

HIGH SPEED HAZARD AVOIDANCE FOR UNMANNED GROUND VEHICLES IN EMERGENCY SITUATIONS

M. Spenko*

Stanford University, Department of Mechanical Engineering
Stanford, CA, 94305

J. Overholt

U.S. Army Tank Automotive Research Development and Engineering Center
Warren, MI, 48089

K. Iagnemma

Massachusetts Institute of Technology, Department of Mechanical Engineering
Cambridge, MA, 02139

ABSTRACT

Rapid hazard avoidance maneuvers will be required for unmanned ground vehicles operating at high speeds in rough changing terrain. Without rapidly decreasing speed in every situation, there is limited time to perform navigation calculations based on detailed vehicle and terrain models. This paper presents a novel method for high speed navigation and hazard avoidance based on the two dimensional “trajectory space,” which is a compact model-based representation of a robot’s dynamic performance limits on natural terrain. Simulation and experimental results on a small high-speed UGV demonstrate the method’s effectiveness

Traditionally, online planning problems have been performed either by selecting from a set of predetermined paths (i.e. search techniques over small spaces), or by reactive behaviors, which evoke a predetermined action in response to specific sensor signals. Since the majority of mobile robots have been designed for use on flat or slightly rolling terrain at speeds that do not excite vehicle dynamics, these techniques have not had to consider vehicle dynamics and vehicle/terrain interaction. This paper addresses the problem of navigation and hazard avoidance on flat, rough, and uneven terrain at speeds that excite the vehicle’s dynamics.

1. INTRODUCTION AND LITERATURE REVIEW

Numerous important military applications including surveillance and supply deployment require an unmanned ground vehicle (UGV) to move at high speeds (loosely defined here as speeds that excite vehicle dynamics such as side slip and rollover) through uneven, natural terrain with various compositions and physical parameters. Techniques to perform this successfully will become increasingly more important as UGVs transition from traversing dirt roads and open plains to more extreme terrain.

UGVs are commonly given a pre-planned path generated from topographical data to follow. In natural terrain at high speeds, it is likely that emergency maneuvers would be required due to unforeseen situations that may result from outdated topographical data, unidentified hazards due to sensor limitations or errors, or unanticipated physical terrain conditions. Despite increasing computation speed, in emergency situations, it is difficult to compute a new dynamically safe path and velocity profile using detailed vehicle and terrain models.

Previous researchers have used search-based techniques to navigate a HMMWV-class vehicle at speeds up to 10 m/s while avoiding large hazards (Coombs, *et al.*, 2000). The method relies on a pre-computed database of approximately 15×10^6 20 to 30 meter long clothoid trajectories. The vehicle is assumed to travel on relatively flat terrain at fairly low speeds, and thus the model used in the calculations does not consider vehicle dynamics. An online algorithm eliminates candidate clothoids that intersect with hazards or are not feasible given the initial steering conditions. From the remaining paths, the algorithm chooses one that follows the most benign terrain. Several contenders of the 2005 DARPA Grand Challenge utilize similar approaches that have been successful at speeds in excess of 17 m/s. Despite this, the technique does not consider the important aspects of terrain roughness, inclination, and vehicle/terrain traction characteristics, all of which will become increasingly more important as autonomous vehicles move from traversing roads and relatively benign terrain to more dangerous and extreme topography.

Other researchers have developed a fuzzy logic-based algorithm for reactive outdoor hazard avoidance (Daily, *et al.*, 1988; Olin, *et al.*, 1991). The approach arbitrates between hazard avoidance and goal seeking and allows

for UGV navigation at speeds up to 1 m/s. Another successful reactive behavior-based technique was developed where the “behaviors” are candidate steering angles, and an arbitrator chooses a steering angle based on hazard and goal locations (Kelly and Stentz, 1998). Other work in the area has focused on problems arising from partially known and dynamic environments (Laugier, *et al.*, 1998) or sensing issues in outdoor terrains (Langer, *et al.*, 1994). Although these techniques have been successful at low to moderate speeds, they do not explicitly consider vehicle dynamics and changing terrain characteristics.

