

Artificial Potential Functions for Highway Driving with Collision Avoidance

Michael T. Wolf and Joel W. Burdick
Mechanical Engineering
California Institute of Technology
wolf@caltech.edu

Abstract— We present a set of potential function components to assist an automated or semi-automated vehicle in navigating a multi-lane, populated highway. The resulting potential field is constructed as a superposition of disparate functions for lane-keeping, road-staying, speed preference, and vehicle avoidance and passing. The construction of the vehicle avoidance potential is of primary importance, incorporating the structure and protocol of laned highway driving. Particularly, the shape and dimensions of the potential field *behind* each obstacle vehicle can appropriately encourage control vehicle slowing and/or passing, depending on the cars' velocities and surrounding traffic. Hard barriers on roadway edges and soft boundaries between navigable lanes keep the vehicle on the highway, with a preference to travel in a lane center.

I. INTRODUCTION

The structure and relative simplicity of automobile highway driving offer an opportunity for greater automation and driver assistance, hopefully increasing the safety and efficiency of these roads. Control frameworks may be applied to highway travel without the complications of street driving (stop lights, stop signs, cross traffic, frequent turns, etc.), and the higher vehicle speeds mean that avoiding collisions is of even greater importance, being more likely to save lives. Additionally, many hypothesize that adding automation on highways will increase throughput, with coordinated traffic moving swiftly even in crowded cases.

However, driving is a complex task, and even highway subtleties make the implementation of a rule-based navigation algorithm cumbersome. Therefore, utilizing an artificial potential field for local trajectory generation offers an elegant alternative. Potential function approaches have already seen considerable use for robot motion planning [1], [2]. This method models the robot as a positively-charged body moving in an electric field. Typically, the potential function $U(q)$ includes a negative (attractive) component driving the robot towards its goal location and positive (repulsive) components around all obstacles in the workspace. Then, the desired motion from any given pose q can be determined by gradient descent, $\dot{q} = -\alpha \nabla U(q)$, or with dynamics,

$$m\ddot{q} + c\dot{q} = F = -\nabla U(q) .$$

The key to employing this technique is creating an appropriate potential function $U(q)$ to generate the desired behavior. Although collision-free robotic navigation has been well studied, distinctive issues and asymmetries apply to highway vehicles to differentiate our problem from that of a

classic omnidirectional robot in an “unbiased” configuration space:

- Vehicles have a preference to stay in well-defined lanes and further to stay in the middle of a lane.
- Direction of travel is limited; that is, vehicles must travel in the direction of the highway.
- Lateral velocity (due to lane change) is generally small relative to longitudinal velocity (along the highway). Accordingly, acceptable tolerances for obstacle proximity vary by direction; that is, being one meter to the side of another car is more acceptable than following one meter behind.
- The dynamic nature of the controlled vehicle and the obstacle vehicles plays a role in acceptable proximity; both the absolute and relative velocities determine desired distance to maintain behind another car, and these parameters clearly change with time.

Thus, we present a potential function for highway driving that encodes avoiding obstacles (primarily other vehicles) and road edges while preferring travel in a lane center and at a desired speed. This potential function may be envisioned as part of an overall control system for autonomous vehicle navigation in the traditional formulation; however, other applications, such as active driver assistance devices, may represent more realistic near term options.

Other researchers have also investigated utilizing potential fields for navigating automobiles. Gerdes and Rossetter have many contributions [3]–[6]; in particular, Rossetter's dissertation [5] focuses on incorporating a lane-keeping potential into a drive-by-wire car, including a detailed vehicle model and lateral stability requirements. Our primary contribution above their work and others [7]–[9] is an appropriate vehicle collision avoidance potential in a full two-dimensional field, including implicit decision-making on whether to pass slow obstacles. We also consider vehicle velocity dependencies and preferences in a more thorough manner.

II. POTENTIAL FUNCTION COMPONENTS

Let us first consider the behavior we wish our highway vehicle to exhibit. In the absence of obstacles, the control vehicle should maintain a target speed v_{des} (e.g. the speed limit) and stay in the center of its current lane. When approaching another vehicle in the same lane, the controlled vehicle should either (a) change lanes to avoid the other car, if the speed of this “obstacle car” is too slow and an adjacent lane is clear; or, (b) slow to match the other car's speed,

which is acceptable if it is near v_{des} or in heavy traffic. In no circumstance should the vehicle leave the road or contact another vehicle.

