Distributed, Evolutionary Learning on Android-powered COTS Robots

Terence Soule, Robert B. Heckendorn
Computer Science Department, University of Idaho, Moscow, ID, 83844-1010
tsoule@cs.uidaho.edu, 1-208-885-7789, heckendo@uidaho.edu,

Abstract

Recent improvements in autonomous robots' capacities have led to robots capable of solving a wide range of problems and operating in diverse environments. However, the most significant successes remain limited to autonomous robots operating within specific pre-defined environments, solving specific problems, and using large, expensive, highly instrumented robots, engineered for specific tasks. Currently there is neither a widely used, sophisticated, and affordable research robot, nor a reliable, widely used method for non-experts to train robots.

We propose combining several complementary novel hardware and software solutions to solve these problems. We use a robot platform we developed with Commodity Off The Shelf (COTS) principles to provide low-cost, connected, peripheral rich, and computationally powerful robots. We use simple learning methods, such as learning by demonstration, to allow non-experts to train robots. We will implement evolutionary "learning in the cloud" to allow many robots to contribute to the learning process via the Internet. This will allow anyone to participate in the distributed robot training. Our ultimate goal is an electronic library of learned behaviors that continues to improve by group participation, much like Wikipedia does.

Problem Statement For autonomous robots to become a part of our daily lives it must be possible to program, or re-program, them quickly and easily with robust, reliable behaviors. I.e. the major problem with modern robotics is that a novice user cannot quickly and reliably teach a robot to solve a novel problem and/or to operate in a novel environment. In addition, a robot for the general public should be comparatively inexpensive, have the computational power to carry out the demanding computation of learning novel behaviors and yet not require special expertise to maintain. Note that these requirements: low-cost, computational power, and ease of maintenance, are essentially the same as for modern smart phones.

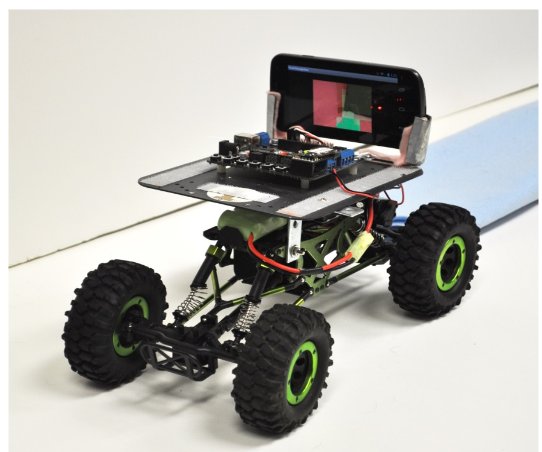

Figure 1: A COTSbot is a solderless assembly of an Android phone, an Arduino microcontroller, and an RC "rock crawler" car chassis. It is capable of several hours of autonomous driving in rough terrain, has a variety of built-in sensors, and significant computational power for on-board image processing and learning.

Project Goals Our goal for this project is to develop a *distributed* robot training framework that non-experts can use to cooperatively train novel and robust robot behaviors. First, we will use a robot platform we have developed that uses a Commodity Off The Shelf (COTS) design, which provides low-cost, connected, peripheral rich, and computationally powerful robots. Second, we will use simple learning techniques, such as learning by demonstration [2]. Third, we will use evolutionary techniques for the learning process. Finally, we will implement distributed learning, i.e. "learning in the cloud", which will allow many robots to share what they learn via the Internet to improve robustness and shorten training time.

The project's outcome will be a framework for group learning by smart phone robots using multiple, geographically distributed trainers. The robots will use on-board, distributed evolutionary computation to learn demonstrated behaviors. Multiple trainers acting in multiple environments will quickly generate large sets of diverse training data. This rich training data coupled with learning distributed across multiple robots will allow the rapid evolution of general, robust behaviors.

The project's largest risk is that training data from such diverse sources will overwhelm the learning process. Expanding the learning process by distributing it should avoid this problem. If necessary we will apply clustering and outlier detection algorithms to refine the training data.

