

Road Runner: An Autonomous Vehicle for HRI Research

Ross Mead^{*}, Jerry B. Weinberg^{*}, Jenna Toennies[±], Jeffrey R. Croxell⁺, Bryan Adams⁺, George Engel⁺, John Hiatt⁺, Nick Italiano⁺, Ryan Krauss[±], Aaron Backs[±], and Matt Gorlewicz[±]

Southern Illinois University Edwardsville
Department of Computer Science^{*}, Department of Electrical Engineering⁺, Department of Mechanical Engineering[±]
Edwardsville, IL 62026-1656
mini-darpa@siue.edu

Abstract

Road Runner is an autonomous electric vehicle that was developed to compete in the Mini Grand Challenge, which requires traversing paved pathways on a college campus. A part from the competition, Road Runner was designed as a platform to explore areas of human-robot interaction between the passenger and vehicle. Road Runner is a one person electric golf cart that was retrofitted to provide basic control services of acceleration, braking, and steering. The states of these services are communicated to and queried by a laptop controller that is mounted for easy viewing and access to the passenger. Road Runner also relies on its extensive sensing capabilities in order to avoid obstacles, follow the path, and interact with the crowd. An integral part of path finding involves using a color webcam to sample the color of the path in front of the cart. By sampling, Road Runner can adapt to different path colors, so it can easily go from blacktop to concrete or any color of paved material. Using the sample color the image is processed to distinguish the paved pathway. Overall, Road Runner provides a unique opportunity for research in both autonomous and shared modes of transportation.

Introduction

Some of the key goals of the DARPA Grand Challenge include an increase in safety and efficiency of automobile travel (Barry 2008, DARPA 2008). With several teams having completed the desert and urban Grand Challenges, these goals are starting to be realized. A next step in the field of autonomous transportation considers a passenger in the vehicle, and how this passenger and the vehicle would interact. Communicating a destination or series of destinations to the robot, receiving feedback and contextual information from the robot regarding the surrounding area, and determining when it is appropriate for a human to take control of the vehicle are just a few of the problems that must be considered.




Figure 1: Road Runner with sensor and actuation subsystems. A Lenovo Thinkpad T61 serves as the controller (not shown).

The motivation for our project began with the Mini Grand Challenge competition held at Penn State Abington (Avanzato 2007). The competition is similar to the DARPA Grand Challenges, but the size of the vehicles and the length of the path have been scaled down to demonstrate similar technologies on low-cost platforms. Each autonomous robot entry must avoid obstacles while traversing a college campus—on both paved walkways and off-pathway—to reach a series of given GPS waypoints. A truly unique aspect of this competition is that of crowd involvement—the event often draws quite a large group of followers, and robots are evaluated based on how they interact with onlookers. While the Mini Grand Challenge aims at low-cost vehicles, entries thus far are too small to transport a person (e.g., see Dodds 2008). Conversely, many of the vehicles in the DARPA Grand Challenge are capable of carrying a significant payload, but are large and expensive. Our entry in the Mini Grand Challenge competition is a small and inexpensive one-person electric golf cart, which has been dubbed Road Runner (Figure 1). Road Runner provides a nice middle-ground mode of transportation allowing a single passenger a comfortable ride with full interaction capabilities via a laptop interface.



Figure 2: Robotics team with completed retrofitted cart. Note that the driver space is maintained with clear access to the laptop.

Robot Platform Overview

The golf cart, a Taotao ATE-810 produced by TaoTao Group (www.zj-trading.com), originally retailed for \$600. The electric cart runs on three 12-volt batteries in series (36-volt total). These were additionally regulated to 12-volts to provide power for various items such as the laptop and speakers. The cart has both forward and reverse drive. The accelerator is electronically controlled through a potentiometer connected to pedal. The cart has a hydraulic, rear wheel disc brake.

The golf cart was retrofitted to provide motor control, braking, and steering. The steering wheel was removed and a gear was mounted to the steering column. A steering motor was mounted to the golf cart frame with a second gear attached to provide position control of the steering column. Removing the steering wheel provided space to mount an automobile laptop stand that would allow a clear view to the screen and easy access to the keyboard for a passenger (Figure 2). Braking was accomplished by mounting a linear actuator to the frame, and acceleration was provided with a direct electrical signal.