For simplicity, assume the highway is straight and of infinite length; for curved roads, the same technique may be applied to longitudinal / lateral highway coordinates. We also assume knowledge of the roadway layout, our own vehicle state, and the positions and velocities of nearby vehicles. Other obstacles may be incorporated into this formulation but here we consider only other vehicles, as this is most common.

It is convenient to choose coordinate frames in which to describe the potential functions. Let the x -direction align with the direction of highway travel and choose the global reference frame origin such that the first (right-most) lane is centered at $y = 0$, as shown in Figure 1. The number of lanes and width of each lane are parameterized by N_{lanes} and δ_{lane} , respectively. All vehicles are modeled as rectangles of width (y -size) W and length (x -size) L , with a body-fixed frame attached in the rear middle (rear bumper). As in the usual formulation, we can consider our controlled vehicle as a point object in configuration space.

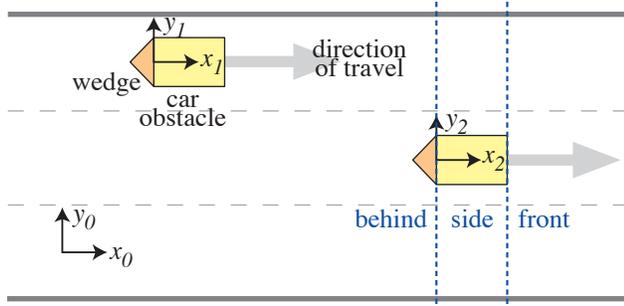


Fig. 1. Highway Setup: coordinate frames and obstacle labels

The total proposed potential U results from the superposition of several potential functions, which each fulfill a particular role in the highway navigation. We label these component functions *lane* for lane-keeping and centering, *road* for prevention of leaving the roadway, *car* for collision avoidance, and *vel* for targeting the desired speed v_{des} :

$$U = U_{lane} + U_{road} + U_{car} + U_{vel} .$$

U_{lane} , U_{road} , and U_{car} are each summations of similar components for each lane division, roadway edge (two), and nearby car, respectively. For example,

$$U_{lane} = \sum_i^{N_{lanes}-1} U_{lane,i} .$$

We address each component type individually below and later discuss their interplay.

A. Lane Potential

The lane potential U_{lane} creates a barrier to lane changing and guides the car into the center of its lane, but also is small enough to be easily overcome in the case that a lane change is necessary (or preferred) for collision avoidance. A Gaussian-like function is used to model the desired behavior between each lane:

$$U_{lane,i} = A_{lane} \exp\left(-\frac{(y - y_{c,i})^2}{2\sigma^2}\right) .$$

A_{lane} determines the maximum amplitude of the lane divider potential, $y_{c,i}$ is the y -position of the i^{th} lane division, and σ determines how quickly the potential rises/falls and is proportional to the lane width. Figure 2(A) shows a plot of this potential for a y -axis cross-section of a three-lane highway.

B. Road Potential

The road potential prevents the vehicle from leaving the highway by becoming infinite at the road edges (positive and negative y -boundaries of the workspace). On each side of the road, we employ a repulsive potential commonly used in robotics motion planning:

$$U_{road,j} = \frac{1}{2}\eta \left(\frac{1}{y - y_{0,j}}\right)^2 ,$$

where η is a scaling factor and $y_{0,j}$ is the j^{th} road edge coordinate, $j \in \{1, 2\}$. See Figure 2(B) for a plot and (C) for the roadway cross-section with the cumulative effect of both lane and road potentials.

C. Car Potential

The car potential aims to keep the vehicle a safe distance from each obstacle car by building a potential that rises to infinite strength approaching any part of the vehicle. Also, the potential must be shaped such that a lane change is encouraged if the controlled vehicle is rapidly approaching another car from behind — the vehicle should be forced to the side (positive or negative y direction) rather than to a stop (negative x direction).

The foundation of the car potential is the Yukawa potential [10]:

$$U_{car,m}(K) = A_{car} \frac{e^{-\alpha K}}{K} ,$$

where K is a distance measurement to the m^{th} obstacle car. A key exploration of this paper is defining this potential, particularly with regards to an appropriate pseudo-distance K , as it play a critical role in governing the key behaviors of our vehicle. A FIRAS potential [1] was also considered, but the Yukawa proved to be better suited to our application based on a slower rate of rise at moderate distances. This property helps govern the transition between maintaining a safe distance (pushing back) when farther away versus forcing a lane change (pushing to the side) when closer to the obstacle.

