

CHAPTER XI. FUTURE RESEARCH

XI.A. Use a sensor to train the controlling intelligence to interpret other sensors

Team 2005-16 used LIDAR sensors to train a single color camera to detect obstacles at a range which exceeded the maximum effective range of LIDAR sensors¹⁴.

Team 2005-16 stated: “To extend the sensor range enough to allow safe driving at 35 mph, [the challenge vehicle] uses a color camera to find drivable surfaces at ranges exceeding that of the laser analysis.” ([51], p. 672). However, this strategy could be extended to other combinations of sensors in simulation. For example:

- GPS/INS/IMU output could be used to train the controlling intelligence to detect “slippage” of steering position and odometry. DARPA stated an “independent technical evaluation team identified the following technology from Grand Challenge 2004 noteworthy”: “Sensor-based slippage detection (conceptual)” ([3], pp. 10 - 11)¹⁵.
- LIDAR sensors could be used to train RADAR sensors to see farther up the road, increasing the maximum effective range of RADAR sensors, or providing a basis for the development of more effective navigation RADAR.
- Position sensors could be used to develop algorithms to integrate incremental distance measurements provided by sensors such as magnetic or optical encoders on axles or the drive shaft, differential odometers, etc. more effectively.

This is similar to the strategy utilized by COTS components. For example, Team 2005-06 stated: “[Team 2005-06] chose to use the RT3000 from Oxford Technical Solutions to provide vehicle localization. ... The integrated INS

allows the RT3000 to survive GPS outages of up to 30 seconds with virtually no performance degradation. Because the GPS and INS are integrated together, each can compensate for problems with the other. For example, if the INS started to drift laterally, the integrated GPS will automatically correct that drift.” ([53], p. 9).

- Distance could be estimated by throttle position for unit time and slope, and integrated over changes in terrain roughness, providing an alternative to dead-reckoning.

In addition, this strategy could be extended to combinations of sensors which are not obviously complementary. For example:

- Team 2004-07 described how the controlling intelligence used information such as time of day, orientation, and lighting conditions to detect obstacles: “Since the system will know the time of day, its orientation, and the lighting conditions, it can employ a shape-from-shading and shape-from-shadow system to determine the approximate position and dimensions of obstacles like large rocks or craters.” ([54], p. 5). However, there is no reason the controlling intelligence would not be able to determine the time of day, orientation, or lighting conditions using the approximate position and dimensions of obstacles.
- Team 2004-09 stated: “Road boundaries and obstacles will be reliably detected when the vehicle is bouncing over rough terrain and turns. We will use a rapid shutter speed of 1/8000 sec. to minimize blurring. We will mount the camera and

other sensors on a platform designed to absorb shock. Inertial data will normalize the image perpendicular to the ground when the vehicle is tilted one direction or the other. In addition, when the vehicle is driving over uneven terrain, the normalization process attempts to use information from previous images to locate the horizon and road. Topographic information may also be used to locate the horizon and road. Images that do not normalize to recognizable data can be skipped because the frame rate of 30 frames/sec. is more than sufficient to allow us to dispose of 'bad frames.' If the vehicle is tilted upward or downward so that the camera is facing images of sky or ground, the autonomous control can use pitch information to discard those frames.” ([55], p. 7).

This is similar in concept to Team 2005-16's later use of LIDAR sensors to train a single color camera to detect obstacles at a range which exceeded the maximum effective range of LIDAR sensors, but using shock, vehicle attitude relative to the horizon, and accelerometer data to normalize data. In their technical proposal, Team 2004-09 does not report their controlling intelligence was trained to normalize the data, but learning to normalize visual processing data is a potential task for a controlling intelligence.

XI.B. Emergence of unexpected behavior

In general human beings observe certain “rules of the road”: they navigate roads with recognizable characteristics such as color, texture, lane markings, and signage which establishes context and allows them to determine what is, and is not, a road; they travel from point to point in lanes, the width of which varies depending on location; and they

must obey posted speed limits. However, human beings are not constrained by the electronic equivalent of an overwhelming compulsion.

A truly autonomous vehicle would be able to evaluate its own objectives within the constraints imposed on it by its programming, and it would violate some rules if necessary to accomplish them, for example, by proceeding on a more direct course to its destination if requested to travel a circuitous course similar to the 2005 GCE course which crosses and overlaps itself in several areas. There is no evidence that the emergence of unexpected behaviors was a goal or outcome of the Grand Challenge. If the development of artificial intelligence is a goal of autonomous vehicle development, the emergence of unexpected behaviors would be a measure of successful development.