This paper presents a hazard avoidance method that considers vehicle dynamics, terrain parameters, and hazard properties. The technique is computationally efficient enough for high-speed applications and has similarities to the dynamic window method for low-speed collision avoidance in structured environments (Fox, *et al.*, 1997). It incorporates features that are critical to UGV navigation, such as vehicle/terrain interaction, the presence of hazards, and terrain roughness and unevenness. The algorithm relies on the trajectory space, a compact framework for analyzing a UGV’s dynamic performance on uneven, natural terrain (Spenko, *et al.*, 2004). In addition, an algorithm is presented here for trajectory replanning after a hazard avoidance maneuver has been enacted. The effectiveness of the proposed hazard avoidance and replanning algorithms is demonstrated through experimental results of a UGV moving at high speeds over flat and sloped terrain. It is shown that the algorithms operate favorably in harsh, real world conditions.

2. PROBLEM STATEMENT AND ASSUMPTIONS

Here a UGV is assumed to be following a nominal preplanned path. The algorithm’s goal is to rapidly plan maneuvers that allow the UGV to avoid unexpected hazards while considering vehicle dynamics, steering dynamics, vehicle/terrain interaction, and vehicle performance limits. After the hazard avoidance maneuver is complete, the algorithm must efficiently resume the nominal path, again considering the above factors. It is assumed that the UGV has a control algorithm that can sufficiently keep it on the desired path.

Hazards are defined as discrete objects or terrain features that significantly impede or halt UGV motion, such as trees, ditches, and areas of poorly traversable terrain (e.g. very soft soil). It is recognized that hazard detection and sensing are important aspects of UGV mobility and an active research topic (Fish, 2003), (Shoemaker & Borenstein, 2000); however, it is not a focus of this work. As such, it is assumed that the hazards can be accurately recognized by sensors once within a given range of several vehicle lengths.

A terrain patch is described by its average roll (ϕ), pitch (ψ), roughness (σ), and traction coefficient (μ). It is assumed that coarse estimates of the tire/ground traction coefficient and ground roughness are known or can be determined online using currently available techniques (Arakawa & Krotkov, 1993), (Iagnemma, Kang, Brooks, & Dubowsky, 2003), (Manduchi, Castano, Talukder, & Matthies, 2005).

It is assumed that the vehicle is equipped with a forward-looking range sensor that can measure terrain elevation; an inertial navigation sensor that can measure the vehicle’s roll, pitch, yaw, roll rates, pitch rates, yaw rates, and translational accelerations with reasonable uncertainty; and a global positioning system that can measure the vehicle’s position and velocity in space with reasonable uncertainty.

3. TRAJECTORY SPACE DESCRIPTION

The hazard avoidance algorithm is based on the trajectory space, a two-dimensional space of a vehicle’s instantaneous path curvature, κ , and longitudinal velocity, v (Spenko, 2003]. Fig. 1 is an illustration of the trajectory space with icons depicting a vehicle’s actions corresponding to various points in the space. For this work, velocities are limited to positive.

The trajectory space is a convenient space for navigation because 1) constraints such as dynamic roll over, side slip, steering mechanism limits, over/understeer, and acceleration, braking, and steering rate limits can be imposed on the space to yield a compact representation of a vehicle’s performance limits over uneven terrain and 2) the trajectory space maps easily to the UGV actuation space (generally consisting of the throttle and steering angle). The following is a summary of the application of the trajectory space to UGV method. A more complete description can be found elsewhere (Spenko 2005).

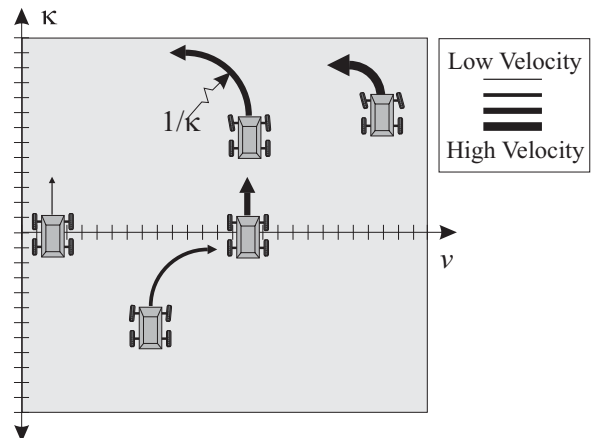


Fig. 1: Representation of vehicle action as described by its location in the trajectory space.




Fig. 2: Dynamic trajectory space limits for varying terrain roll angles (UGV wheelbase = 2.5 m).

3.1. The Dynamic Trajectory Space, A

The dynamic trajectory space consists of curvature and velocity pairs (v, κ) that do not cause excessive side slip or rollover and are attainable considering vehicle over/understeer effects.