As noted earlier, the variable K acts as a pseudo-distance from the obstacle. We have taken some liberties with this

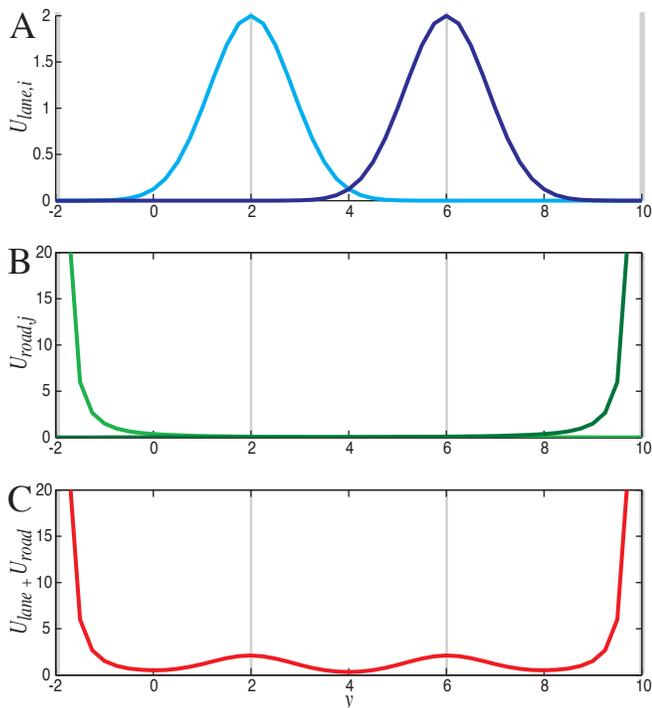


Fig. 2. Lane and road potentials for a cross-section of a three lane highway. (A) Two lane-dividing potentials. (B) Road edge potentials. (C) Superposition of all above components. Gray lines mark lane dividers and road edges.

measurement, artificially modifying it to account for the inherent asymmetries of highway driving. Consider the following reasoning. The level of danger associated with obstacle proximity depends also upon several other factors, governed both by the vehicle dynamics and rules of the road. For example, the relative position of cars on the highway plays a strong role; a vehicle should stay far behind a leading car in its lane, but is allowed relatively near the obstacle to the side. Thus, in terms of pseudo-distance, one may intuitively think of being 30 meters behind another vehicle “as close as” 2 meters to the side of that vehicle. By altering the pseudo-distance for this purpose, we can avoid the need to have a different potential function for the region behind an obstacle.

Another factor governing safe distances from obstacle vehicles is speed, particularly when following a vehicle. For high *relative* speed, the potential function influence should occur at larger distances, as (a) it will take longer to slow to match speeds or (b) there will be less distance over which to change lanes. Additionally, higher *absolute* velocity should also increase influence-distances, as, even when matching speeds with the leading vehicle, more space would be required for evasive maneuvering.

Accordingly, we describe below how K is determined for two regions: forward/side and behind the obstacle vehicle (see Figure 1).

1) *K for Forward and Side of Obstacle:* In front and directly to the side of the obstacle, K is the Euclidean

distance to the *nearest point* on the obstacle car:

$$K = \min_{b \in B} \|q - b\| ,$$

where B is the set of points comprising the obstacle. This approach results in lines of constant distance on immediate sides and front of the obstacle, with circular quarter-circle arcs connecting them (see Figure 3). The side-of-obstacle potential should have only a slight influence on adjacent vehicles, avoiding the initiation of a lane change into the obstacle. It is unlikely that the front-of-obstacle potential actually influences our vehicle given the responsibility of the trailing vehicle to avoid leading vehicles.

2) *K for Behind Obstacle:* The pseudo-distance K behind an obstacle is adjusted from the Euclidean distance in two important ways. First, a triangle, or wedge, is appended behind the vehicle to guide lane-changing behavior (Figure 1). The location of the rear-most vertex of this triangle is parameterized by Δ_t — moving it farther back gives a more pointed potential behind the car but also results in a larger area behind the car that is impossible to enter (Figure 3). This impossible-to-enter region acts as an additional safety zone behind an obstacle car. Note that far behind the vehicle, it is desirable that the potential field is fairly flat in the y -direction, getting more pointed as the vehicle (wedge) is approached.