XI.C. Development of novel sensor technologies

Several teams attempted to use low-cost photoelectric, ultrasonic, or short-range RADAR sensors to provide useful information to the controlling intelligence. Because of their limited utility in practice, these sensors were discounted by the author. However, DARPA stated an “independent technical evaluation team identified the following technology from Grand Challenge 2004 noteworthy”: “Extended range of low-cost, ultrasonic sensors” and “Single-point laser rangefinder as a low-cost distance sensor” ([3], pp. 10 - 11).

Simulation might enable the identification and development of novel sensor technologies, such as a SONAR sensor array that provides a 3D point map as accurate as that provided by a LIDAR sensor, but using SONAR returns, the effective use of non-scanning LIDAR sensors, or the development of a goniometer (direction-finding antenna)

for providing accurate position information. As a minimum, the use of simulation might provide an environment in which the practical applications of such sensors could be explored.

In addition, combined sensor strategies in use by teams which participated in the 2004 and 2005 GCE included the use of LIDAR in combination with high-quality STEREO or RADAR, but alternate strategies were in use. Each strategy was specifically tailored to a challenge vehicle. Simulation might increase the likelihood the generic application of the combined sensor strategies in use by most teams would be adequately explored and potential commercial applications identified.

XI.D. Use simulation to train the controlling intelligence to recover from a loss of sensor data or other sensor failure

XI.D.1. Primary obstacle and path detection sensor

Several teams reported a single sensor was in use by the team as the primary obstacle and path detection sensor:

- A proprietary stereo camera pair was in use by Team 2004-06.
- One SICK LMS 291-S05 was in use by Team 2004-12.
- One Epsilon Lambda ELSC71-1A was in use by Team 2004-21.
- A proprietary video system was in use by Team 2004-22.
- A proprietary LIDAR sensor was in use by Team 2005-03.
- A Point Grey Bumblebee stereo camera pair was in use by Team 2005-12.

Neither Team 2004-06, 2004-12, 2004-21, 2004-22, 2005-03, nor 2005-12

reported how the challenge vehicle controlling intelligence would respond to the loss of the single primary obstacle and path detection sensor.

In general, teams which reported multiple obstacle and path detection sensors were in use by the team also did not describe how the controlling intelligence would respond to the loss of a sensor. Three teams which participated in the 2005 GCE reported a sensor, type of sensor, or array of sensors was “redundant” in the sense that it provided obstacle and path detection information in the event a sensor failed. The author considers this to be functional redundancy. For example:

- Several obstacle and path detection sensors were in use by Team 2005-08, including three Delphi Forewarn ACC3 RADAR. Team 2005-08 stated: “[The Delphi Forewarn ACC3 RADAR] can act as a redundant sensor for the [challenge vehicle].” ([56], p. 9).
- Although the author concluded ultrasonic sensors were not in use by Team 2005-15, Team 2005-15 stated: “...the ultrasound sensors act as additional redundant sensors, which are less susceptible to dust or fog.” ([22], p. 9).
- Team 2005-20 stated: “Our goals were to... develop a sensor array that contains redundancy for accuracy and reliability...” ([29], p. 2).

The author considers it likely teams selected multiple complementary obstacle and path detection sensors by necessity and to have functional redundancy. For example, Team 2005-10 stated: “There does not appear to be any one sensor that can 'do it all'. Each sensor has its strengths and its weaknesses.” ([57], p. 7).

However few teams reported how the controlling intelligence would respond to the loss of sensor data or other sensor failure, perhaps because DARPA did not explicitly request teams provide such information. In contrast, DARPA explicitly requested teams determine how the controlling intelligence would respond to “GPS outages”. As a result, teams generally reported how the controlling intelligence would respond to the loss of GPS data or GPS failure. See paragraph XI.D.2.

Several teams acknowledged the loss of sensor data or other sensor failure would affect challenge vehicle performance. For example:

- Team 2004-01 stated: “Speed setting algorithms will take into consideration the following and reduce speed appropriately: ... Sensor obstruction ... Sensor disagreement, Data discontinuities or gaps ... Component failure” ([58], pp. 6 - 7).
- Team 2004-02 stated: “Component failure testing: Since [the challenge vehicle] cannot operate without power, testing will be done to insure that the vehicle has power the whole race. These tests will include cutting power to individual sensors, computers, and support electrical units.” ([59], p. 13).
- Team 2005-04 stated: “These sensors are monitored for changes in their operating state, validated using both dynamic and rule based tests, and finally fused using a Kalman filter based approach to provide continuous position and orientation information even [*sic*] the presence of individual sensor dropouts, reduced accuracies, or complete failures.” ([13], p. 11).
- Team 2005-11 stated: “Hardware and software have been designed to minimize the impact of temporary failed components. However, limited redundancy in

components means that permanent outages of sensors will have a detrimental effect on [the challenge vehicle's] performance.” ([18], p. 7).