The sensor suite included a sonar array, differential GPS, shaft encoders, and a color camera. The motor services and the sensing capabilities are described in more detail in the following sections.

Motor Control Services

To autonomously navigate the campus, Road Runner functions from a basic services model with a communication interface via I²C bus. The basic control services are acceleration, braking, and steering. The states of these services are maintained by on-board PSoC microcontrollers, made by Cypress Semiconductor, using shaft encoders with resolution of 1000 counts per revolution for feedback control. For the steering of the cart,

the encoder count is used to determine the position of the drive shaft. This position can then be translated into a steering angle. A DC-540 motor was used in conjunction with a pair of custom spur gears, radii 2.5 inches and 0.58 inches, to turn the drive shaft in either direction. This resulted in a gear ratio of 4.3:1 which provided a maximum amount of torque output given the size constraints of the cart's frame. By applying proportional control to the motor, a desired steering angle can be given, and Road Runner can reach its desired position within $\pm 5\%$. The speed of the cart is also determined through encoder counts. An optical encoder mounted on the rear shaft is used to determine an accumulated count during a given time interval. Using the circumference of the wheels, this count can be translated into a distance. Thus, Road Runner's velocity can be tracked at all instances. It is also essential that this velocity can be adjusted to allow Road Runner to traverse different grades. Thus, velocity control was implemented such that the brake is either compressed or eased depending on if the cart is traveling too fast or too slow, respectively. A 35-lb. force, 6-inch stroke linear actuator was mounted parallel to the existing hydraulic brake to provide the necessary braking needs; a flat, steel lever approximately 18 inches in length connects the actuator arm to the brake itself. Using feedback from an internal encoder, the actuator arm moves in or out, allowing for acceleration or deceleration of the cart. Not only does Road Runner possess efficient control services, but it also has querying capabilities. Through these services, the cart can request both its velocity and its steering angle over the I²C bus. In this manner, Road Runner can keep track of its relative position, as well as its progress in reaching a new position at its discretion.

Sensing Capabilities

The Mini-Grand Challenge (Avanzato 2007) requires that an electric vehicle traverse the paved campus pathways to reach five given GPS waypoints; a sixth waypoint is given that requires off-path traveling. To make path-finding easier, orange cones are placed at intersections blocking wrong turns so that only the correct choice is clear of these obstacles (Figure 3).

In addition to reaching the goal waypoints, there are rules to ensure human and robot safety. The campus pathways are not closed off to pedestrian traffic, so the rules require that the cart stop within three feet of any obstacle until the pathway is cleared. Further, the vehicle's speed must not exceed 5 miles per hour.

As noted in the Motor Control Services section, a shaft encoder was placed on the steering motor to provide position control for the steering column; the linear actuator for the braking had a built-in encoder also for position control. A shaft encoder was added to the rear axle to provide velocity control.



Figure 3: Example of campus pathway intersection (Avanzato 2007).

Two Garmin 17x HSV GPS units (www.garmin.com) were mounted to the top of the cart to provide robust differential waypoint detection (i.e., “intelligent redundancy”). In addition to detecting waypoints, these units were also used for crowd entertainment. Part of the scoring in the competition requires entertaining crowd interaction. Points of interests were added to the system, such as buildings or statues. Whenever the cart detected it was near a point of interest, it would stop, announce a relative bearing (e.g., “ahead and to my right”), and play a voice-synthesized message about it to the crowd (i.e., name of the point, its creator, and other contextual information).

An array of four sonars was mounted to the front bumper for obstacle detection; a fifth sonar was mounted on the rear bumper for the potential of adding other types of crowd interactions. (A SICK outdoor laser rangefinder was recently purchased to replace the front sonars.)