The second modification to the pseudo-distance K in the region behind the vehicle is to shrink the effective distance as motivated above. Before determining the Euclidean distance to the closest point on the wedge, the x -position relative to the vehicle is scaled by a factor $\xi \in (0, 1]$, simulating that the current position is actually much closer to the wedge than it really is. As noted earlier, however, we wish ξ be dependent on the vehicle speeds. For the m^{th} obstacle car, we define

$$\xi_m(v) = \xi_0(v) \cdot e^{-\beta(v-v_m)} ,$$

where v is the controlled vehicle’s speed and v_m the speed of the obstacle car. The exponential factor therefore depends on the relative speed, scaled by parameter β . The quantity $\xi_0(v)$ then represents the scale factor for when the vehicles are traveling at the same speed and can depend on the vehicle’s absolute velocity. To set $\xi_0(v)$, we have derived inspiration from the “three second rule,” which dictates a following car should allow a distance to the leading car equal to the time traveled in three seconds:

$$\xi_0(v) = \begin{cases} \frac{d_0}{T_f v} & v \geq \frac{d_0}{T_f} \\ 1 & \text{otherwise} \end{cases} ,$$

where T_f is the desired following time (e.g., three seconds) and d_0 is essentially the maximum distance at which the potential U_{car} exerts an influence, as determined from the other system parameters. The case where $v < \frac{d_0}{T_f}$ should rarely be invoked since d_0 is small compared to vT_f but is included to ensure $\xi(v)$ does not exceed 1.

In summary, the procedure to calculate the potential (or gradient) when behind the k^{th} obstacle is: (1) scale down the x coordinate by $x' = \xi_m(v)x$ (in obstacle frame

coordinates); (2) compute K , the Euclidean distance from $q' = [x'y]^T$ to the nearest point on the wedge appended behind the obstacle; (3) use K to compute the potential (or gradient); and (4) for the gradient vector direction, scale back by $\xi_m(v)$.

Figure 3 displays the total resulting potential around a car. The scaling behind the car clearly expands its influence in this region. Additionally, the wedge appended behind the car artificially extends the obstacle and modifies its shape (compare with the front of the car) to both prevent the controlled vehicle from entering this zone and encourage lane change behavior when appropriate.

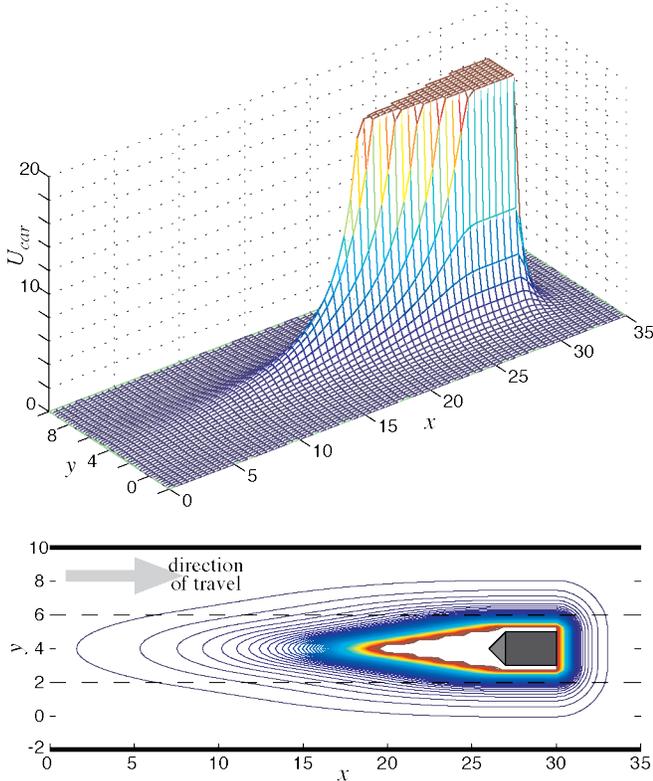


Fig. 3. Complete *car* potential, with isopotential contours. Large values are truncated.