The roll over limit is a function of vehicle tire and suspension characteristics, center of mass location, and terrain roll and pitch. The side slip limit is a function of the tire/terrain traction coefficient and terrain roll and pitch. The steering limits are a function of the maximum steering angle, center of mass location, and tire properties. Equations for deriving these constraints and techniques for modifying these constraints for rough terrain have been omitted (Spenko, 2005).

Fig. 2 illustrates the effect of terrain inclination on the dynamic trajectory space rollover limits. This example corresponds to a vehicle traversing a side slope with the fall line perpendicular to the vehicle's heading. As expected the vehicle can safely execute downhill turns (negative curvature) with greater velocity than it can execute uphill turns, since gravity counters the centripetal acceleration.

3.2. The Reachable Trajectory Space, B

The reachable trajectory space consists of velocity and curvature pairs that can be transitioned to in a given time. It is a function of the current UGV curvature and velocity as well as actuator, acceleration, braking, and steering characteristics. Fig. 3 shows a reachable trajectory space overlaid on the dynamic trajectory space for a HMMWV size vehicle with a current location in the trajectory space of $(v = 20.0, \kappa = 0.01)$. Steering rate limits are here fixed such that $\dot{\kappa}_{\max} = 0.05$.

3.3. The Admissible Trajectory Space, N

The admissible trajectory space (ATS) is the




Fig. 3: Reachable trajectory space

intersection of the dynamic trajectory space and the reachable trajectory space, $N = A \cap B$.

3.4. The Hazard Trajectory Space, H

Hazards can be generally classified as belonging to one of two types: trajectory independent hazards and trajectory dependent hazards. A trajectory independent hazard, such as a boulder or water trap, is one that a vehicle cannot safely negotiate independent of approach velocity and direction. For a trajectory dependent hazard, safe traversal depends on the vehicle approach velocity and/or direction. An example might be a shallow ditch where at high velocities a UGV can achieve ballistic motion and successfully “jump” the ditch.

The hazard trajectory space consists of curvatures and velocities that, if maintained from the current UGV position, would lead to intersection with a hazard (see Fig. 4). Here a point vehicle representation is employed. Note that there are no limitations as to the number of hazards that can appear in the trajectory space. The hazard trajectory space is generated by evaluating a pre-computed library of clothoidal paths that connect the current location in the trajectory space to other locations.




Fig. 4: Illustration of hazard trajectory space.




Fig. 5: ATSS defined with hazard present.

4. HIGH SPEED HAZARD AVOIDANCE

During high-speed navigation, emergency situations are likely to occur that require a UGV to rapidly perform a hazard avoidance maneuver. The two fundamental issues discussed below are 1) hazard detection, and 2) hazard avoidance maneuver selection.

4.1. Hazard Detection

A scenario similar to that illustrated in Fig. 5 is assumed. A UGV attempts to follow a pre-planned nominal trajectory given by a high-level path planner, $\tau_{nominal} \equiv (v(\mathbf{x}), \kappa(\mathbf{x}))$, where \mathbf{x} designates the UGV position in space. If a hazard detected by a range sensor poses a threat, the UGV enacts an emergency hazard avoidance maneuver. The sensor scan is divided into n discrete vehicle-sized patches and an ATS corresponding to each patch is computed. The size and number of these patches, sensor accuracy, and throughput are important issues, but are beyond the scope of this paper.

Let N_i denote the ATS for a patch that $\tau_{nominal}$ intersects. Let N_{raj} be defined as the intersection of all N_i , i.e. $N_{raj} \equiv N_1 \cap \dots \cap N_m$, where m is the number of patches that $\tau_{nominal}$ intersects. A maneuver is enacted when a hazard lies on the vehicle's current desired path or when a part of $\tau_{nominal}$ will violate a constraint on N_{raj} (i.e. a UGV is commanded to follow a dynamically inadmissible trajectory for a given terrain). Since the trajectory space gives a snapshot of the UGV's safe configurations for given terrain properties, as the terrain profile or composition change, the trajectory space limits also change.

4.2. Hazard Avoidance Maneuver Selection

To determine which maneuver to enact, let the total admissible trajectory space be defined as the intersection of all ATSS in the sensor scan minus the hazard space, H :

$$N_{total} \equiv (N_1 \cap \dots \cap N_n) - H \quad (1)$$

Let τ describe the UGV velocity and curvature at the current position \mathbf{x} . The goal of hazard avoidance is to find $\tau^*(\mathbf{x}) | \tau^*(\mathbf{x}) \in N_{total}$ where τ^* represents the hazard




Fig. 6: Curvature diagram

avoidance maneuver. The maneuver thus transitions the vehicle from a location that violates an ATS constraint to one that does not. There are numerous techniques for finding a τ^* that results in a “good” maneuver. The following method was adopted for its simplicity.