Finally, a few teams reported specific action to be taken to resolve a loss of sensor data or other sensor failure. None of these teams described how the controlling intelligence would respond to a loss of sensor data or other sensor failure:

- Team 2005-12 stated: “The emergency brake’s pneumatic system is setup such that any failure of the [the challenge vehicle's] software or hardware will result in an emergency brake application.” ([60], p. 3).
- Team 2005-20 stated: “A failure of any individual sensor results in no information being broadcast from that specific sensor.” ([29], p. 7).
- Teams 2004-13, 2004-14, and 2005-15 reported an emphasis on the isolation of hardware and software modules from each other so that a failure in one module does not cause an overall failure, and Teams 2005-13, 2005-14, 2005-16, and 2005-19 reported an emphasis on restarting modular hardware and software components.

Three teams reported a loss of obstacle and path detection sensors or other sensor failure during the 2005 GCE: Teams 2005-14, 2005-15, and 2005-18. Team 2005-14 successfully completed the 2005 GCE. Team 2005-15 reported a loss of all LIDAR sensor and internal state data due to a “USB hub” failure. Team 2005-18 reported a loss of “midrange” LIDAR sensor data.

The author concluded the failures were preventable system integration failures. Because Team 2005-14 had significant experience but neither Team 2005-15 nor 2005-18 had significant experience, the author proposes the use of simulation may have helped “level the playing field”, by enabling teams without significant experience to learn how to recover from a loss of sensor data or other sensor failure as well as an experienced team and eliminate the causes of the preventable system integration failures which resulted in their failure to complete the 2005 GCE.

XI.D.2. GPS sensor failure

GPS “drift” or “jumps” were consistently reported by teams which participated in the 2004 QID or GCE or 2005 GCE. For example, Team 2005-05 stated: “Very often, especially when the vehicle would drive near a wall or approach a tunnel, there would be highly erratic jumps in the GPS measurements due to multipath reflections.” ([15], p. 542).

In addition, GPS sensor failure was directly implicated in the failure of five teams to complete the 2005 GCE: Teams 2005-02, 2005-09, 2005-15, 2005-18, and 2005-19. For example, Team 2005-02 stated: “... it appears that the calculated GPS position drifted by approximately 20 feet causing the vehicle to want to move to the right of the actual road.”, which caused “a corresponding shift of the boundary smart sensor that eliminated the actual sensed road as an option to the planner.” ([12], p. 621).

DARPA, via 2004 SQ 1.g.2 and 2005 SQ 2.2.1 requested teams describe how they would handle “GPS outages”¹⁶. In general, teams described how the challenge vehicle controlling intelligence would continue to determine position reliably in the absence of

GPS data. A few teams described test and evaluation to determine the effect of GPS outage on the challenge vehicle controlling intelligence. For example:

- Team 2004-17 stated: “We have tested the ability of various materials to block antenna reception. Flat sheets of aluminum and Lucite were unable to block the GPS, as multi-path reflections off of the ground still reached the antenna. Wrapping the antenna in aluminum foil cut off reception (we can selectively cut off satellites and simulate GPS outages).” ([61], p. 12).
- Team 2005-06 stated: “Another extremely effective test involved manually steering the vehicle off course at high speed and then switching back to autonomous mode. This simulated a GPS jump, which can occur rather frequently. After noticing that the navigation system abruptly turned the steering wheel to counteract this jump, the navigation system was updated to eliminate this abrupt movement.” ([53], p. 12).

In addition, the military deployment of autonomous ground vehicles will result in the development of countermeasures to preclude their use¹⁷. For example:

- Strong magnetic fields may confuse magnetometers, causing the vehicle's controlling intelligence to incorrectly interpret compass headings.
- The U. S. government's ability to control the accuracy of GPS position information using “selective availability” is a strategic limitation on the use of GPS. Although U. S. military ground vehicles would not likely be affected by selective availability, an effective controlling intelligence should be able to

identify the problem if it occurs and adjust the weight of other sensors appropriately or take corrective action to determine geolocation using some other method, such as dead-reckoning. Alternate strategies, such as using beacons or reflectors delivered by artillery, or aerial drones, to provide stable “known” geolocation similar to survey markers may also be successful.

