

Autonomous driving in structured and unstructured environments

Conference Paper**Author(s):**

Kolski, Sascha; Ferguson, Dave; Bellino, Mario; Siegwart, Roland

Publication date:

2006

Permanent link:

<https://doi.org/10.3929/ethz-a-010079280>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/IVS.2006.1689687>

Autonomous Driving in Structured and Unstructured Environments

Sascha Kolski*, Dave Ferguson[†], Mario Bellino*, and Roland Siegwart*

*Ecole Polytechnique Federale de Lausanne (EPFL)
Lausanne, Switzerland

{sascha.kolski, mario.bellino, roland.siegwart}@epfl.ch

[†]Carnegie Mellon University (CMU)
Pittsburgh, PA, USA

dif@cmu.edu

Abstract—Autonomous vehicles are being used increasingly often for a range of tasks, including automated highway driving and automated parking. Such systems are typically either specialized for structured environments and depend entirely on such structure being present in their surroundings, or are specialized for unstructured environments and ignore any structure that may exist. In this paper, we present a hybrid autonomous system that recognizes and exploits structure in the environment in the form of driving lanes, yet also navigates successfully when no such information is present. We believe such an approach is more flexible and more robust than either of its sub-components alone. We demonstrate the effectiveness of our system on both marked roads and unmarked lots.

I. INTRODUCTION

In Europe, over forty thousand people were killed through road accidents in the year 2000, with another 1.7 million injured [1]. The great majority of these accidents were due to human error, with only 5% caused by defective vehicles. Such staggering findings motivate the use of driver assistant systems and fully automated vehicles to increase driver and passenger safety.

Driver assistant systems can help drivers to identify dangerous vehicle states and traffic scenarios and reduce the risk of accidents. These driver assistant systems are widespread in all categories of vehicles and range from anti-lock brakes to radar based adaptive cruise control. The development of these systems has been accelerated by integrated drive-by-wire components such as electronic gas pedals, brakes, and steering systems.

The development of such components has also hastened the arrival of autonomous passenger vehicles. In 1997, the NavLab vehicles travelled ‘no hands’ across the United States, requiring only accelerator and brake pedal interaction from the driver [2]. In 2005, 23 autonomous vehicles started a race across the Nevada desert in the DARPA Grand Challenge race [3], with 5 of them finishing the 211.1 Km distance.

Most of these systems depend on environmental structure like driving lanes or dense sets of GPS points. However, in many common driving scenarios neither of these sources of information will be available, for example, when leaving a road and entering a parking lot.

Autonomous navigation in unstructured environments is an active research area in field robotics, and a number of effective approaches have been developed that address this task [4]–[7]. A common technique is to maintain a map of the environment

and use this to plan safe paths to a desired goal location. As the vehicle traverses the environment, it updates its map and path based on its observations. Such an approach works well when dealing with reasonably small areas, but storing and planning over maps of the entire environment is impractical when traversing long distances. Further, without taking into account non-spatial information such as road markings, these approaches are unable to ensure that the vehicle stays within its lane (or even on the road) when navigating through highway or urban environments.

In this paper we present a hybrid navigation system that combines the benefits of existing approaches for driving in structured environments (e.g. roads) and unstructured environments (e.g. parking lots). When driving on detectable roads, the system uses visual lane detection and laser range data to generate a local map, which is processed by a local planner to guide the vehicle down the lane while avoiding obstacles. When driving in unstructured environments, the system employs a global map and planner to generate an efficient trajectory to a desired goal. The combined system is capable of navigating a passenger car to a given goal position without relying on road structures, yet it makes use of such structure when it is available.

We begin by introducing our test vehicle and its sensing and computing hardware. We then describe our system for navigating structured and unstructured environments. In Section VI we present results from runs performed in road and parking lot scenarios. We conclude with discussion and future work.

II. VEHICLE AND SENSORS

Our vehicle is a Smart *fortwo* passenger car that has been modified for autonomous operation. Firstly, we have interfaced the Smart’s controller area network (CAN) bus to access data on the dynamic state of the vehicle, specifically the wheel speed and the steering angle. We have also added actuators to the accelerator, the brake pedals, and the steering column. Finally, a number of sensors (discussed below) have been added to provide vehicle and environmental information.

At the core of our control system is a standard laptop computer running the Linux operating system. It is connected to the entire sensor setup and the CAN-bus. Our software platform is based on GenoM [8], a real-time framework for Linux. Data acquisition is performed at a rate of 100 Hz for



Fig. 1. Our autonomous Smart car platform at EPFL. The white SICK laser range finder used for mapping and obstacle detection can be seen attached to the front of the vehicle. The camera used for lane detection is mounted below the rear-vision mirror inside the car (and is attached to the yellow/light-gray cable).

the wheel encoders and IMU, 20 Hz for the laser scanner and 25 fps for the camera.