D. Velocity Potential

The *velocity* potential function effectively creates a linear feedback law in the x direction in order to guide the vehicle towards a desired speed v_{des} :

$$U_{vel} = \gamma(v - v_{des})x .$$

This function results in a force propelling the vehicle forward when it is under its desired speed and a retarding force when the vehicle exceeds its desired speed.

III. SIMULATION RESULTS

Based on the potential function components above, we simulated several scenarios on a multi-lane, populated highway, employing the gradient-descent law described earlier. The scenario and parameter values used are shown in Table

I. To investigate the validity of the potential functions, we modeled the vehicle as a simple omnidirectional point forced by $-\nabla U$ with light damping. For more sophisticated applications of potential functions to car-like vehicles, in which our potential functions may also play a strong role, see [5].

TABLE I
SETUP AND FUNCTION PARAMETER VALUES

| | | | |
|----------------|--------------------|-----------------------|--------------------|
| Highway Setup | N_{lanes} | number of lanes | 3 |
| | δ_{lane} | lane width | 4 |
| | W | car width | 2 |
| | L | car length | 3 |
| Lane Potential | A_{lane} | max. height | 2 |
| | σ | width parameter | $0.3\delta_{lane}$ |
| Road Potential | η | scale | 3 |
| Car Potential | A_{car} | Yukawa amplitude | 10 |
| | α | Yukawa scale | 0.5 |
| | Δ_t | wedge vertex location | -0.5 |
| | β | ξ exp. scale | 0.6 |
| | T_f | ξ_0 follow time | 3 |
| | Velocity Potential | γ | slope scale |
| | v_{des} | desired speed | 25 |

We describe the resulting autonomous behavior in our simulations, presenting still images showing the potential fields from some of the more interesting cases. In these figures, a yellow square denotes the starting position and a green circle the end position. However, recall that these potential fields are dynamic, especially changing as a result of the controlled vehicle's velocity. For greater clarity on the resulting vehicle behavior, see the corresponding video of these simulations included in the proceedings.

Without obstacles present, we have tested many initial conditions to verify vehicle behavior to center in its lane, stay on the road, and travel at the desired speed. Perturbations from lane center and desired speed are quickly corrected, with perhaps some small oscillations before steady state, depending on damping values used in the simulation. The controlled car never leaves the roadway, successfully blocked by the *road* potential.

The addition of obstacles (other cars) in adjacent lanes effectively prevents lane changes in that direction without unduly forcing the controlled vehicle from its adjacent lane center. The most interesting test cases, however, consider how the controlled vehicle acts when approaching an obstacle in its lane — specifically whether the vehicle changes lanes.

Scenario 1: Figure 4 shows the total potential with one obstacle, meant to represent a leading car traveling nearly the desired velocity. In this case a controlled vehicle traveling at v_{des} in the middle lane slows into the local minimum behind the obstacle car. Note the obstacle car in the first lane is traveling faster and therefore has a “shorter” potential influence behind it.

Scenario 2: Figure 5 portrays a scenario in which the obstacle vehicle and controlled vehicle are going significantly slower. In this case, the speed of the leading vehicle is implicitly below threshold, under which it is then deemed an “obstacle.” (Consider a vehicle having mechanical difficulty or perhaps immediately following a collision; we show a

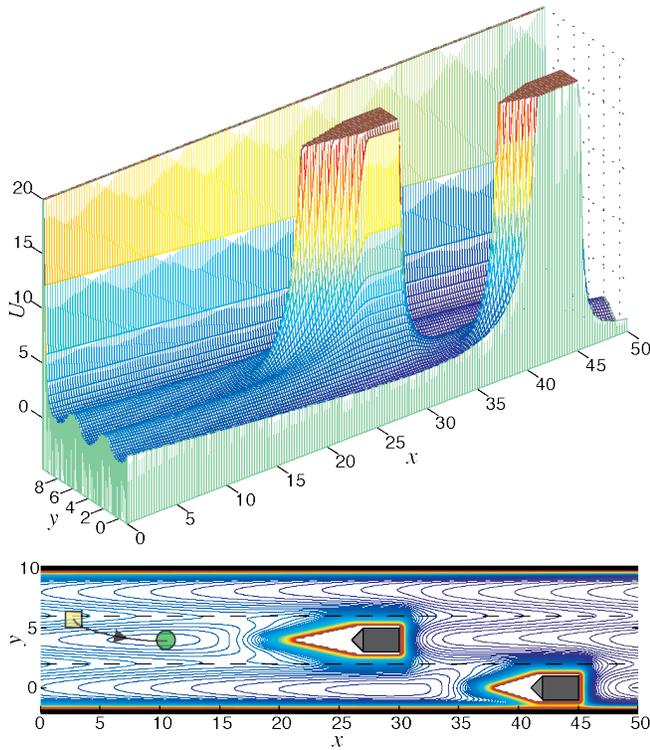


Fig. 4. Scenario 1: Local minima encourage vehicle to slow slightly and stay in lane.