Although effective, the test method employed by Team 2005-06 represented a real risk to the team challenge vehicle. Effective simulation may have allowed teams to develop strategies to mitigate the effects of GPS drift or jumps and to gracefully recover from a temporary or permanent loss of GPS sensor data by allowing a model to be driven off course, then “switching back to autonomous mode” in a manner similar to that reported by Team 2005-06, but without risk to the team challenge vehicle.

XI.E. Standardization and standard references

XI.E.1. Standard dictionary, acronyms, and abbreviations

Develop a standard dictionary of terms and their associated acronyms and abbreviations for use in future research similar to the Grand Challenge to be maintained as a set of user-defined dictionaries for various word-processing applications. The author acknowledges that current word-processing software is limited in the amount of customization that it provides. For example, OpenOffice.org Writer (version 3.0.1) supports custom dictionaries, but does not yet provide a way to import or export custom dictionaries. In addition, there is a finite limit to the number of words allowed in a custom dictionary.

XI.E.2. Standard reference terrain

Develop a library of standard reference terrains using available sensors to gather complete data using environment and geolocation sensors consistent with the state of the art. For example, using a research platform with roof-mounted cameras, and LIDAR, RADAR, and GPS sensors visit:

- the Pennsylvania Turnpike, to record mountainous terrain, including several extremely long tunnels during which GPS reception will be lost
- the Mojave Desert, to record desert terrain, including the “negative obstacles” typically encountered in desert terrain such as wadis
- the California coast on US-1 (the Pacific Coast Highway), to record coastal highway, extending north through the Redwood National Forest
- Interstate 40, to record a long traversal across the United States with many different reference terrains

While traversing reference terrain, record the precise geolocation on a continuous basis. Use existing technologies to subtract vehicles and other obstructions from the reference terrain as recorded by the environment sensors in use. Correlate GPS position with the terrain in simulation. Use the LIDAR data to produce a “point map” of the reference terrain, and map the return from camera sensors onto this point map as a trimesh, providing simulated cameras with more realistic data.

This would make it possible to add vehicles and other obstructions as desired, or to test the controlling intelligence in an environment completely devoid of risk to other

vehicles while still allowing it to perceive at the limit of available environment sensor technology.

In addition to recording the standard reference terrain in different locations, record the standard reference terrain at different times of the day and year. Although there may be little difference to a LIDAR sensor from night to day, the difference to a camera will be significant. In addition, there will be a significant difference between the efficiency of a camera or LIDAR sensor pointed into the sun at sunrise, heading east, or sunset, heading west, and at other times of the day. Terrain details may be obscured by snow during the day, or brought into sharper contrast at night. All of this is useful information to the controlling intelligence.

Simulation environments such as the Player Project could be modified to use simulated reference terrain for real-time testing.

XI.E.3. Standard obstacle and position problems

Develop a library of standard obstacle and position problems (herein “standard problems”), and acceptable responses based on human driving tests. These problems should first be implemented in simulation to support the development of algorithms and acceptable responses. Acceptable responses should then be verified during real world testing. For example:

Every state has established a standardized program of driver education which requires a minimum level of competency to be demonstrated by drivers prior to licensure. For example, in Virginia, this program is called “The Driver Education Standards of Learning and the Curriculum and Administrative Guide for Driver Education in Virginia”

(herein “Guide”) ([62]). The Guide describes a series of “Modules” presenting required course content. Module 11 is titled “Laboratory Instruction – Behind-the-Wheel and In-car Observation”. Module 11 describes a series of “Lessons”, “Basic Skills”, and “Driving Procedures”, which ensure the driver has achieved a minimum level of competency ([63]). Successful completion of the 2007 Urban Challenge was determined, in part, by the challenge vehicle's controlling intelligence's ability to obey California state traffic laws.

It is not unreasonable to require an autonomous vehicle's controlling intelligence to meet or exceed the basic minimum level of competency expected of a human driver, in effect making the standard problems, and acceptable responses, a “Turing test” for autonomous vehicle controlling intelligences.

It is unreasonable to expect the public to be forgiving of an autonomous vehicle which loses contact with a GPS signal, for example, and unexpectedly stops in a tunnel during rush-hour traffic, or to accept the loss of life and property damage that may be caused by an autonomous vehicle that loses the ability to distinguish between the road and terrain in the rain, and crosses the center line of a divided highway with disastrous consequences. As a result, standard problems must also evaluate the controlling intelligence's ability to meet or exceed the basic minimum level of competency expected of a human driver in similar situations.