First, the trajectory space is discretized into i closely spaced grid points. τ^* is chosen as the location in the trajectory space that minimizes the distance, Δ , from the current location in the trajectory space, $\tau = (v_0, \kappa_0)$, to a candidate point:

$$\Delta = \sqrt{\frac{K_1}{\kappa_{max} - \kappa_{min}} (\kappa_0 - \kappa_i)^2 + \frac{K_2}{v_{max}} (v_0 - v_i)^2} \quad (2)$$

where K_1 and K_2 are static non-negative gain factors. These factors affect the relative weighting of changes in velocity and curvature. The minimum distance Δ over N_{total} can be found using a variety of search techniques.

The resulting τ^* represents a dynamically admissible curvature and velocity pair that avoids hazards in the current sensor scan. A low-level control algorithm is then employed to command the UGV along the new trajectory.

5. PATH RESUMPTION

After a hazard avoidance maneuver is executed, the UGV must plan a kinematically and dynamically feasible path to return to the pre-planned nominal path. Assuming constant velocity, v , the state of a front-steered rear-drive wheeled vehicle can be described by the following coupled nonlinear equations.

$$\begin{aligned} \kappa(s) &= u(s) & \theta(s) &= v \int_0^L \kappa(s) ds \\ x(s) &= v \int_0^L \cos \theta(s) ds & y(s) &= v \int_0^L \sin \theta(s) ds \end{aligned} \quad (3)$$

where s is the vehicle distance along a path, $u(s)$ is the steering input, and $\theta(s)$ is the vehicle heading angle.

Consider the situation illustrated by the plot shown in Fig. 6. The solid line represents a pre-planned nominal maneuver's curvature in path coordinates. A hazard avoidance maneuver is executed at a , and the maneuver

ends at b . The curvature of the nominal desired path, hazard avoidance maneuver, and path resumption maneuver are defined as $\kappa_1(s)$, $\kappa_2(s)$, and $\kappa_3(s)$ respectively. The goal of the path replanning problem is to find $\kappa_3(s)$ in a computationally efficient manner such that:

$$(\kappa(c), \theta(c), x(c), y(c))_1 = (\kappa(d), \theta(d), x(d), y(d))_3 \quad (4)$$

where c is the desired “meeting point” of the replanning maneuver and the nominal trajectory, and d is the terminal point of the replanning maneuver.

A computationally efficient replanning method termed the ‘curvature matching method’ is presented here. Comparisons of this approach with others can be found elsewhere (Spenko, 2005). An outline of the method is presented below:

1. Make an initial choice of the “meeting point” on the nominal trajectory. Here c is initially chosen such that $(c-b) = (b-a)$. The initial value of d is chosen to be the smallest value such that it is possible to transition from $\kappa_2(b)$ to $\kappa_3(d)$ without violating $|\dot{\kappa}| \leq \dot{\kappa}_{\max}$.
2. Find $\kappa_3(s)$ such that:

$$\int_a^c \kappa_1(s) ds = \int_a^b \kappa_2(s) ds + \int_b^d \kappa_3(s) ds \quad (5)$$

This ensures that $\theta_1(c) = \theta_3(d)$. The curvature, κ_3 , must also stay within the boundaries of the total admissible trajectory space. Details are given in (Spenko, 2005).

3. Calculate $x_3(d)$ and $y_3(d)$ using (3).
4. The algorithm ends if $x_3(d)$ and $y_3(d)$ are within the acceptable threshold. If not, c and d are adjusted as:

$$\begin{aligned} c_{i+1} &= c_i - k_c(e_{lon}) \\ d_{i+1} &= d_i - k_d(e_{lat}) \end{aligned} \quad (6)$$

where k_c and k_d are adjustable gains and e_{lon} and e_{lat} are the longitudinal and lateral error respectively and are defined by:

$$\begin{aligned} e_{lon} &= (x_1(c) - x_3(d))\cos\theta_1(c) + (y_1(c) - y_3(d))\sin\theta_1(c) \\ e_{lat} &= (x_3(d) - x_1(c))\sin\theta_1(c) + (y_1(c) - y_3(d))\cos\theta_1(c) \end{aligned} \quad (7)$$

Due to the fact that the equations of motion are coupled and nonlinear (see Equation 3) algorithm convergence cannot be guaranteed. However, the convergence properties have been studied numerically and have yielded excellent results (Spenko, 2005). A ten thousand trial simulation using a PIII 1.5 GHz computer showed the curvature matching method generating a path with a median time of 10 ms and a mean time of 44 ms, which indicate the algorithm is sufficiently fast for use in high-speed situations.