moving vehicle to demonstrate behavior closer to the lane-change threshold and to compare better to Scenario 1.) Thus, the desired behavior is to change lanes. This does indeed happen, and one can see the local minimum behind the obstacle is not present as it was in the preceding scenario. The vehicle changes lanes to avoid the the obstacle and continue forward at v_{des} . Note that if an obstacle vehicle were in either adjacent lane, the vehicle would change lanes the opposite direction due to effects of the neighboring *car* potential.

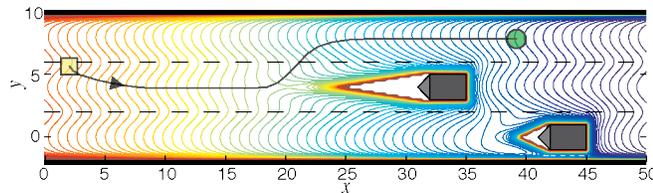


Fig. 5. Scenario 2: Vehicle changes lanes to avoid obstacle car and speed up.

Scenario 3: In Figure 6, multiple obstacle vehicles are present, representing a case with heavier traffic, with cars in both adjacent lanes and the leading obstacle car moving at a speed well below v_{des} . In this case, a local minimum exists once again, similar to the first scenario, so the vehicle will stay in its lane and slow to match the speed of the leading car.

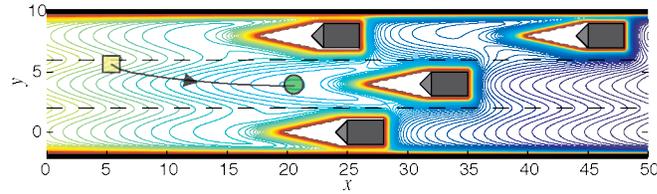


Fig. 6. Scenario 3: In heavier traffic, the vehicle slows and stays in its lane.

IV. DISCUSSION

Our simulation results display the vehicle behavior we designed the potential functions to elicit. With regards to lane changing, we are content to follow an obstacle provided it is moving reasonably close to our desired speed — we may consider the “effort” of changing lanes not worth it. However, if the leading vehicle’s speed is below tolerance, the automatic behavior is to change lanes, but only if adjacent lanes are clear of other obstacles.

The *lane*, *car*, and *velocity* potentials all work together to effect the above behavior. The question of whether a lane change will occur is a matter of considering if the lateral forces from the *car* potential’s wedge shape are large enough to overcome the barrier of the lane dividers. The slower the controlled vehicle is going, the greater the force resulting from *velocity* potential; in effect, this cancels out the gradual-sloped periphery of the *car* potential and thus the controlled vehicle engages the (stronger) region closer to the obstacle, where the lateral forces will be greater. Additionally, a larger difference in velocities between the two vehicles would increase the region of influence (decreases the pseudo-distance K) of the *car* potential, so the lane change would be initiated earlier.

This interplay, however, necessitates that the potential field parameter values must be chosen holistically. This fact, along with the sheer dimension of the parameter space, is a weakness (albeit surmountable) of our approach. One procedure for parameter selection is to select a few of the parameters as “independent,” with others assigned progressively. For example, starting with A_{lane} sets a scale for potential values. Then η , A_{car} , and α can be chosen based on the desired overlap of road edges and car sides into adjacent lanes. Next the behind-obstacle parameters Δ_t and β may be selected, and so on. Other approaches, such as automated computational parameter explorations, may be useful.

Unlike traditional potential field robotic navigation, there is no true goal location. Rather, any lane center is acceptable, with sufficient buffer from obstacles. Importantly, local minima are thus generally acceptable and in fact play an important role in the case of staying behind leading vehicles, especially in traffic. In this scenario we would not wish to navigate in between lanes to get to a lower potential, so staying in the resulting local minima is certainly preferable.