Commodity Off-The-Shelf (COTS) design advocates the use of low-cost, mass produced products to build powerful, novel devices. E.g. the assembly of commodity PCs to make the first cluster computers. COTS leverages mass production and competition to achieve low cost and a rich set of desirable features, such as durability, reliability, integrated peripherals, and connectedness. We have used COTS approaches to combine Android smart phones with commercial RC cars to create our robots (COTSBots) [6]. Because COTSBots use a familiar product, smart phones, they are approachable and easy to use. Thus, another benefit of this project should be a significant increase in the number of people participating in cutting edge research with autonomous robots.

Evolutionary Computation (EC) is a stochastic optimization technique based on the principles of natural evolution. Analogs of biological reproduction and selection are used to evolve solutions to problems [1]. An evolutionary algorithm consists of four fundamental elements: a population, a fitness function, a selection function, and one or more genetic operators. The population is a set of individuals that represent potential solutions. The fitness function is used to measure the quality of the solutions. The selection function is used to select solutions with better fitness, i.e. survival of the fittest. Finally, the genetic operators are used to generate new solutions by recombining and mutating the selected solutions. EC has proven to be capable of learning robust robotic behaviors from very large, very noisy data sets (see for example [3]). However, it is a computationally intensive learning technique requiring large training sets and computational resources significantly beyond those found on other research robots such as ePucks, webots, and kheperas.

Distributed learning distributes the evolutionary process across multiple robots. Research shows it improves on-board evolution of robotic controllers [4]. Our robots will use the smart phones' built-in communication to implement distributed evolution and to share training data collected from each trainer. Each robot maintains its own evolving population of behaviors, but will share the most promising, i.e. highest fitness, behaviors with the other robots. Multiple trainers working with multiple robots in varying environments rapidly creates rich training sets supporting robust behaviors and distributed evolution allows successful behaviors to evolve much more rapidly.

With support from BEACON (an NSF funded Science and Technology Center for the study of evolution in action http://beacon-center.org/) we have built COTSBots (Figure 1) which can be trained via demonstration and EC to follow a color or path [7]. For example, 5 minutes of human driving followed by 15 minutes of evolution time is sufficient to evolve a neural network capable of reliably identifying and following a distinctly colored path. We have code for direct inter-bot communication using either Bluetooth or wireless communication and for indirect communication through a web server. Thus, the only remaining task is to implement the current evolutionary algorithm in a distributed fashion using the established communication protocols and EC techniques.

Our test problem will be learning to drive on outdoor, wilderness trails. Wilderness trails are poorly defined, with highly variable surfaces and edges. Thus, this problem will test our system's ability to generate, and reliably learn from, large, diverse training sets. In addition, this problem has a practical component, using the phone's built in GPS we will record the robots' movements allowing us to autonomously build GPS maps of wilderness trails. Outdoor activities and industries are an important aspect of the Idaho economy and being able to autonomously create GPS maps of wilderness regions will be a significant benefit to the state. Other applications of this problem include search and rescue, autonomous wilderness monitoring, and industrial security.

Learning Process The initial learning process has two stages: human operation (demonstration) and learning. During human operation, multiple operators drive multiple robots, following multiple trails, and the robots build a training set of terrain images and behaviors. During learning distributed EC is used to train neural networks to generate the correct behaviors for a given image. Each robot maintains its own evolving population. Promising networks (i.e. ones with high fitness) are exchanged between robots to increase diversity and improve the evolutionary process. When evolved behaviors match the human decisions the robots can operate autonomously.

Once testing is complete with the discrete approach the training and learning steps will be merged. EC will run continuously in the background during human operation. Additionally, during autonomous operation, if a robot seems to be performing poorly or moves into a novel environment, a human operator can resume control to mentor the bot.

Outcomes This project will lead to the development of hardware and software for simple, fast, and reliable training of robust robotic behaviors, which will make robots much more accessible and useful to the general public. Although our research is involves small sets of robots (5-10), the process is distributable across thousands of robots. Networks of users facing similar problems can share training data and evolved solutions creating learning populations. In addition, this work will further demonstrate the usefulness of the COTS approach to robot design and will make low cost computationally powerful robots widely available.

DATA PLAN Our COTSBots designs, build instructions, and sample code are available through our website at cotsbots.wordpress.org and via Bitbucket under the project name UIdahoLair. We will continue to update both sites as the research progresses. We will publish scientific results in academic venues and will attend machine learning conferences to present tutorials on building and using COTSBots. Dr. Soule presented such a tutorial at GECCO 2013 [5].