A color webcam was mounted beneath the roof of the cart for path-finding. Using the camera, points immediately in front of the camera were sampled for the color of the pathway. This allowed path finding to adapt to changes in the path color. For example, in some places on our campus (SIUE), pathways transition from blacktop to white concrete. Once the color of the path was determined, then a series of filters were applied to the camera image to find the blob that pertained to the path immediately in front (Figure 4). This was accomplished using the Open Source Computer Vision Library (OpenCV). Similar filtering techniques have been used in previous entries (Dodds 2008). Detection of the orange cones was used to eliminate sections of the frame.

After the image is processed, the detected pathway is segmented into rows of a pre-specified height (Figure 4). The centroid of each segment of the path is used as a local waypoint, and a path trajectory is determined by connecting these centroids. The cart is then steered based on its immediate trajectory—the relationship between the bottom-center of the frame and the next-closest centroid.

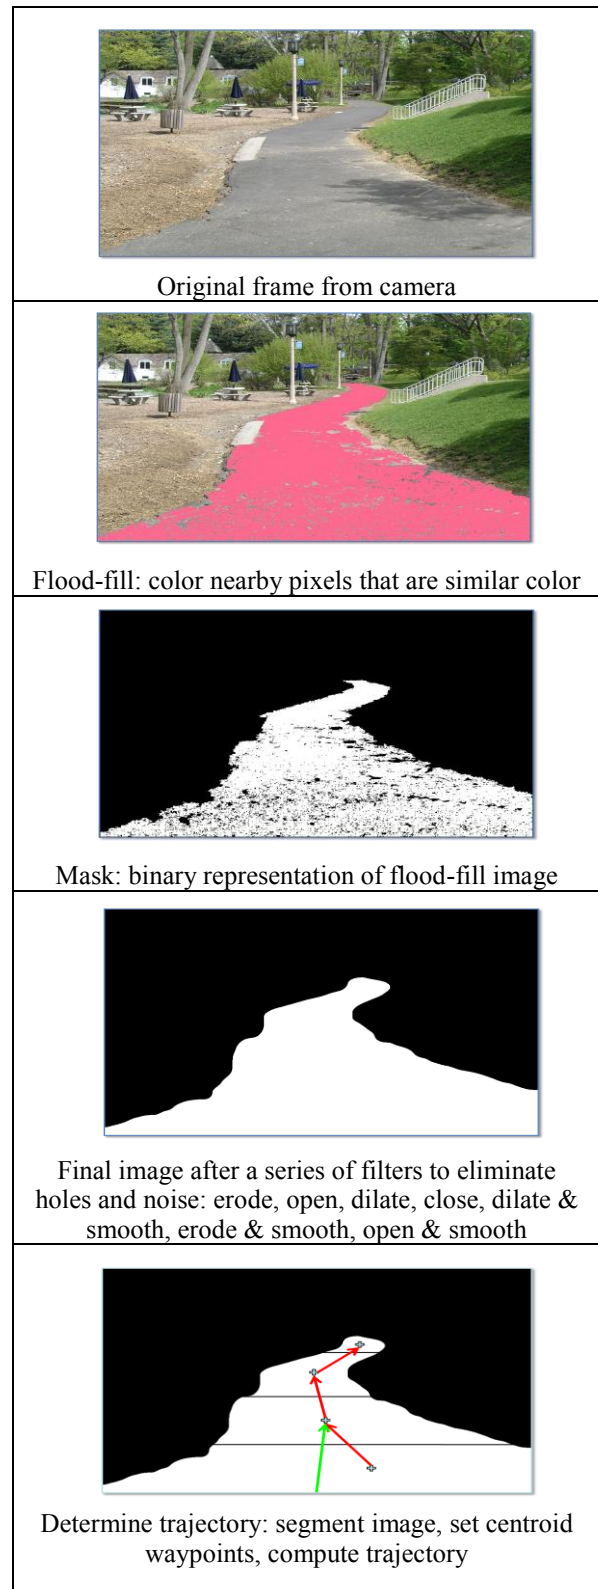


Figure 4: Image processing using Open-CV to find the path and then to determine the trajectory.

Orange cones are identified using simple color segmentation and blob detection. An avoidance trajectory is determined by connecting the centroids of the lowest and highest orange blobs in the image, where the relative angle between them produces a proportional steering output.