Known issues exist for attempting autonomous navigation via artificial potential fields. Although we have pointed out how local minima are actually helpful in some scenarios for

our application, they can also be a hindrance; for example, they prevent more complex maneuvers to get around another car. Additionally, care is required to ensure the commanded force is realistic, respecting the vehicle dynamics. Another consideration is time delay or time step size, as implementations occur in discrete time and large time steps can lead to unstable behavior and collisions.

Most likely, many limitations of autonomous potential field navigation will be overcome by using the proposed potential function as part of a hierarchical or hybrid control system. Such systems provide more convenient means of trajectory control, perhaps using a set of motion primitives that incorporate both car dynamics and highway protocol. Since decision logic in these systems is often difficult, a potential function provides a simple yet holistic view of environment to assist in evaluating the appropriate direction of motion. For example, perhaps instead of following the dynamics of a particle in the potential field to change lanes, a certain vector resulting from the potential field may trigger the stored procedure (motion primitive) of how to make the lane change, along with setting the necessary parameters.

To further refine the potential field presented in this paper, alternative or additive potential functions could also be explored in the future. The front-of-obstacle potential is one candidate for change. Although we do not wish for trailing cars to “push forward” our vehicle (and so we want to keep this influence small), expanding the influence of front-of-obstacle potential for cars traveling in *adjacent lanes* may be useful to prevent changing lanes immediately in front of a faster-moving car. For the velocity potential, other means may be explored for moving at desired velocities, as in slow conditions the resulting ramp can be quite strong — small perturbations in position or the obstacle layout can result in relatively large forces, creating a “tension” in which the vehicle is always trying to move at v_{des} (perhaps not unlike human drivers).

V. CONCLUSION

We have presented a set of potential function components useful for modeling highway driving. The *lane* and *road* potentials keep the vehicle on the roadway and preferably in the center of a lane and present a barrier to lane changes. The *car* potential results in appropriate collision avoidance behavior, with important velocity-dependent adjustments included for the region behind the obstacle. Finally, the *velocity* potential encourages the vehicle to move at a desired speed. Although we have demonstrated the usefulness of the combined potential function as the sole navigation means of an autonomous vehicle, several other uses may exist, including for decision making in a motion primitive based controller, or as additive input in driver assistance devices.

REFERENCES

- [1] O. Khatib, “Real-time obstacle avoidance for manipulators and mobile robots,” *The International Journal of Robotics Research*, vol. 5, no. 1, pp. 90–98, 1986.
- [2] H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, 2005.

- [3] J. C. Gerdes and E. J. Rossetter, “A unified approach to driver assistance systems based on artificial potential fields,” *Journal of Dynamic Systems Measurement and Control – Transactions of the ASME*, vol. 123, no. 3, pp. 431–438, Sep 2001.
- [4] J. C. Gerdes, E. J. Rossetter, and U. Saur, “Combining lanekeeping and vehicle following with hazard maps,” *Vehicle System Dynamics*, vol. 36, no. 4–5, pp. 391–411, 2001.
- [5] E. J. Rossetter, “A potential field framework for active vehicle lane-keeping assistance,” Ph.D. dissertation, Stanford University, 2003.
- [6] E. J. Rossetter, J. P. Switkes, and J. C. Gerdes, “Experimental validation of the potential field lanekeeping system,” *International Journal of Automotive Technology*, vol. 5, no. 2, pp. 95–108, 2004.
- [7] B. Schiller, V. Morellas, and M. Donath, “Collision avoidance for highway vehicles using the virtual bumper controller,” in *Proceedings of the IEEE International Symposium on Intelligent Vehicles*, 1998.
- [8] D. Reichardt and J. Shick, “Collision avoidance in dynamic environments applied to autonomous vehicle guidance on the motorway,” *Intelligent Vehicles '94 Symposium, Proceedings of the*, pp. 74–78, 1994.
- [9] M. Hassoun, Y. Demazeau, and C. Laugier, “Motion control for a car-like robot: potential field and multiagent approaches,” in *Proc. of the Int. Workshop on Intelligent Robots and Systems*, 1992.
- [10] R. Volpe and P. Khosla, “Manipulator control with superquadric artificial potential functions: theory and experiments,” *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 20, no. 6, pp. 1423–1436, 1990.