

By L.D. Jackel, Douglas Hackett,
Eric Krotkov, Michael Porschbacher,
James Pippine, *and* Charles Sullivan

HOW DARPA STRUCTURES ITS ROBOTICS PROGRAMS TO IMPROVE LOCOMOTION AND NAVIGATION

*Research teams
use common
robots and machine
learning to teach the
robots outdoor
navigation and
locomotion skills.*

The two Defense Advanced Research Projects Agency (DARPA) robotics programs we discuss here are designed to further the development of autonomous locomotion and navigation. Each addresses challenging problems that have shown steady but slow progress to date. Now, however, a combination of machine-learning techniques [5] and smart development techniques has begun to accelerate the pace of autonomous systems development.

One innovative feature is the way the programs' research teams have been supplied with common hardware so they are better able to focus on autonomous behavior. In addition, the programs

use frequent testing to monitor success and uncover and resolve deficiencies. This promotes rapid innovation, allowing both the government and the teams to modify their methods as they see the results. Finally, the programs foster collaboration among the teams. Working together, they have built on successful approaches, avoided dead ends, and made more progress than would be expected in an isolated, competitive environment. The programs have thus far achieved autonomous locomotion by a robotic quadruped over terrain that was completely impassible two years ago and have more than doubled the speed of autonomous navigation through complex terrain.

People expect robots to get from place to place on their own, much as people do. That is, using their on-board sensors, they should be able to know where they are and where they are trying to get to and understand what is in between and how to control their bodies to traverse the distance. In general, this is an unsolved problem.

If we constrain the problem to on-road navigation, with waypoints defining the route to an accuracy of a few meters, with few obstacles on the prescribed path, then robotic vehicles can successfully travel at high speed for distances exceeding 100km. Such a milestone was achieved in the 2005 DARPA Grand Challenge autonomous vehicles race (www.darpa.mil/grandchallenge/) [2, 3]. However, if we require autonomous navigation in rough off-road, complex terrain with widely spaced waypoints, then robotic systems fare poorly. To address this need, DARPA created the two programs: Learning Applied to Ground Vehicles (LAGR) in 2004 and Learning Locomotion (L2) in 2005. Here, we describe the experimental methodology they use; the algorithms developed for the programs are described elsewhere [5].

The traditional approach to autonomous robot navigation for the 10 years preceding the programs was to map out the 3D environment in the vicinity of the vehicle through either laser range finders (LADAR) or stereo cameras, then use a form of rule-based system to determine which regions were traversable and which were not. A path-planning algorithm would then direct the vehicle to its destination. Motor control was usually accomplished through a combination of rule-based commands. While progress was made with this approach, the sys-

tems tended to show brittle, scripted behavior, consistently making the same errors again and again. A report commissioned by the National Academy [1] showed only a doubling of autonomous off-road speed in complex terrain over the course of a decade. DARPA expected that by introducing machine learning into the picture, the rate of progress would accelerate significantly.

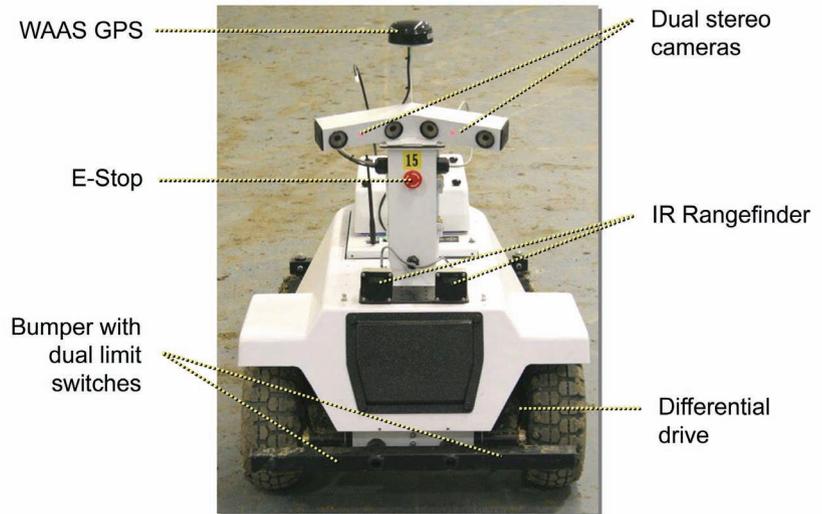


Figure 1. hHerminator vehicle.