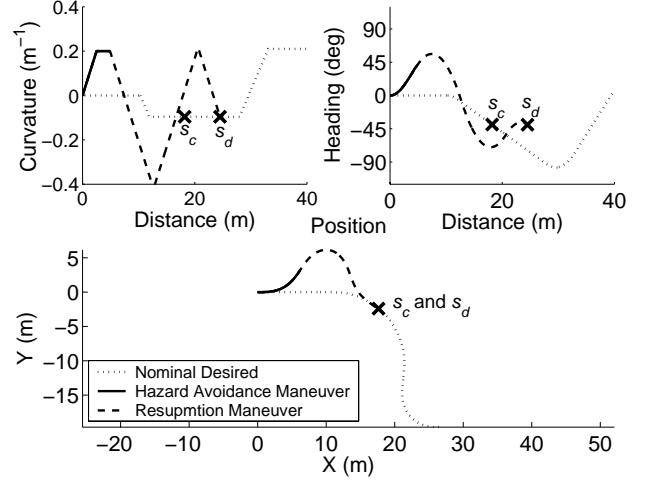


Fig. 8: Example of curvature matching method

Fig. 7 shows an example path resumption maneuver generated using the curvature matching method. Note that the nominal path’s curvature and heading and the path resumption curvature and heading profiles are identical at points s_c and s_d (upper left and upper right subplots), and points s_c and s_d are coincident along the path (lower subplot).

6. EXPERIMENTAL RESULTS

Experimental trials were conducted on the Autonomous Rough Terrain Experimental System (ARTEmiS); see Figure 8. ARTEmiS is a front-steer rear-wheel drive UGV that measures 0.88 m long, 0.61 m wide, and 0.38 m high. It has a 0.56 m wheelbase and 0.25 m diameter pneumatic tires. It is equipped with a 2.5 Hp Zenoah G2D70 gasoline engine, Crossbow AHRS-400 inertial navigation system (INS), Novatel differential global positioning system (DGPS) capable of 0.2 meter circular error probable resolution, Futaba S5050 servos for steering, brakes, and throttle, and a PIII 700 MHz PC104 computer. ARTEmiS is not equipped with forward-looking range sensors. Instead, using knowledge of ARTEmiS’ position, hazard locations are only revealed once they are within the range of a “virtual sensor.” Simulation results were obtained using MSC.ADAMS/Car software.

The steering angle and throttle were controlled using proportional-derivative control. For the steering angle,



Fig. 7: ARTEmiS experimental UGV

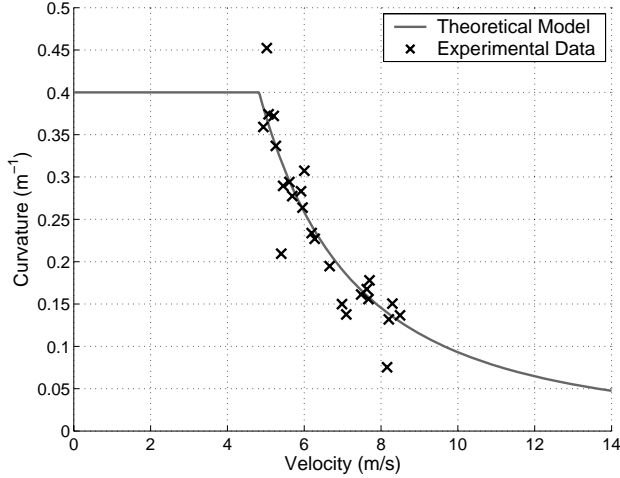


Fig. 9: Experimentally validated trajectory space constraints for flat ground

the gain was inversely proportional to the vehicle longitudinal velocity. Sufficient path following results were obtained using previously developed algorithms (Canudas de Wit, Siciliano, & Bastin, 1996).

Because ARTEmiS exhibits only slight oversteer, for the purpose of the experiments presented here the steering constraints were considered to be derived from a neutral-steered vehicle. Also, the center of mass of ARTEmiS does not bisect the track width of the vehicle, and thus the rollover constraints are not symmetric about zero curvature.