BUDGET (Total: $49,416)
PhD Student Salary ($25/hr, 40hrs/week summer, 20hrs/week academic year + fringe): $31,685
Student in-state tuition (full time RA's receive an out-of-state tuition waiver): $10,311
Android phones (10 Nexus 4's or similar): $3,000
Robots (12 nitro RC cars ($120) + 12 DFRobot Romeo microcontrollers ($40)): $1,920
Additional supplies (Batteries, chargers, etc.): $1000
Travel (attend 4 conferences to present results and tutorials): $1500

References

[1] J. A. Foster. Evolutionary computation. *Nature Reviews Genetics*, 2(6):428–436, 2001.

[2] E. Guizzo and E. Ackerman. The rise of the robot worker. *IEEE Spectrum*, 49(10):34–41, 2012.

[3] I. Harvey, E. Di Paolo, R. Wood, M. Quinn, E. Tuci, and E. T. Iridia. Evolutionary robotics: A new scientific tool for studying cognition. *Artificial life*, 11(1-2):79–98, 2005.

[4] R.-J. Huijsman, E. Haasdijk, and A. Eiben. An on-line on-board distributed algorithm for evolutionary robotics. In *Artificial Evolution*, pages 73–84. Springer, 2012.

[5] T. Soule. Designing and building powerful, inexpensive robots for evolutionary research. In *Proceeding of the 15th annual GECCO conference companion*, pages 919–934. ACM, 2013.

[6] T. Soule and R. B. Heckendorn. Cotsbots: computationally powerful, low-cost robots for computer science curriculums. *Journal of Computing Sciences in Colleges*, 27(1):180–187, 2011.

[7] T. Soule and R. B. Heckendorn. A practical platform for on-line genetic programming for robotics. In *Genetic Programming Theory and Practice X*, pages 15–29. Springer, 2013.

Curriculum Vitae

Terence Soule

Professional Preparation

Reed College	Physics	B.A., 1991
Washington State University	Physics	M.S., 1994
University of Idaho	Computer Science	Ph.D., 1998

Appointments

Director, Neuroscience Program, University of Idaho, 2009 - 2013

Associate Professor, Computer Science, University of Idaho, 2006 – present

Adjunct Professor, Bioinformatics and Computational Biology, University of Idaho, 2004 - present

Assistant Professor, Computer Science, University of Idaho, 2000 - 2006

Assistant Professor, Computer Science, St. Cloud State University, 1998 – 2000

Recent Relevant Publications

1. T. Soule and R. B. Heckendorn. "A practical platform for on-line genetic programming for robotics." *Genetic Programming Theory and Practice X*, pages 15-29. Springer, 2013.
2. *Leidenfrost, H. T.,, *Tate, T. T., Canning, R. J., Anderson, M. J., Soule, T., Edwards, D. B., Frenzel, J. F., "Autonomous Navigation of Forest Trails by an Industrial-Sized Robot" Transactions of the ASABE, in Press.
3. Soule, T. "A Projected Based Introduction to C++", ISBN 978-1-4652-1328-0, KendallHunt, 2012 (Text book).
4. Soule, T. and Heckendorn, R. "COTSBots: computationally powerful, low-cost robots for Computer Science curriculums" Journal of Computing Sciences in Colleges, 27:1, 180-187, 2011.
5. *Solomon, Michael and Soule, Terence and Heckendorn, Robert B., "A comparison of a communication strategies in cooperative learning", "Proceedings of the Fourteenth International Conference on Genetic and Evolutionary Computation Conference", pp. 153-160, ACM, 2012. (Best paper award)
6. Soule, Terence, "Evolutionary Dynamics of Tag Mediated Cooperation with Multilevel Selection" *Evolutionary Computation*, 19:1, MIT Press, 2011.
7. Soule, T. and Heckendorn, R. B., "Improving Performance and Cooperation in Multi-Agent Systems," *Genetic Programming, Theory and Practice V*, Chapter 13, Springer, 2007.
8. Hallin, J., H. Egbo, P. Ray, T. Soule, M. O'Rourke, D. Edwards, "Enabling Unmanned Underwater Vehicles to Reason Hypothetically," Proceedings of Oceans '09 MTS/IEEE Biloxi, Biloxi, Mississippi, October 26-29, 2009.
9. *Thomason, Russell, Heckendorn, R. B., and Soule, T., "Training Time and Team Composition Robustness in Evolved Multi-agent Systems", Genetic Programming: Proceedings of the 6th European Conference on Genetic Programming, EuroGP 2008, Lecture Notes in Computer Science 4971, Springer-Verlag, ppg. 1-12, 2008.
10. *Gotshall, Stan, Browder, K., *Sampson, J., Wells, R., and Soule, T." Stochastic Optimization of a Biologically Plausible Spino-neuromuscular System Model - A Comparison with Human Subjects", Genetic Prgoramming and Evolvable Machines, Special Issue on Medical Applications, Vol 8:4, pg. 355-380, 2007.

