley. Redwood City, CA. pp. 549–578.
- Koza, J. R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. The MIT Press. Cambridge, Massachusetts.
- Miglino, O., H. H. Lund and S. Nolfi (1995). Evolving mobile robots in simulated and real environments. *Artificial Life* 2(4), 417–434.
- Roy, N. and S. Thrun (1999). Online self-calibration for mobile robots. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.
- Thrun, S. (1998). Bayesian landmark learning for mobile robot localization. *Machine Learning* 33(1).