6.1. Experimental Validation of Trajectory Space Constraints

The accuracy of the model-derived trajectory space roll over constraints was studied experimentally on flat terrain at speeds up to 8 m/s. The vehicle was commanded on a desired path consisting of a straight line followed by a clothoid segment. Roll over was defined as occurring when $a_y \geq gh/d$, where a_y is the lateral acceleration of the vehicle, g is gravity, h is the height of the vehicle center of mass and d is one-half the axle width. This simple metric is commonly used for rollover studies in the passenger vehicle industry. Due to the high traction coefficient ($\mu \approx 1.3$), roll over occurred before excessive side slip.

The experimental data matches the predicted dynamic limit well. The most prevalent source of error is the calculation of the path curvature, which can be highly sensitive to the GPS and INS position estimates.

6.2. Validation of Hazard Avoidance Maneuver Algorithm

The hazard avoidance maneuver algorithm was

validated through both simulation and experimental analysis. Over 80 hours of experimental data was collected on a variety of terrain surfaces, profiles, and conditions, at speeds ranging from 3.0-9.0 m/s. This section provides results from five experiments.

ARTEmiS was placed in an initial starting location, (x_0, y_0) , and commanded to follow a nominal desired trajectory, $\tau_{nominal}$, with a corresponding path, $\mathbf{x}_{nominal}$. Hazards were represented by traffic cones placed in various configurations. The range of the sensor varied among experiments from 12 m to 20 m (21 to 35 times the vehicle wheelbase). Once a hazard was in range it was assumed that the hazard geometry was known. All experiments used the curvature matching method to generate a path resumption maneuver. All experiments also used the maneuver selection cost function given in Equation 2 with $K_1 < K_2$ unless otherwise noted.

Note that other experiments were conducted to investigate the effects of a reduced sensor range on resulting hazard avoidance maneuvers, but due to length constraints the results are not included here. As expected as the sensor range is reduced, the resulting hazard avoidance maneuvers are usually more severe (sharper turning and harder braking) than similar experiments conducted with longer-range sensors.

6.3. Multiple Hazard Experimental Results

Results from two experimental trials are presented that illustrate the ability of the algorithm to avoid multiple hazards.

Figure 10(left) shows three “snapshot” subplots of the GPS trace from an experiment for high speed avoidance of two hazards. The experiment was performed on a field of mixed grass and dirt, at a desired velocity of 6.0 m/s. The nominal desired path was a 100 m long straight path. ARTEmiS detected the first hazard at $x = 16.4$ m. This is shown in the top subplot of Figure 10. At this point a hazard avoidance maneuver was executed. ARTEmiS followed the modified path until a second hazard was detected at $x = 43.2$ m. This is shown in the middle subplot of Figure 10. A second maneuver was then executed and ARTEmiS successfully resumed the nominal path, as shown in the lower section of Figure 10.

Figure 10(right) shows the trajectory spaces at the instant that the first hazard was detected. An \mathbf{x} marks ARTEmiS’ location in the trajectory space. Here, ARTEmiS modified its trajectory from $\tau_0 = (6.0, 0.00)$ to $\tau_f = (6.0, -0.03)$, i.e. it executed a sharp turn to avoid the hazard.