11. *Katarzyna H. Kucharzyk, Ronald L. Crawford, Andrzej J. Paszczynski, Terence Soule, Thomas F. Hess, "Maximizing microbial degradation of perchlorate using a genetic algorithm: Media optimization ", Journal of Biotechnology, Volume 157, Issue 1, January 2012, Pages 189-197
12. Edlund , A., Soule, T., Sjöling, S. and Jansson, J.K. "Microbial community structure in Baltic Sea sediment along a pollution gradient," Microbial Ecology, V8:2, pp. 223-232, February, 2005.

Synergistic Activities

- Member, SIGEVO Executive Board (ACM Special Interest Group for Genetic and Evolutionary Computation), 2013 to present
- Editor in Chief, Genetic and Evolutionary Computation Conference (GECCO) 2012.
- Co-Chair, Genetic Programming: Theory and Practice workshop VI, Ann Arbor, MI, 2006 - 2008
- Director of the Neuroscience Program, University of Idaho, August 2009 - August 2012
- Member of the Education and Human Resource Development steering committee and Education and Human Resource Development lead for the University of Idaho for the BEACON NSF Science and Technology Center for the study of *Evolution in Action*.
- Appointment Bioinformatics and Computational Biology Programs – University of Idaho
- Co-chair, European Conference on Genetic Programming, 6th and 7th Annual Conferences, 2003, 2004.

Recent Grants and Contracts

- "Digital Innovation Generating New Information Technology (Dig' n IT)", Micron, $14,550, August 2013-May, 2014, (PIs: Julie Amador, Terence Soule).
- "Distributed, Onboard Evolution in a Robotic Cloud", BEACON NSF Center for the Study of Evolution in Action, $168,231, June, 2013 - May 2014, (Lead PI: Soule, T., Co-PIs: Heckendorn, R.,McKinley, P., Zhan, J., Harrison, S.), (Internal BEACON Award).
- "Evolution Curriculum for Elementary Classrooms: Implementation and Assessment of LadyBug and Supporting Activities," $87,696, June, 2012-May, 2013, (PIs: Thomas Getty (MSU) and Terence Soule (UI)).
- "Evolutionary Games for K-6th," BEACON NSF Center for the Study of Evolution in Action, $11,108, June, 2011, May 2012, (PI: Soule, T. (Internal BEACON Award)).
- "Algorithmic Improvement of Calls and Reads in 454 Sequencing", Center of Biomedical Research Excellence (COBRE), $33,287,May 1, 2010 through April 30, 2011, (PIs: Soule, T. and Heckendorn, R. (Internal Award).)
- "Magnetic Signature Assessment System using Multiple Autonomous Underwater Vehicles (AUVs) Phase 3", Office of Naval Research, $1,111,600, 2009-201, (PI: Dean Edwards, Investigators: Soule, et al.)

Biographical Sketch for Robert B. Heckendorn

Professional Preparation

The University of Oklahoma	Mathematics	B.A.	1977
University of Arizona	Computer Science	M.S.	1979
Colorado State University	Computer Science	Ph.D.	1999

Appointments

2006-present	Associate Professor, Computer Science, University of Idaho
1999-present	Adjunct Professor, Bioinformatics and Computational Biology, University of Idaho
1999-2006	Assistant Professor, Computer Science, University of Idaho
1995-1999	Research Assistant, Computer Science, Colorado State University
1979-1995	Research and Development Engineer, Hewlett-Packard Company