Software Architecture

Knowledge Representation

Road Runner considers the world as of a collection of state vector interfaces, each describing a property of the cart relative to an aspect of the environment. A state vector stores a representation of the measured quantity (e.g. distance reading, trajectory, displacement, angle, etc.), a timestamp, a confidence value, and a flag (used primarily to indicate valid data). States are updated asynchronously by threads, which query and preprocess the relevant sensors (i.e., converting raw sensor data into the proper interface representation) when appropriate; for the most part, the state will only be updated if the data is valid (i.e., unless an error flag is thrown), in which case this new data is timestamped. The state interfaces described here are based on the implementation for the Mini Grand Challenge; however, note that interfaces can easily be added or removed, depending on the application domain.

For each waypoint provided, a state interface exists, containing a distance and bearing based on the cart's relative pose to the destination. Two GPS devices are used to determine the robot's global pose (position and orientation). Each GPS device is updated in its own thread. Data is then combined through weighted averaging based on confidence, as well as a moving average based on previous data; if a device returns invalid data, its value is ignored. Once a pose is determined with high confidence, the robot then calculates its relative geodesic distance and bearing to each waypoint; this value is stored in the corresponding state, along with a timestamp and confidence value (conveniently provided by each device). States corresponding to waypoints are sorted in the desired sequential order of arrival (the reasoning for this will be described in the Behavior-based Control section).

Five state interfaces exist for each sonar. A single thread exists to update all of the sonars; each sonar is updated asynchronously and in an alternating order to eliminate crosstalk. The thread sends an update command (i.e., ping) to each sonar over I²C, and then later polls the sonar for its response (i.e., reads its data from a register). These distances are converted to inches and are factored into a moving average for storage.

Though there is currently only a single camera input (and a single thread processing it), multiple interfaces are used for processed camera data. One state is dedicated to the resultant orange cone avoidance trajectory described above, while other states represent as much of the pathway as possible. In the Sensing Capabilities section, we

discussed generating a path trajectory by segmenting an image into a series of rows. For each of the connected line segments generated by this method, a state interface is used. Each line corresponds to a proportional steering angle and length of travel along a section of the sensed path. Though Road Runner primarily reacts to the immediate trajectory, these additional trajectories improve accuracy and robustness in spite of camera limitations and failures, providing a sort of "memory" of what the path looks like ahead. This will be discussed further in the Behavior-based Control section.

There are three states that maintain the robot's measured velocity, steering angle, and brake position. A single thread queries the onboard microcontrollers for the appropriate data. This information is then used for state estimation across all other states. Using linear and angular velocities, and the difference between the current time and the timestamp of each state, the thread produces a new hypothesized value of the state given the movement action. This was an integral part of Road Runner's success, as it provided a rather accurate model of how the world changed based on how the robot was moving. Time between sonar updates was short, but fundamentally flawed in that an object that was previously considered far enough away before might not be far away as the cart is quickly approaching it; time between GPS updates could cause jerky steering behavior as the robot responds to old pose data; slow camera response could cause the robot to steer off course. By incorporating the cart's sensed actions into the state, these problems are eliminated.

A final collection of three states is dedicated to velocity, steering, and braking commands from a human operator. At any time during the cart's progress, the user can take full control. A computer interface offers acceleration and deceleration (in the form of both braking and natural velocity reduction), cruise control, and steering with auto-centering. User input sets the appropriate state to the desired value, overriding any form of autonomous control. In regards to basic HRI, we are considering a shared approach, where the user may specify a velocity for Road Runner to operate at as it approaches its goals, as well as a steering aid, in the case of unforeseen obstacles and for purposes of learning.

Behavior-based Control

Once raw sensor data has been processed and converted into the appropriate sensor interfaces, Road Runner must make decisions in terms of what actions should be taken. Conceptually, a movement action takes the form of a linear velocity and an angular velocity; at the actuator interface level, this corresponds to commands to the drive, steering, and brake motors. In addition, Road Runner may enqueue phrases into the speech system, which are processed and executed asynchronously in their own thread.