DARPA identified three major challenges for autonomous navigation in complex, unstructured environments:

- Detecting objects robustly to discriminate among objects (such as rocks and other obstacles and compressible bushes and other non-obstacles);
- Overcoming limitations of nearsighted sensing (LADAR or stereo) to minimize the probability of the vehicle getting caught in a cul-de-sac; and
- Walking, using a high degree-of-freedom vehicle in extremely rough terrain.

Looking to make quick progress, DARPA decided to factor the problem. The first two items are primarily issues in machine perception and are addressed by the LAGR program. The third item concerns motor control and is addressed by the L2 program.

Like other DARPA efforts, prior robotics programs funded multiple teams to explore the problem in different ways. In most of them, researchers would first build or customize an existing robot, then write low-level “housekeeping” software to control the robot’s systems, and finally write the code that advances the research objectives. In addition to diverting the research effort and driving up costs, this approach meant that each research team had a unique vehicle, making it extremely difficult to compare results from

one group to another.

To circumvent this problem in LAGR and L2, DARPA instead contracted with external suppliers to build small fleets of identical vehicles that were supplied to each research team. For the LAGR program, Carnegie Mellon University's National Robotics Engineering Center built the 100kg hHerminator vehicle (www.rec.ri.cmu.edu/projects/lagr/index.htm) shown



Figure 2. Learning locomotion test site; the Little Dog vehicle is near the center of the terrain board.

in Figure 1. The hHerminator's onboard sensors include two stereo camera pairs, a wide-area augmentation system GPS antenna and receiver, an inertial measurement unit (IMU), a bumper impact sensor, and short-range IR sensors, all controlled by four Pentium M computers. While the hHerminator is sensor-rich, its motor control is simple; its two front electrically driven wheels can be moved in tandem to go forward or in reverse or differentially to turn. The rear wheels are passive casters. hHerminators are meant to drive off-road but only over terrain with noncompressible obstacles no greater than about 5cm in height.

For the L2 program, Boston Dynamics built the 3kg Little Dog vehicle (see Figure 2) (www.bostondynamics.com/content/sec.php?section=LittleDog), a 12-degree-of-freedom quadruped robot with onboard IMU, foot touchdown sensors, and short-range IR sensor; otherwise, it is blind. The control system relies on an external Vicon motion-capture (mocap) system to measure the position and pose of the robot in real time. DARPA intentionally removed the extremely challenging perception problem from the program, allowing researchers to concentrate on locomotion.

The L2 program also built a series of terrain boards

on which the Little Dog walks. These boards were laser-scanned to obtain precise surface geometries that are also supplied to the control system in real time. Thus, by using the mocap systems in conjunction with the scanned boards, a Little Dog has nearly perfect knowledge of its environment and of itself. This means the L2 research teams can focus on learning and control systems, without having to also tackle perception.

Figure 2 shows a Little Dog on a terrain board surrounded by a mocap system; most terrain boards created to date mimic the natural environment.

Without resorting to special-purpose computing hardware, Little Dog is too small to carry significant onboard computer power. Most of the complex computing used to control Little Dog is provided by an off-board processor linked to the robot by a wireless connection.

Each of the eight teams selected for the LAGR program was supplied with two hHerminators. In the L2 program, each of the six teams was supplied with a Little Dog robot and mocap system. In both programs, after three months of familiarization with their systems, the teams were required by DARPA to submit code for monthly tests at a government site.

The monthly tests were a key aspect of both programs and served several functions. First, they gave quick feedback to both DARPA and the teams on the progress of their research. This allowed DARPA to adjust subsequent tests to push them toward maximum performance. Second, making the results of the tests known among the teams fostered competition, further encouraging them to excel.