A. Proprioceptive Sensors

As with many other passenger cars, the Smart is equipped with a variety of sensors which are linked using the vehicle's CAN bus. By interfacing this bus it is possible to access the sensor data and measure the vehicle's dynamic state precisely.

a) *Wheel Encoders*: The overall vehicle speed is derived from the four wheel encoders with a resolution of 0.5 revolutions/minute. The steering wheel angle is available with a resolution of 0.04° .

b) *IMU*: We have added a 6 degree of freedom IMU to the Smart that is able to measure angular rates up to $100^\circ/\text{sec}$ at a resolution of 0.025° . Lateral accelerations in all three dimensions can be measured up to $2g$ with a resolution of 0.01 m/s^2 .

B. Exteroceptive Sensors

c) *Laser Range Finder*: We use a SICK LMS 291 laser range finder for sensing the spatial structure of the environment. This is a configurable laser range finder based on time of flight measurements, with an angular resolution of 1 or 0.5° , an angular range of 180° and a measuring range up to 80 meters. The laser is mounted directly on the front bumper of the vehicle to detect obstacles.

d) *Monocular Camera*: An automotive gray-scale camera is mounted inside the vehicle at the top of the windscreen for observing the area in front of the vehicle and detecting lane information. The resolution of the camera is 750×400 pixels and it delivers information at 25 frames per second.

III. POSE ESTIMATION AND MAPPING

For accurate navigation and mapping it is necessary to have a precise estimation of the vehicle's position at all times. To estimate the position of our vehicle we have two main

sources of information. The Smart's built-in sensors provide an accurate approximation of the steering wheel angle and the translational velocity of the car. Additionally, our IMU provides the angular and translational accelerations of the vehicle.

Our position estimation is based on two Kalman filters [9] - one for the translational speed of the vehicle and one for the angular velocity. Both filters use the output of the IMU for the prediction step. The first filter uses the output of the wheel encoders to provide the translational speed update, while the second uses the steering angle to provide the angular velocity update. Combining these filters provides accurate position estimation over reasonable distances. See [10] for more details on this approach.

Given the vehicle's position, we can convert measurements from our laser range finder into world coordinates. This enables us to produce maps of the environment that can be used for navigation. Figure 5 shows a map of a rural road and parking lot created online during an autonomous traverse.

IV. DRIVING IN STRUCTURED ENVIRONMENTS

When driving in structured environments such as roads or highways, it is important for safety that vehicles abide by traffic rules and stay in their own lanes. For autonomous vehicles, such structure is useful because it constrains the available actions of the vehicle and reduces the complexity of the navigation task. For instance, if an autonomous vehicle is traveling down a road, it knows it must stay within its current lane so the lane can be used as a guide for where the vehicle must travel to next. Such an approach can be coupled with a standard commercial navigation unit that provides higher-level guidance on when to turn down which street.

However, to ensure safe navigation, it is not enough to just follow the current lane. The vehicle must be alert at all times and able to avoid other cars and objects that may unexpectedly place themselves in its path, such as cars pulling out from driveways or pedestrians crossing the street, for example. To achieve such behavior in our Smart, we construct a local map representing the immediate surroundings of the vehicle and then plan a collision-free path through this map. Both the map and the plan are updated frequently (at 20 and 10 Hz, respectively, for vehicle speeds up to 5 m/s). With both the local obstacles and lane information encoded in the local map, the vehicle is able to plan trajectories that keep it within the current lane and also avoid any obstacles.

A. Local Planning

We use the information from our laser range finder to construct a local grid-based cost map specifying the nearby obstacles and difficult areas to traverse for the vehicle. Each cell in the grid is a square of width 20 cm. Cells containing obstacles are assigned an infinite cost, representing untraversable areas, with cells corresponding to less difficult terrain assigned less-expensive cost values. We perform a configuration space expansion on this map, which has the effect of 'growing' the obstacles and other expensive areas out by the width of the