Movement commands are generated by a series of independent behaviors, which were developed based on the

robot’s anticipated interactions with the environment in this competition. A behavior is given a name and a processing function with optional parameters. Each behavior has mutually exclusive access to all state interfaces, but will only consider data that is deemed relevant by its function definition. The function determines whether or not the behavior is active, returning the appropriate motor action during each step of the robot’s execution.

A behavior engine manages all behaviors and their respective outputs. As these behaviors are independent of each other, it is highly likely that conflicts in desired motor output will occur. To resolve these discrepancies, the engine utilizes a subsumption architecture. Figure 5 shows the behaviors used for the Mini Grand Challenge in increasing levels of competence (described below).

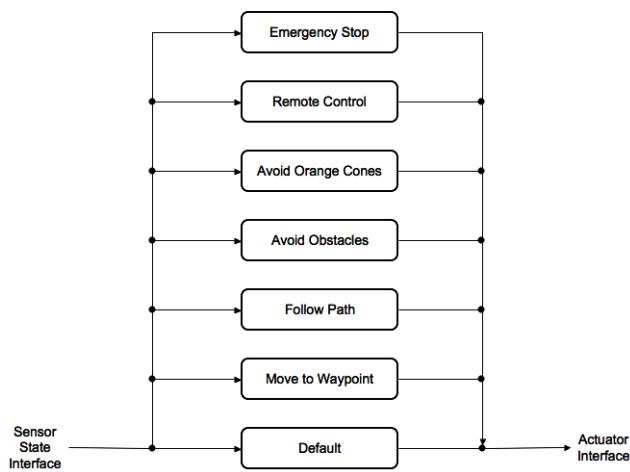


Figure 5: Subsumption architecture used for resolving behavior output.

0. Default. This was the first and most basic behavior created. Originally, it would simply try to keep the cart stationary (without locking the brake). However, for the purposes of the competition, we decided that, if any waypoints remained unreached, it always be best that the robot do *something*—with any luck, a few slight movements would cause other behaviors trigger; the worst that could happen is that we would be disqualified (as opposed to sitting there and running out of time). Thus, we opted to have the robot wander very slowly in the direction of “least resistance”, based on the states of the front sonars. Likewise, we nicknamed this the “Do Something” behavior. Fortunately, during the competition, Road Runner never exhibited this primitive behavior, as it would have been an indication of serious subsystem failure.

1. Move to Waypoint. As mentioned in the previous section, states that correspond to waypoints are sorted in the desired order of arrival. Thus, Road Runner is able to generate linear and angular velocities proportional to its preprocessed relationship to the next waypoint. This behavior will output the appropriate movements to get the cart to the waypoint, assuming no obstructions or terrain restrictions. Once the destination is reached, it will be removed from the list of waypoints, and Road Runner will seek the next one. If the robot senses that it has arrived at a waypoint that comes after one or many of those that it should have already arrived at, it assumes that it has, indeed, already arrived at those waypoints as well, and removes them from the waypoint list. This was crucial in the Mini Grand Challenge, as part of the path took Road Runner through an area with dense tree coverage, resulting in invalid GPS data and the potential to “miss” a destination; by assuming that previous points have been reached if a waypoint is in proximity, the robot is able to continue along the path and even transition off-road when appropriate. However, it is noted that, for other applications, this might not be as desirable.

2. Follow Path. This behavior outputs drive and steering commands proportional to the most immediate local trajectory of the state interfaces. Recall that the linear and angular velocities of the robot are factored into each trajectory (state). In the case of high camera latency, this “memory” of the path ahead generates a sort of “reactive pipeline” for the cart—it can still react based on its estimate of how its relationship to the path has changed given the robot’s movement. This is analogous to a person observing a living room (i.e., with furniture and, most likely, other clutter) from one corner, closing his or her eyes, and then moving to the opposite corner, relying completely on spatial memory.