The LAGR program was structured as two 18-month phases; L2 has had a single 15-month phase

followed by two 12-month phases. In order to advance to the next phase in each program, teams are required to pass preestablished metrics.

DARPA decided at the inception of both programs that these criteria would be fixed and not relative, so the teams would not compete with one another for second-phase contracts. It was thus possible that all or none of the teams would advance to the second phase.



This policy contrasts with a down-selection in which only a limited number of first-phase teams advance. The absence of a down-select means that teams are more willing to share methods and even code with one another. DARPA encourages such sharing and allows each team to contribute in areas in which it is most capable. For the third phase of the L2 program, scheduled to begin in the second half of calendar year 2008, DARPA will down-select from six to three performer teams.

Objectively measuring progress in autonomous robotics research has always presented challenges, including a lack of a standard vehicle and an a priori measure of the difficulty of a course (which in turn depends on the mechanical capability of the vehicle), as well as difficulty comparing results from one test site to results from other sites. Testing often takes the form of a large, complex demonstration at the end of a development program. Developers do not have the opportunity to learn from their mistakes. At its worst, testing is performed under carefully selected conditions to ensure a high probability of success. Such tests do not adequately explore the robustness of a system or assess how it might perform in the real world.

In both programs, all teams are supplied identical (within manufacturing tolerances) vehicles. They

develop code to control their vehicles at their own facilities, then send the code to DARPA for testing. The DARPA test team then independently evaluates the code onto either an hHerminator or a Little Dog on identical robots at the program test site and runs a series of trials to try to determine the effectiveness of the code.

The hHerminators in the LAGR program were shipped with a modular “baseline” code developed at the National Robotics Engineering Center, which was state-of-the-art in 2004 and a legacy of the completed DARPA PerceptOR program [4]. Thus the teams were able to examine how the baseline system performed in their own environments. The baseline code included modules for stereo analysis, obstacle detection, and path planning. Teams were able to replace individual modules as they developed their software; they were also able to compare the performance of their own code against that of the baseline system to readily determine if their modifications were indeed improvements.

In the LAGR tests, the baseline code also served an additional function: calibrating the difficulty of the test courses (each about 100 meters long) devised by DARPA. That is, the average speed of an hHerminator vehicle running baseline code on a test course was defined as the course’s baseline speed. A team’s code’s speed on the course was normalized by the baseline speed, enabling consistent comparison with the baseline. This

process helped compensate for variations in difficulty among different courses and let DARPA measure progress from one test to the next.

By design, LAGR courses are changed each month so teams cannot “memorize” the features of a particular course. About 70% of the tests have been conducted at various locations at the U.S. Army’s Fort Belvoir in Virginia. The other tests were conducted at test sites in Hanover, NH, and San Antonio, TX (see Figure 3). These locations were chosen for their ability to provide a variety of terrain and vegetation types.

In order to advance to the second phase of the LAGR program, teams had to demonstrate an average speed 10% faster than the baseline system on two of the three final tests in the first phase. This rather modest metric was chosen to allow teams to attempt risky but promising approaches that might not be fully developed during the first 15 months of the pro-

Figure 3. LAGR testing at the Southwest Research Institute, San Antonio, TX.

gram. All eight teams achieved this metric on the phase-end tests, with speeds from 1.2 to 2.2 times the baseline performance. The objective for the second phase, which is still in progress, is three times the baseline speed under more robust conditions. Compared to the National Academy report, LAGR has compressed the pace of doubling performance from 10 years to less than 36 months.

In the L2 tests, courses took the form of manufactured terrain boards. Each new board was used for a DARPA test, then distributed to the teams for further in-house testing. New boards were not distributed prior to a test so the teams would not be tempted to memorize a script for traversing the board. Because boards are expensive to manufacture, it was not always possible to test on a "virgin" board. In these circumstances DARPA would usually change the orientation of a course across a board or through the tilt of a board to foil attempts at scripted runs.