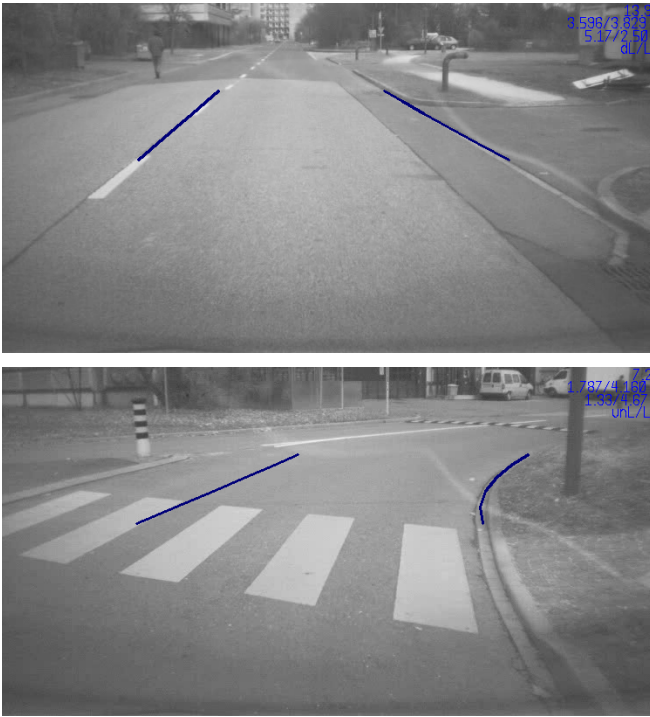


Fig. 2. Example results from our lane detection approach applied to images from a straight (top) and curved (bottom) section of road.

vehicle. This allows us to treat the vehicle as a single point during planning.

Given this local map and the current vehicle position and orientation within this map, we can then project potential vehicle actions onto the map and check the cost of these actions. We use a discrete set of variable-length arcs for our vehicle actions, corresponding to different steering angles and vehicle speeds [4]. Each of these arcs represents an action that is feasible from the current vehicle position, orientation, and velocity. We then choose the best of these arcs according to their costs and perhaps also some general objective, such as the amount of distance the arc takes us in our desired direction of travel (e.g. down the road). This arc can then be directly executed by the vehicle.

B. Lane Detection

Because the laser only recovers spatial structure from the environment, it is unable to provide information not manifested in spatial structure but e.g. color, such as lane markings. As a result, local maps constructed from only the laser will not indicate the boundary of the current lane or road.

To extract lane information, we use a monocular gray-scale camera designed for automotive use and a real-time lane detection algorithm running on a separate computer equipped with a frame grabber. Our approach combines hypotheses from several lane detection algorithms, each designed to detect different types of lanes, such as the closest lane to the vehicle, straight lanes, or curved or symmetric lanes. These algorithms rely mainly on the spatial gradient of the image to extract their hypotheses. The results of the individual algorithms are




Fig. 3. **Local Planning in Structured Environments.** The vehicle projects a set of feasible arcs through the local map from its current position and orientation (arcs for a single speed are shown in red/gray). The cost of each of these arcs is computed, based on the cost of the cells the arc travels through (darker areas are more expensive, with black cells representing obstacles). The best arc is shown in blue/black. The current lane information is shown as a series of expensive (dark gray) cells, while an obstacle appears in front of the robot as a series of black cells.

then combined to determine the most probable lane. Example results from our lane detection algorithm are shown in Figure 2 and more details on the algorithm can be found in [11].

Once the lane has been extracted, it is added to the local map as a high-cost area so that the vehicle will prefer to stay within the center of the lane but may move towards one side if it has to in order to avoid an obstacle.

An illustrative example of the combined approach is shown in Figure 3. Here, the darker an area, the more costly it is to traverse. Each small square in this map represents 20 cm and the vehicle can be treated as a point because the map represents its configuration space. The lane information is shown as a series of dark-gray cells. In this example, an obstacle also appears at the right side of the lane (this could be a person stepping into the lane). The obstacle is shown in black and must be avoided at all costs. The area adjacent to the obstacle is shown as high-cost area that should be avoided if possible. The robot picks the best arc (shown in blue/black) from its available set of arcs (shown in red/gray) according to the cost of the arcs and some higher-level objective (such as distance traveled parallel to the road).

V. DRIVING IN UNSTRUCTURED ENVIRONMENTS

In unstructured environments where there is no lane information to guide or constrain the actions of the vehicle, the vehicle must use a more general approach for navigation. For instance, imagine our vehicle has arrived at its intended destination address and now wants to park in a specified goal location within the parking lot. To do this, we could use the local planning component of our system and modify our higher-level objective so that the planner selects arcs that minimize the distance between the vehicle and its goal location. However, such an approach is susceptible to local minima, meaning that it can cause the vehicle to get ‘stuck’ behind obstacles that reside between its initial position and the




Fig. 4. **Global Planning in Unstructured Environments.** The vehicle projects a set of feasible arcs through the local map from its current position and orientation (arcs for a single speed are shown in red/gray). The cost of each of these arcs is computed, based on the cost of the cells the arc travels through (darker areas are more expensive, with black cells representing obstacles). A global path is planned from the end of each arc to the goal (shown as a filled circle on the right side of the map) and the cost of this path is added to the cost of the arc. The best arc is shown in blue/black, along with the global path from the end of this arc to the goal.

goal.