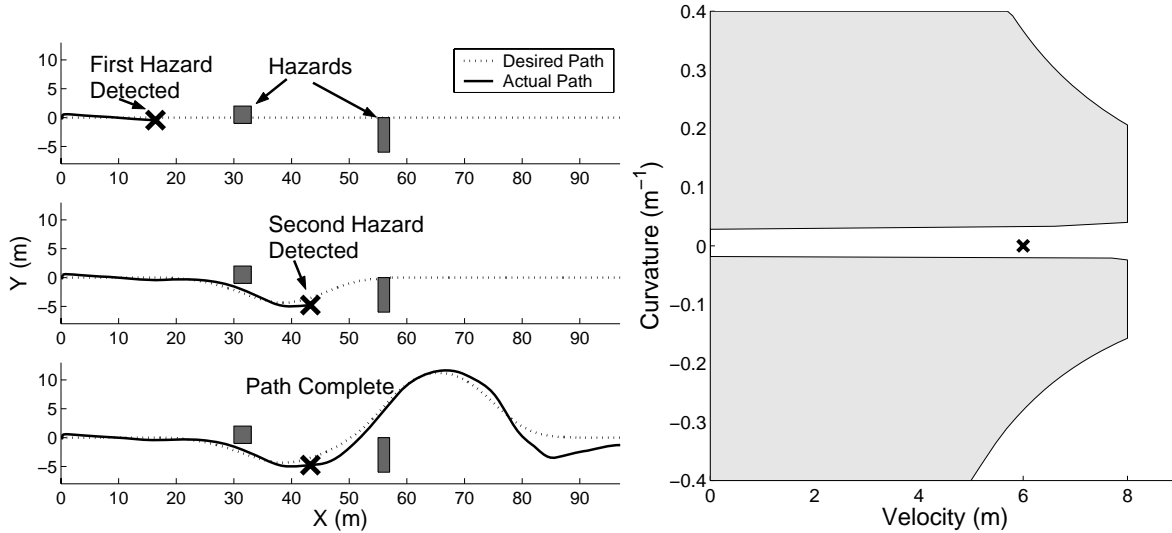


Figure 10. Hazard avoidance maneuvers executed for multiple hazards (left), and corresponding trajectory space (right).

Path tracking errors in the experimental system were due to position estimation errors and mechanical limitations of ARTEmiS’s steering mechanism, which are backdrivable and slightly underpowered. Thus terrain roughness caused substantial disturbances to the steering system.

6.4. Rough Terrain Experimental Results

Experiments on rough terrain were performed at Minute Man National Historic Park. The terrain consisted of a bumpy, uncut grass field. Physical terrain features tended to be on the order of one-half the wheel radius. Figure 15 illustrates the roughness of the terrain by comparing experimentally-measured UGV vertical acceleration measured on both smooth and rough terrain at the experiment site. Data was gathered while ARTEmiS traveled at 7 m/s.

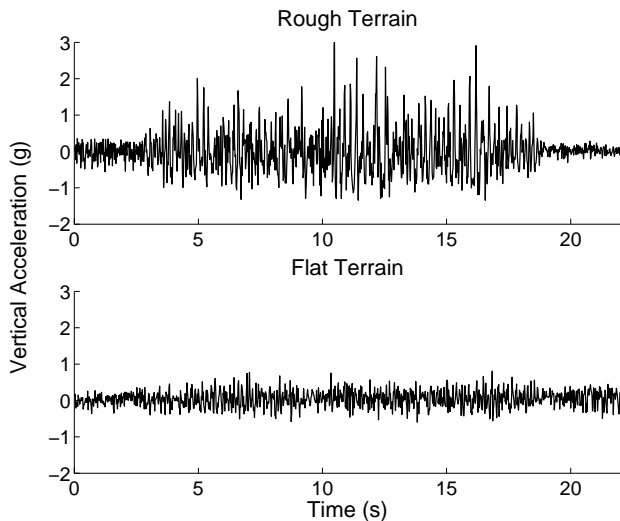


Figure 11. Rough and flat terrain vertical acceleration comparison.

Figure 16 shows the experimental site. The nominal desired path is a 100 m long straight path. ARTEmiS is pictured at the start of the path. The goal location is obstructed from view by the hazard. The hazard consists of a cluster of tall brushes, and small trees.

Figure 17 shows three “snapshot” subplots of the experiment. The experiment was performed at a speed of 7.0 m/s. ARTEmiS detected the first hazard at $x = 10.4$ m. This is shown in the top subplot of Figure 17. At this point hazard avoidance and path resumption maneuvers were executed, as shown in the middle subplot of Figure 17. The lower section of Figure 17 shows the completed path.

Figure 18 shows the trajectory space at the time the hazard was detected. The dynamic rollover limits included an empirically-determined “safety margin” to compensate for the effects of terrain roughness. When the hazard was detected, ARTEmiS modified its trajectory from $\tau_0 = (7.0, 0.00)$ to $\tau_f = (7.0, 0.03)$.




Figure 12. Rough terrain experimental setup.

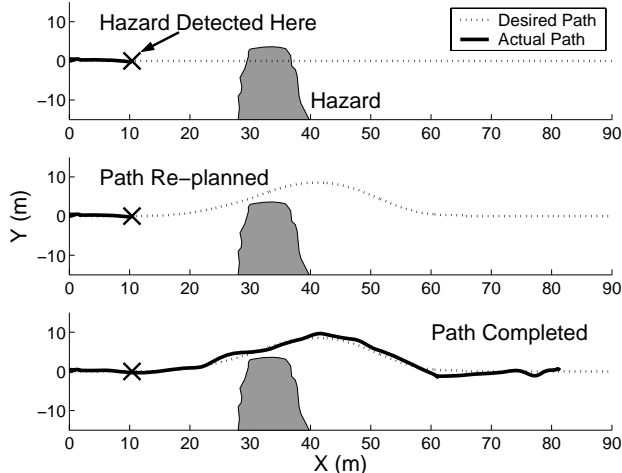


Figure 14. Rough terrain experimental results.