In both programs, DARPA gave teams complete logs of the test runs to recreate test conditions and analyze their performance. Teams are still mastering the complex terrain boards, but it is already clear that L2 has significantly advanced autonomous legged locomotion on extreme terrain (www.cs.cmu.edu/~cga/leg-learn/).

In order to advance to the second phase of L2, teams were required to make their Little Dogs traverse obstacles 0.4 leg lengths in height (4.8 cm) and move at an average speed of 0.1 leg lengths per second (1.2 cm/sec). For the second and third phases, the metrics for speed are defined as 4.2 and 7.2 cm/s, respectively, and obstacle heights are defined as 7.8 and 10.8 cm, respectively.

A key cost-benefit of this standardized testing approach is that it reduces testing costs to the government. By using a common platform and standardized software process, it is able to test frequently at a lower cost per test. In prior DARPA robotics programs, the government's test group would invite research teams to a central location to compare the performance of several systems on a standard course. The government would incur substantial shipping and travel costs for both the test and the research teams. In addition to financial costs, the extra travel and logistics takes valuable time from research.

In DARPA's new model for testing, the government has reduced its costs in several areas. First, by providing a common platform, teams focus on software development and do not have to dedicate a team member to vehicle maintenance. Second, because the teams transmit their code to the government rather than travel to an event, the government is able to evaluate progress on its own timetable. Third, research

team personnel save time by not having to travel to each test, instead monitoring performance by viewing real-time video and streaming data. Finally, disseminating data logs to all teams gives each one a much larger data set than it could reasonably expect to collect alone.

CONCLUSION

These two DARPA programs in learned autonomous locomotion and navigation have developed clear and simple experimental methods for measuring progress and for encouraging cross-team cooperation. Even though both are still ongoing, they have generated considerable scientific knowledge and insight. By providing each team with a standard vehicle, DARPA has focused research on producing new science, rather than on nursing a semi-custom robot. The framework for regular, objective evaluation of relative performance was intended to promote innovation, not hinder it. Midway into these programs, initial results indicate the strategy is working. 

REFERENCES

1. Committee on Army Unmanned Ground Vehicle Technology, National Research Council. *Technology Development for Army Unmanned Ground Vehicles*. National Academies Press, Washington, D.C., 2002; www.nap.edu/catalog/1.php?record_id=10592#description.
2. Iagnemma, K. and Buehler, M., Guest Eds. Special issue on the DARPA Grand Challenge, Part 2. *Journal of Field Robotics* 23, 9 (Sept. 2006), 655–835.
3. Iagnemma, K. and Buehler, M., Guest Eds. Special issue on the DARPA Grand Challenge, Part 1 *Journal of Field Robotics* 23, 8 (Aug. 2006), 461–652.
4. Krotkov, E., Fish, S., Jackel, L., Perschbacher, M., and Pippine, J. The DARPA PerceptOR evaluation experiments. *Autonomous Robots* 22, 1 (Jan. 2007), 19–35.
5. Mulligan, J. and Grudic, G., Guest Eds. Special issue on machine-learning-based robotics in unstructured environments. *Journal of Field Robotics* 23, 9 (Sept. 2006), 655–835.

L.D. JACKEL (larry.jackel@att.net) was a program manager at DARPA from 2003 to 2007 when he created the programs described here and is now president of North-C Technologies, Holmdel, NJ.

DOUGLAS HACKETT (dhackett@griffin-technologies.com) is a principal of Griffin Technologies, Inc., Wynnewood, PA.

ERIC KROTKOV (ekrotkov@griffin-technologies.com) is president of Griffin Technologies, Inc., Wynnewood, PA.

MICHAEL PERSCHBACHER (mike@rovnotech.com) is president of RovnoTech, Inc., Springfield, VA.

JAMES PIPPINE (jpippine@sysplan.com) is a project manager at System Planning Corporation, Arlington, VA.

CHARLES SULLIVAN (csullivan@mapcorp.com) is a mechanical engineer at Maritime Applied Physics Corporation, Baltimore, MD.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.