Also, this approach would allow the controlling intelligence to be trained to respond to situations in a manner uncharacteristic of human drivers. For example, a human driver reacting to a vehicle entering the lane next to his or her vehicle might react

out of fear, pulling the steering wheel suddenly to the right or left to avoid collision, and entering the next lane, unintentionally causing an accident. An autonomous vehicle's controlling intelligence would be able to more effectively estimate the position of the autonomous vehicle in relation to its surroundings, and decide not to attempt to avoid a collision if attempting to avoid the collision will cause a collision with another vehicle and if the autonomous vehicle will not be seriously damaged. However, if the vehicle pulling into the lane next to it is a 40-ton tractor-trailer, the autonomous vehicle's controlling intelligence might conclude a collision is unavoidable, and decide to collide with a lighter vehicle, due to the tractor-trailer's greater damage potential.

XI.E.4. Team descriptions of standard reference terrain and standard problems

Several teams described attempts to gather standard reference terrain or proposed the implementation of standard problems. However, no team proposal was comprehensive. For example:

- Teams 2004-13 and 2004-14

Teams 2004-13 and 2004-14 were co-competitors during the 2004 GCE, and stated: “During field trips to the Mojave desert, we have recorded more than 7 hours of video from a vehicle-mounted camera, recording the path ahead. We have run parts of these video sequences through our path tracking software.” ([64], p. 6 and [65], p. 7). However, this approach was not comprehensive, in that it did not allow the teams to adjust the mounting of the camera to optimize the performance of their path tracking software or experiment with different types of cameras.

- Team 2004-20

Team 2004-20 stated: “The road-follower software has been tested against video recordings of desert roads, with marginally satisfactory results. The imagery used was too narrow. The road follower is being revised and will be retested with wider-field imagery.” ([52], p. 9). As noted by Team 2004-20, this approach was not comprehensive, in that it did not allow the team to adjust the field-of-view of the camera to optimize the performance of their road-following software.

- Team 2004-23

Team 2004-23 described a special type of terrain called “Robot”, and stated ([34], p. 6):

“Robot” is a special terrain/location where the vehicle has to go through a specific exercise, possibly with a set of predetermined operations, to go past an obstacle or through a narrow constrained passage.

Examples where Robot behavior may be needed include underpasses, gates, sharp turns at roadway intersections and possible passage through mazes of natural and synthetic obstacles.

- Team 2005-01

In response to 2005 SQ 2.4.1, Team 2005-01 stated ([66], p. 11):

Extensive testing in the field has led to extensive development of these corner cases. [The challenge

vehicle] does not return to missed waypoints, since in many cases the road is not wide enough to make a full turn to reach the missed waypoint. The vehicle will continue along the assigned path in this case.

When the vehicle is "stuck", this may occur with wheels slipping, and the vehicle is not actually driving forward. For this case, we detect this condition in the National Instruments software, and reverse a few meters to free ourselves from this condition.

If the vehicle travels out of bounds, the "boundary" voter immediately pushes us back into bounds by providing a strong negative weight along any path that continues out of bounds. If an obstacle is detected in the path, the vehicle detects this with either the four LADAR sensors or the five bumblebee cameras. Upon detection, the vehicle's path is adjusted to pass the obstacle by with a safety margin.

- Team 2005-04

Team 2005-04 described a special case for braking or starting on a hill: "The speed set point is generated regardless of the slope of the ground. The speed controller

has the 'integration' part that keeps increasing the throttle if the vehicle is slower than the speed set point so that we can climb a hill. In order to stop short in some situations, the vehicle applies the maximum brake pressure.” ([13] , p. 13).

At least one team failed to complete the 2004 GCE due to an inability to increase throttle sufficiently to climb a steep hill. Team 2005-05 participated in the 2004 GCE as Team 2004-07. Team 2005-05 later stated: “[The Team 2004-07 challenge vehicle] traveled 5.1 miles in the 2004 Challenge... before stopping on a steep slope because of an excessively conservative safety limit on the throttle control.” ([15], p. 528). As a result, the author considers this problem a potential standard problem.

- Team 2005-08

Team 2005-08 stated: “...in December 2004 a team of engineers with two sensor instrumented platforms drove large segments of the course, collecting navigation, image, and laser data for algorithm development and design validation for components such as the shock isolation sled.” ([56], p. 22). However, this effort was not comprehensive.