To avoid this limitation of purely local planners, our approach combines a global planner with the local planning capability mentioned in Section IV-A. This enables the vehicle to plan efficient trajectories through cluttered areas (such as parking lots) since the vehicle takes into account the entire map and all the obstacles during planning. This allows for early avoidance of obstacles and much better behaviour.

Our global planner is based on the Field D* algorithm, which has been incorporated into several fielded robotic systems [12]. This algorithm provides very low-cost paths through grid-based representations of an environment. These paths do not take into account the heading restrictions of the vehicle and instead approximate the least-cost path to the goal for a vehicle that can turn in place. Because Field D* does not encode the mobility constraints of the vehicle, it cannot be used alone for accurate trajectory planning for the vehicle. Consequently, we combine it with our local, vehicle-specific arc-based planner to provide feasible paths. Our combined system maintains a global map of the environment containing all the observed obstacles and high-cost areas. Then, every planning cycle, the vehicle projects out its set of available arcs into this map and computes the cost of each arc based on its distance and the cost of the cells it travels through, exactly as described in Section IV-A. This gives the cost of traversing the arc itself. To this value we then add the cost of a global path from the end of the arc to the goal. This cost is provided by our global Field D* planner. Then, the arc with the minimum combined cost is selected and executed by the vehicle.

Figure 4 shows an illustrative example of this combined approach. The set of available arcs are shown in red/gray, with the best arc shown in blue/black. Here, the best arc was selected based on a combination of the cost of the arc itself and the cost of a global path from the end of the arc to the goal




Fig. 6. Results from our lane detection and mapping in a structured environment. Data was gathered from roughly 100 meters of traverse down a road (travelling from left to right). The top image shows the combined local maps created by the vehicle during the traverse, with lane information shown as dark gray areas and obstacles shown in black. Notice that the obstacle information does not provide any real indication of the location of the lane or even road, and so does not suffice for safely guiding the vehicle. The bottom image shows a satellite map of the area.

(the goal is shown as a filled circle at the right of the figure). The global path from the end of the best arc to the goal is also shown in blue/black. In this example, a purely local planner would have selected the straight arc leading directly to the right, as this brings it closest to the goal in terms of straight-line distance. However, such an arc could cause it to get stuck behind the clump of obstacles in the middle of the map.

VI. EXPERIMENTS

We have tested our system in both structured and unstructured environments. For structured environments, we had the vehicle drive down a road and record the resulting local maps. Figure 6 shows the combined cost map constructed from the series of local maps and highlights both obstacles and lane information. Since the laser range data does not contain any information about the lane markings, the vision-based lane detection system is necessary to keep the vehicle in its lane.

To test our vehicle in unstructured environments, we gave it a more complex task. We began on a road and tasked it with autonomously navigating to a goal location in a nearby parking lot. Because there were large shrubs between its initial position and its goal, it was forced to travel down the road until it observed an opening through which it could enter the parking lot. At this point it entered the parking lot and navigated to its goal location.

Figure 5 shows the resulting map built by the vehicle and the vehicle's traverse. Figure ?? shows a series of images taken




Fig. 5. Results from global planning and mapping in an unstructured environment. Shown here is the map created from the laser during an autonomous traverse from an initial position on a rural road to a goal position inside a large parking lot. Also shown is the path (in blue/black) traversed by the vehicle. The vehicle began from the position marked in green/gray at the top of the map, and navigated to the goal position marked in red/gray at the bottom.

from a video of the traverse. Overall the vehicle travelled about 140 meters in 62 seconds, i.e. at average speed of roughly 2.3 m/s. It traversed clear parts of its path (including portions of the turn) at speeds of up to 4 m/s and slowed down significantly near obstacles.

The vehicle trajectory seen in figure 5 illustrates the advantage of our arc-based local planner. Because these arcs accurately represent the actual driving abilities of our vehicle, the resulting path is very smooth, especially in the wide turn the vehicle takes to enter the parking lot.

Together these experiments illustrate our vehicle's ability to navigate through both road and non-road environments. Our vehicle effectively avoids obstacles to reach a defined goal position without relying on an a priori model of the environment.

These experiments also showed that the complex tasks of data acquisition, vehicle localization, environment mapping, and local and global planning can be performed on a single standard laptop computer.