This experiment demonstrates that the proposed hazard avoidance algorithm can be applied in to UGVs operating at high speeds on rough terrain. These conditions are expected to be similar to actual operating conditions for many practical applications.

7. CONCLUSIONS

This paper has presented an algorithm for hazard avoidance for high-speed unmanned ground vehicles operating on rough, natural terrain. The algorithm accounts for dynamic effects such as vehicle sideslip, rollover, and over/understeer, as well as vehicle steering dynamics, drive train properties, terrain geometry, and vehicle/terrain interaction. The method is computationally efficient (operating on the order of milliseconds), and thus suitable for on-board implementation. Extensive simulation and experimental results have been presented that demonstrate the algorithm's effectiveness. The hazard avoidance algorithm based on the trajectory space is only one of many that could be implemented, and future work focuses on expanding this area.

ACKNOWLEDGEMENTS

This research was supported by the U.S. Army Tank-automotive and Armaments Command and the Defense Advanced Research Projects Agency. Thank you to Shingo Shimoda, Guillaume Morel, and Dariusz Golda for their help with ARTEmiS.

REFERENCES

- Coombs, D., et al., 2000: Driving autonomously off-road up to 35 km/h. *Proc. of the IEEE Intelligent Vehicle Symposium*, 186-191.
- Daily, M. et al., 1988: Autonomous cross-country navigation with the ALV. *Proceedings of the IEEE ICRA*, **2**, 718-726.
- Dudgeon, J. and R. Gopalakrishnan, 1996: Fractal-based

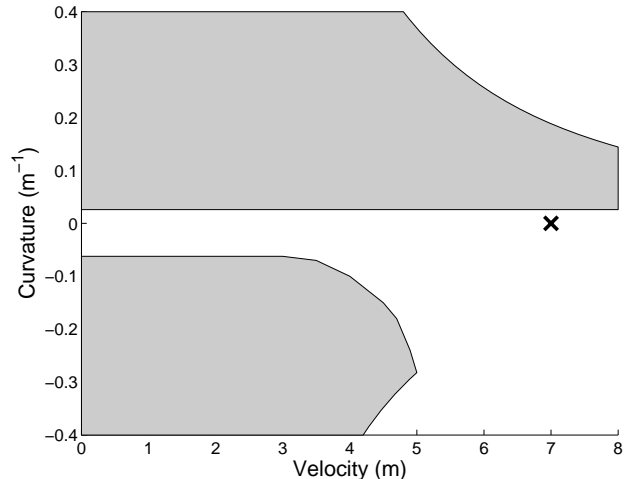


Figure 13. Rough terrain trajectory space.

modeling of 3D terrain surfaces. *Proc. of the IEEE Conference on Bring Together Education, Science, and Technology*, 246-252.

- Fox, D., W. Burgard, and S. Thrun, 1997: The dynamic window approach to collision avoidance. *IEEE Robotics and Automation Magazine*, **23**, 23-33.
- Iagnemma, K., S. Kang, C. Brooks, and S. Dubowsky, 2003: Multi-sensor terrain estimation for planetary rovers. *Proc. of the 7th, i-SAIRAS*.
- Iagnemma, K., and S. Dubowsky, 2004: *Mobile Robots in Rough Terrain, STAR Series on Advanced Robotics*, Springer.
- Kelly, A., and A. Stentz, 1998: Rough terrain autonomous mobility – part 1: a theoretical analysis of requirements. *Autonomous Robots* **5** 129-161.
- Langer, D., J.K. Rosenblatt, and M. Hebert, 1994: A behavior-based system for off-road navigation. *IEEE Transactions on Robotic and Automation*, **10.6**, 776-783.
- Laugier, C. et al., 1998: Sensor-based control architecture for a car-like vehicle. *Proc. of the IEEE ICRA* 216-222.
- Olin, K. and D. Tseng, 1991: Autonomous cross-country navigation: an integrated perception and planning system.” *IEEE Expert*, **6.4**, 16-30.
- Spenko, M., K. Iagnemma, and S. Dubowsky, 2004: High speed hazard avoidance for mobile robots in rough terrain. *Proc. of the SPIE Conf. on Unmanned Ground Vehicle Technology VI*, **5422**
- Spenko, M., 2005: Hazard Avoidance for High-Speed Rough-Terrain Unmanned Ground Vehicles *Ph.D. Thesis. Massachusetts Institute of Technology*.