Although Team 2005-08 collected standard reference terrain similar to that expected to be encountered during the 2005 GCE, the development of fully autonomous vehicles will require a greater library of reference terrain be available.

- Teams 2005-13 and 2005-14

Teams 2005-13 and 2005-14 stated: “Standardized tests must be developed that measure a robot’s ability to sense and accurately localize obstacles of varying size. These tests should account for differing perception sensing modes. Standard tests that measure an autonomous vehicle’s ability to safely and reliably interact with other vehicles and

humans are needed. These tests and others are required in order to move autonomous ground vehicles from technological curiosities to common tools used by people everywhere.” ([21], p. 499).

However, a library of standard obstacle detection tests is not enough. Terrain affects obstacle detection and avoidance. Autonomous vehicles must also be taught to recognize degraded sensor performance not caused by simple failure of the sensor, such as lack of calibration or misalignment. For example, Team 2005-05 stated: “The virtue of ladars used in this vertical-plane configuration is that the ground profiles are easy to interpret, and are not particularly prone to confusion due to rolling, pitching, or bouncing motion of the vehicle. (Of course, a six-degree error in pitch could make a marginally-traversable 27-degree slope appear to be a marginally-untraversable 33-degree slope, or vice versa.” ([48], p. 6).

In this particular example, the purpose of the standard pose estimation problem would be to teach the challenge vehicle controlling intelligence to recognize the error in pitch is caused by sensor misalignment, and compensate accordingly, and not treat the error as a permanent change in slope resulting in a determination that traversable terrain is not traversable, thus overcoming sensory input that is contra-indicative of the challenge vehicle's capabilities.

Multiple solutions to such a problem exist, depending on available sensors. In this example, the controlling intelligence may be able to estimate the slope of a road by measuring the distance known acceleration moves the challenge vehicle in a given time; the controlling intelligence may be able to utilize an altitude sensor or the elevation

reported by commercial GPS to arrive at an independent estimate of the slope of the path of travel; or the controlling intelligence may be able to navigate the challenge vehicle over terrain with known characteristics, such as alternating high and low “striping”, i.e., asphalt or concrete of alternating heights, to determine the error in pitch of the LIDAR sensors in use.

XI.F. Time- and space-shifting

Player and Gazebo provide a “passthrough” construct which allows a client program connecting to the Player server to receive sensor output from a client program connecting to another Player server, and to effectively “see through their eyes”. The author proposes using this or a similar construct to allow notional vehicles (vehicles with no density or which do not implement ODE collision callback functions) to be “stacked” in time or space, allowing the controlling intelligence to receive sensor output from the simulation at some time offset in the future. This would allow the controlling intelligence to use the future results of current decisions to make more informed decisions, and effectively give the controlling intelligence the ability to “see” into the future as a training tool.

XI.G. Acclimation

Develop a process of “acclimation”, whereby the controlling intelligence queries a hardware- and software-independent abstraction layer to discover available sensors, and then uses standard reference terrain and standard problems to acclimate itself to their use. The acclimation process would require the controlling intelligence to learn how its outputs correlate with inputs to the abstraction layer, and vice versa, in effect calibrating

itself¹⁸. This would make the controlling intelligence portable between vehicles, and allow one team to install their controlling intelligence in another team's challenge vehicle. As a result, teams would be competing not on the basis of hardware available to the team, but on the basis of their use of information available from standard interfaces.

For example:

- a challenge vehicle controlling intelligence could determine its own braking profile in a manner consistent with the method used by the U. S. Department of Transportation if visible markers with known spacing for VISION or STEREO sensors were painted, or vertical markers for LIDAR or RADAR sensors were placed, on a stretch of asphalt where they could be detected by a challenge vehicle's sensors.
- a challenge vehicle could similarly determine its own turn radius, or calibrate control of the steering wheel, gas pedal, or brake pedal, allowing a controlling intelligence using a hardware- and software-independent abstraction layer to calibrate itself to the specific vehicle in which it is installed.

XI.H. Least free energy state

Develop a process for correlating “desirability” or “traversability” maps to a concept such as the Gibbs free energy, to allow already-existing concepts to be used to describe the cost associated with moving from one metastable state to another.

Obstacles would be represented as local maxima, regardless of whether they were “positive” or “negative” obstacles. The height of the obstacle could be correlated to the potential damage the vehicle that would result in the event of a collision. Each sensor or

combination of sensors would contribute an individual state map.