Products

Publications most closely related to the proposed project

- Terence Soule, Robert B. Heckendorn, **COTSBots: Computationally Powerful, Low-Cost Robots for Computer Science Curriculums**, *Journal of Computing Sciences in Colleges*, 27(1), ACM, 2011
- Solomon, Michael and Heckendorn, Robert B and Soule, Terence, **A Comparison of a Communication Strategies in Cooperative Learning**. *GECCO 2012: Proceedings of the Genetic and Evolutionary Computation Conference*, pp153-160, ACM, 2012
- Soule, T. and Heckendorn, Robert B, **Developmental Scalable Hierarchies for Multi-agent Swarms**. *GECCO 2011: Proceedings of the Genetic and Evolutionary Computation Conference*, pp207-208, ACM, 2011.
- Soule, T. and Heckendorn, R.B, **A Developmental Algorithm for Multi-agent Swarms with Scalable Hierarchies**. *GECCO 2010: Proceedings of the Genetic and Evolutionary Computation Conference*, pp647-648. ACM, 2010.
- Joshua Rubini, Robert B. Heckendorn, and Terence Soule, **Evolution of Team Composition in Multi-Agent Systems**, *GECCO 2009: Proceedings of the Genetic and Evolutionary Computation Conference*, ACM Press, New York, NY, pp 1075-1082, ACM, 2009

Other significant publications

- Russell Thomason, Robert B. Heckendorn, and Terence Soule, **Training Time and Team Composition Robustness in Evolved Multi-Agent Systems**, *Genetic Programming, Proceedings of the 11th European Conference, EuroGP 2008*, pp 1-12, 2008
- Terence Soule, and Robert B. Heckendorn, **Environmental Robustness in Multi-Agent Teams**, *GECCO 2009: Proceedings of the Genetic and Evolutionary Computation Conference*, ACM Press, New York, NY, pp 177-184, ACM, 2009
- Terence Soule, Robert B. Heckendorn, **Improving Performance And Cooperation In Multi-Agent Systems**, Genetic Programming Theory and Practice 2007, Center for Complex Systems, University of Michigan Springer, Chapter 13, 2007
- Robert B. Heckendorn **Partitioning, Epistasis, and Uncertainty**, *Foundations of Genetic Algorithms - 7*, Morgan Kaufmann Publishers, Inc., Palo Alto, CA., pp27-44, 2002

- R. B. Heckendorn, **Building a Beowulf: Leveraging Research and Department Needs for Student Enrichment via Project Based Learning**. *Journal of Computer Science Education*, 12(4), pp. 255-273, Swets & Zeitlinger Publishers, The Netherlands, 2002

Synergistic Activities

- Designed and built affordable robots out of smartphones and RC cars. Taught multiple classes in the design and programming of these robots. Used the robots to do recruiting at Women in Engineering day among other events.
- Took a real budget of $45K and organized a class around designing, buying parts, and building a 70 node Beowulf cluster in one semester. Then wrote an article for a CS education journal on the educational benefits.
- Organized a Ph.D. program for a student that involved not only simulation work on cancer but work in a cancer researcher's wet lab at another university to make a true interdisciplinary degree.
- Wrote software to support my classes by providing a way for students to check in their homework and get it tested and results mailed back to them in 5 minutes. The suite of tools include web interface, grading tools, and the ability to quickly load, compile, and test a student's code to efficiently and comprehensively answer questions they might have.
- Currently working with Drosophila researcher at Michigan State and Biologist at Brown on new statistical models and algorithms for epistasis.

Collaborators and other Affiliations

Dr. Ahmed Abdel-Rahim, Civil Engineering, University of Idaho
Dr. Lisa Carlson, Political Science, University of Idaho
Dr. Raymond Dacey, Business and Economics, University of Idaho
Dr. Ian Dworkin, Department of Zoology, Michigan State University
Dr. Erik Goodman, Director of BEACON Center for the Study of Evolution in Action
Dr. Arend Hintze, Department of Microbiology and Molecular Genetics
Dr. Kay Holekamp, Department of Zoology, Michigan State University
Dr. Richard Lenski, Department of Microbiology and Molecular Genetics
Dr. Terence Soule, Computer Science, University of Idaho
Dr. Daniel Weinreich, Department of Biologist, Brown University

Graduate Adviser

Dr. Darrell Whitley, Computer Science, Colorado State University
Advisees in the last 5 years

Students in progress: Max McKinnen, Joshua Rubini
Students completed: Armand Bankhead, Ph.D. 2006, Oregon Health Sciences University; Smitha Kara M.S. 2005, unknown; Jian Shen M.S. 2004, unknown