VII. CONCLUSIONS AND FUTURE WORK

In our paper we have presented a hybrid approach for autonomous navigation in structured and unstructured environments. As long as road structure like driving lanes are available our approach is able to integrate this structure into the mapping and path planning process. Whenever lane structure is not available, as is the case for example in parking lots, our system

is able to navigate based on observations from its laser and safely reach a given goal while avoiding obstacles like curbs and parking cars. We have provided results demonstrating the operation of the vehicle in both structured and unstructured environments.

We are currently working on a number of extensions to this approach. Firstly, we are incorporating a scan matching algorithm for improved position estimation and mapping [10]. This algorithm matches two consecutive laser scans to provide an approximation of the distance travelled by the vehicle between the scans. We are also working on an improvement to classical scan matching that can detect and model dynamic obstacles in the environment, which is useful when navigating through environments in which other vehicles and people are moving.

Secondly, although our current system can successfully navigate to some specified goal point, it cannot yet turn from road to road to arrive at an address-specified destination. Consequently, we are installing a GPS-enabled commercial navigation unit into the Smart car to provide high-level directions to the vehicle when navigating over roads. By coupling this unit with our local planner, our vehicle will be able to drive to an arbitrary destination using local roads and highways.

Another interesting research direction we are pursuing is using a 3D laser scanner to obtain a 3D model of the car's environment. With full 3D information we can create local and global maps for the robot that encode the slope of the terrain as

well as holes and other obstacles that cannot be seen with our current 2D laser [13]. This allows for navigation over rough outdoor terrain. We can also combine this with vision systems to produce textured 3D models of outdoor environments [14].

ACKNOWLEDGEMENTS

We thank the European Commission for partially funding this work through the SPARC project (Secure Propulsion using Advanced Redundant Control) under IST-507859. Dave Ferguson is partially supported by a U.S. National Science Foundation Graduate Research Fellowship.

REFERENCES

- [1] M. Shell, "Final report of the european esafety working group on road safety, online available," 2003. [Online]. Available: <http://europa.eu.int/information society/activities/esafety/indexen.htm>
- [2] C. Thorpe, T. Jochem, and D. Pomerleau, "The 1997 automated highway demonstration," in *1997 International Symposium on Robotics Research*, 1997.
- [3] "Darpa grand challenge race website." [Online]. Available: <http://www.darpa.mil/grandchallenge>
- [4] A. Kelly, "An intelligent predictive control approach to the high speed cross country autonomous navigation problem," Ph.D. dissertation, Carnegie Mellon University, 1995.
- [5] A. Stentz and M. Hebert, "A complete navigation system for goal acquisition in unknown environments," *Autonomous Robots*, vol. 2, no. 2, pp. 127–145, 1995.
- [6] O. Brock and O. Khatib, "High-speed navigation using the global dynamic window approach," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1999.
- [7] S. Singh, R. Simmons, T. Smith, A. Stentz, V. Verma, A. Yahja, and K. Schwehr, "Recent progress in local and global traversability for planetary rovers," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2000.
- [8] S. Fleury, M. Herrb, and R. Chatila, "Genom: A tool for the specification and the implementation of operating modules in a distributed robot architecture," in *International Conference on Intelligent Robots and Systems*, vol. 2. Grenoble (France): IEEE, Sept. 1997, pp. 842–848.
- [9] R. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME – Journal of Basic Engineering*, pp. 35–45, March 1960.
- [10] B. Jensen, J. Weingarten, S. Kolski, and R. Siegwart, "Laser range imaging using mobile robots: From pose estimation to 3d-models," in *Proc. 1st Range Imaging Research Day*, Zürich, Switzerland, 2005, pp. 129–144.
- [11] M. Bellino, Y. Lopez de Meneses, P. Ryser, and J. Jacot, "Lane detection algorithm for an onboard camera," *SPIE proceedings of the first Workshop on Photonics in the Automobile*, 2004.
- [12] D. Ferguson and A. Stentz, "Field D*: An Interpolation-based Path Planner and Replanner," in *Proceedings of the International Symposium on Robotics Research (ISRR)*, 2005.
- [13] M. Montemerlo and S. Thrun, "A multi-resolution pyramid for outdoor robot terrain perception," in *Proceedings of the AAAI National Conference on Artificial Intelligence*. San Jose, CA: AAAI, 2004.
- [14] H. Zhao and R. Shibasaki, "Reconstructing textured cad model of urban environment using vehicle-borne laser range scanners and line cameras," in *ICVS '01: Proceedings of the Second International Workshop on Computer Vision Systems*. London, UK: Springer-Verlag, 2001, pp. 284–297.