Road boundaries would be represented as continuous maxima “walls” of varying height, depending on how tolerant the terrain is to the controlling intelligence deciding to leave the road.

The difference between the height of the road and obstacle height maxima would determine, for example, whether the autonomous vehicle would attempt to leave the road to avoid an obstacle.

A route would be represented as a continuously decreasing “valley” in the local terrain map.

Local maxima representing obstacles detected by LIDAR sensors would be added to local maxima representing obstacles detected by RADAR, road boundaries, and the route to produce a final traversability map.

The autonomous vehicle's controlling intelligence would always seek to travel from one potential energy state to another, always moving from a greater potential energy state to a lower one, like water flowing downhill.

For example:

- Team 2004-07

Team 2004-07 stated: “...a nominal minimum-cost route from each waypoint to the next will be computed based on map data using a wavefront-propagation path planner.” ([54], p. 5).

- Team 2004-15

Team 2004-15 described a “desirability map” ([67], p. 9) that suggests the

controlling intelligence was using a map similar to a free energy diagram, with geolocation represented by the x- and y-axes, and “desirability” by the z-axis. This suggests a decrease in desirability represents a positive slope in the free energy diagram, or negative reinforcement to the controlling intelligence, and that an increase in desirability represents a negative slope, or enticement. However, this model suggests the controlling intelligence would not be able to enter an area which represents a temporary increase in free energy (or lower desirability) to cross to an area at a net decrease in free energy (or higher desirability). As a result, the controlling intelligence might become stuck in a metastable state, from which it would not be able to free itself.

In addition, this approach would eliminate the potential problem of long-term “statelessness” described by Team 2004-15 as the “heading circle”, and as a result of which the controlling intelligence might be unable to ascertain if it is moving back and forth between two positions of high desirability.

- Team 2004-20

Team 2004-20 maintained an extensive online repository which contained several revisions of their technical paper prior to the final version accepted by DARPA ([52]), including DARPA responses to their first and second revisions indicating that DARPA requested Team 2004-20 report: “How will the potential field path planner escape from local minima?”. No 2004 SQ contains the words “local maxima” or “local minima”. Team 2004-20 stated: “Escaping from local minima is the job of the 'higher level' processing...” ([52], p. 3).

- Team 2005-13

Team 2005-13 stated: “Fusion of perception data is via a terrain cost map and binary obstacle map. Terrain cost maps are generated by evaluating the relative height of a sensed area to its neighbors and assigning a cost of 0 to 255 to that area. Binary obstacle maps are created in a two step process. First, an object detection algorithm, customized for each sensor group, detects and localizes obstacles. Second, detected obstacles are written into a map at the detected location.” ([19], p. 10).

In addition, most vehicles have a rollover threshold, a slope on which the vehicle will roll. For example, Team 2004-23 stated: “The vehicle can traverse a 60% grade and a 30% side slope.” ([34], p. 1). Typically, the left-right rollover threshold is much less for “side slope” than the front-back threshold for “grade”. Therefore, any solution utilizing a desirability or traversability map should assign a higher traversability to a sloped surface it will be required to traverse parallel to the slope, versus a sloped surface it will be required to traverse perpendicular to the slope.

XI.I. Experiment with different LIDAR configurations

By not orienting the sensors so that they intersected the ground at a fixed distance from the vehicle, Team 2005-06 was able to make effective use of LIDAR sensors by detecting obstacles as far from the vehicle as possible, and by using an oscillating mount, Team 2005-06 was able to reduce the number of sensors to the minimum necessary to accomplish this with some redundancy. The author considers this a key distinguishing factor which differentiated Team 2005-06 from all other teams which participated in the 2004 QID or GCE or 2005 GCE, and which contributed to Team 2005-06 successfully

completing the 2005 GCE. See paragraph VIII.E. The use of simulation might allow alternate mounting configurations to be objectively evaluated, revealing which are of interest to further study.

XI.J. Extend the maximum effective range of high-quality sensors

Extend the obstacle detection range of high-quality sensors to enable the controlling intelligence to detect obstacles at ranges consistent with speeds an autonomous vehicle may reasonably be expected to travel. For example, in general highway speed limits in the United States are between 60 and 70 mph. However, the maximum effective ranges of sensors in use by teams participating in the 2004 and 2005 GCE correspond to a maximum speed of 47.6 mph (VISION sensors), 40.2 mph (RADAR sensors), 36.0 mph (long-range LIDAR sensors), and 25.5 mph (short-range LIDAR sensors).