3. Avoid Obstacles. Campus paths remained open for the duration of the Mini Grand Challenge. Thus, competition rules mandated that, while on the paved pathway, the robot stop within two feet of an obstruction (likely, a judge); however, while traversing off-road to reach the last waypoint, the robot must circumvent any obstacles in its way. This avoidance behavior was accomplished using the array of sonars on the front of the cart. While on the path, Road Runner will decrease its speed proportionally as it

approaches an obstruction, and eventually stop once it is within proximity. Once Road Runner travels off the beaten path, it uses sonar distance information to generate a proper turning angle and, if necessary, a decreased drive speed. While adequate for purposes of the competition, it is expected that the integration of the SICK laser rangefinder will improve performance.

4. **Avoid Orange Cones.** A line of orange cones illustrates a special case of obstacle avoidance. These cones are meant to be an aid to the robot by blocking or directing parts of the pathways that are not part of the campus tour, such as side-paths and forks. A state interface is dedicated to the suggested perceived avoidance trajectory; the behavior uses this trajectory to generate a steering angle that will guide the robot parallel to the cones. Thus, Road Runner is able to avoid the obstacles while staying on the path.
5. **Remote Control.** In the case that a movement command is given by the user, this behavior supersedes all autonomous behaviors beneath it.
6. **Emergency Stop.** This behavior becomes active only if the emergency stop button is pressed. All robot activity ceases and the brake is fully engaged until the button is pressed again.

Future Work

While the motivation for the Road Runner project was competing in the Mini Grand Challenge, another goal of the project was to provide a platform for exploring human-robot interaction in autonomous vehicles. The project was successful in this regard. Road Runner maintained most of the original passenger space. Lower-level services of braking, acceleration, and steering were successfully automated, with higher-level decision making, such as navigation and obstacle avoidance, also implemented.

Questions we are exploring in HRI are how does a passenger specify their location and their goal, how does a passenger interrupt the trip for whatever reason and then resume it, how does the passenger intervene in a potentially dangerous situation (such as avoiding a pothole), and when is it (if ever) appropriate for the robot to ignore the control of the human?

Our initial work in this area is to leverage Google Earth as an interface to specify the goal or even the route (Figure 6). The passenger can indicate the path they prefer to take between two points. GPS points can be extracted to determine intermediate waypoints as well as the goal point. These would be used to navigate the pathways. Currently Road Runner assumes that there will only be one unblocked path at an intersection, so additional behaviors will need to be added to decide the correct pathway to take.

A simple approach to this would be to extract waypoints for the entrance and exit of an intersection.

Beyond specifying where to go, the passenger may want to interrupt the trip to talk with a friend, indicate to the vehicle that there is a pothole to go around, or let the vehicle know that the route is blocked due to construction and a new route must be plotted. This type of HRI puts the passenger in a type of supervisory role. It will require exploring a mix-initiative interface where the human passenger can communicate commands or advice to the vehicle (Adams, et al 2004, Hong et al. 2007).




Figure 6: A route specified by the passenger on Google Earth between the Engineering Building (point A) and the Morris University Center (point B) on the SIUE Campus.

References

- Adams, J.A., Rani, P., and Nilanjan, S. 2004 “Mixed-Initiative Interaction and Robotic Systems”, Technical Report of the AAAI-04 Workshop on Supervisory Control of Learning and Adaptive Systems (WS-04-10), pp. 8-13, 2004.
- Avanzato, R. 2007 “Mini Grand Challenge Contest for Robot Education”, Technical Report of *The 2007 American Association of Artificial Intelligence Spring Symposia (SS-07-09)*, pp. , Stanford, CA, March 2007
- Barry, K. 2008 “DARPA Urban Challenge”, *Robot Magazine*, Spring 2008, (10), pp.24-28
- DARPA: <http://www.darpagrandchallenge.com/>. June, 23rd, 2008.
- Dodds, Z. 2008 “Leveraging Laptops: Resources for Low-Cost Low-Level AI”, *Proceedings of the Twenty-First International FLAIRS Conference*, Miami, pp. 234-239, FL, May 15-17, 2008.
- Hong, J., Song, Y., and Cho, S. 2007 “Mixed-Initiative Human-Robot Interaction Using Hierarchical Bayesian Networks”, *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 37, No. 6, pp. 1158-1164, November 2007.