In addition, no team reported a maximum speed greater than 38.0 mph. The maximum speed reported by Team 2004-10 during the 2004 GCE was 36 mph ([68], p. 31) and the maximum speed reported by Team 2005-16 during the 2005 GCE was 38.0 mph ([51], p. 688). The maximum reported speed corresponds to a maximum effective range of 44.6 m, between the maximum effective ranges for RADAR and long-range LIDAR sensors.

Extending the maximum effective range of high-quality sensors will be necessary before an autonomous vehicle will be able to achieve speeds consistent with general highway speed limits.

XI.K. Use alternate speed setting strategies

Implement a controlling intelligence that wants to drive as fast as it can, and in the most direct bearing to goal. The nodes in this example would exert a “negative pressure”, that is, they would exert the equivalent of a braking force to the autonomous vehicle as it attempts to drive with the throttle wide open, or the equivalent of a third hand on the steering wheel providing a change in bearing. The resistance “felt” by the steering wheel or gas pedal to negative pressure would be tuned to circumstances in the local environment. For example, under normal driving conditions at high speed, the controlling intelligence would resist minor pressure at high speeds, but not low speeds; under normal driving conditions at low speeds, the controlling intelligence would experience the equivalent of a driver in the passenger seat reaching across to suddenly grab the steering wheel and change course to radically alter bearing, or to prevent the controlling intelligence from turning into an obstacle.

XI.L. Make provisions to maintain the published record

The 2004 and 2005 GCE made extensive use of the Internet to solicit participation, provide access to team resources, publish requirements, and present results. This was a deliberate decision on the part of DARPA. DARPA stated: “DARPA developed a website devoted to providing information about the Grand Challenge... Interested participants and entrants used the website to communicate directly with DARPA. The website contained a discussion forum that participants used to share ideas about technical approaches for autonomous ground vehicles, including obstacle detection, navigation and position location, sensing, control software, and vehicle components.” ([3], p. 3).

In general, this was a successful strategy. DARPA used the Internet effectively to communicate with teams and the public prior to the 2004 and 2005 GCE. However, the published record is rapidly disappearing. For example:

- DARPA made resources and references available to teams participating in the 2004 QID or GCE or 2005 GCE via the Grand Challenge Website, such as several versions of the 2004 GCE rules, a “description of the mandatory subjects to be addressed” in the team technical proposal, and the 2004 QID and GCE RDDFs. The Grand Challenge Website was substantially redesigned prior to the 2005 GCE. DARPA re-published portions of the Grand Challenge Website as the Archived Grand Challenge 2004 Website, but did not retain all published records. As a result, the Archived Grand Challenge 2004 Website is itself an incomplete record of events.

- Some teams which participated in the 2004 or 2005 GCE have since disappeared entirely from the Internet, leaving traces only in resources and references published by DARPA, or press about the Grand Challenge. Some of the teams which have since disappeared and which participated in the 2005 GCE did not publish their results via the Journal of Field Robotics. As a result, published records of their activity are practically non-existent.
- Some companies formed at the time of the 2004 and 2005 GCE to provide engineering or other services to teams participating in the 2004 QID or GCE or 2005 GCE have since disappeared.

At best, the Internet is an ephemeral resource. Future research which makes extensive use of the Internet should establish requirements for the maintenance of a permanent record of events as part of the published record.

In addition, DARPA established no requirement to publish in an academic journal or similar publication, and there is no evidence that DARPA required teams which participated in the 2004 and 2005 GCE to maintain records of their activities that would allow future researchers to re-construct team challenge vehicles. The author considers it likely other teams, in particular teams with a primary group identity of “Academic”, maintain repositories similar to the repository maintained by Team 2004-20, but these are of limited utility as they were not published.

DARPA intended team technical proposals to be the official published record of the 2004 and 2005 GCE. Prior to the 2004 QID or GCE, DARPA stated: “Publication of

the technical completion of [*sic*] papers after completion of Challenge [*sic*] will ensure they become part of the legacy of this event. They will be the primary mechanism from which knowledge gained from this event is utilized in future research and development. The technical paper does not need to be so detailed that someone could immediately build the vehicle themselves, but it should be detailed enough to teach an interested individual about the design.” ([69]). However, 2004 and 2005 team technical proposals provided insufficient technical detail and contained many errors, omissions, and inconsistencies which caused the author to conclude that they were unreliable as records let alone the “primary mechanism from which knowledge gained from this event is utilized in future research and development”.