

Assisted Highway Lane Changing with RASCL

**Rich Frankel, Olafur Gudmundsson, Brett Miller, Jordan Potter, Todd Sullivan, Salik Syed
Doreen Hoang, Jae min John, Ki-Shui Liao, Pasha Nahass, Amanda Schwab, Jessica Yuan
David Stavens, Christian Plagemann, Clifford Nass, Sebastian Thrun**

{frankel, olafurg, brettm, jordan.potter, todd.sullivan, ssyed}@cs.stanford.edu
{doreenh, jaemin, kslliao, pnahass, aschwab, jmyuan}@stanford.edu
{dstavens, plagemann, nass, thrun}@stanford.edu

Abstract

Lane changing on highways is stressful. In this paper, we present RASCL, the Robotic Assistance System for Changing Lanes. RASCL combines state-of-the-art sensing and localization techniques with an accurate map describing road structure to detect and track other cars, determine whether or not a lane change to either side is safe, and communicate these safety statuses to the user using a variety of audio and visual interfaces. The user can interact with the system through specifying the size of their “comfort zone”, engaging the turn signal, or by simply driving across lane dividers. Additionally, RASCL provides speed change recommendations that are predicted to turn an unsafe lane change situation into a safe situation and enables communication with other vehicles by automatically controlling the turn signal when the driver attempts to change lanes without using the turn signal.

1. Introduction

Driving is one of the most ubiquitous aspects of modern life, yet it is also one of the most dangerous. According to the National Center for Health Statistics, motor vehicle traffic is the leading cause of unintentional injury and death among all Americans [1]. In 2007 alone, 41,059 people were killed in motor vehicle crashes [2]. One of the most dangerous and stressful parts of driving is lane changing. Indeed, 730,000 car crashes (1000 of them fatal) happened during “lane change, merging, and sideswipe” maneuvers [3].

An enduring goal of autonomous vehicle research is to reduce these numbers by removing the human element from driving entirely. The recent successes of the 2005 DARPA Grand Challenge and subsequent 2007 DARPA Urban Challenge suggest this future is closer than ever before, and yet there remains a divide. The transition to fully autonomous driving is unlikely to happen all at once for legal, technological, and psychological reasons. During the transition,

there will necessarily be interaction between an increasingly sophisticated and semi-autonomous car and the human driver.

Indeed, this transition is already underway. Many companies have begun implementing assistive driving systems. Currently on the market are a lane change warning system from Hella [4], a lane departure warning system from AutoVue [5], and a blind spot detection system from Mobileye [6], all discussed in greater detail in the next section.

The demands of fully autonomous driving have led to the development of even more sophisticated systems. Stanford’s entry in the Urban Challenge, Junior, is an autonomous driving platform of unprecedented sophistication [7]. It is equipped with both state-of-the-art sensors to perceive the world and cutting-edge artificial intelligence software to make sense of sensory inputs and navigate the world. These same features that made it such a success also allow us to explore new human-computer interfaces that take advantage of its underlying technology in either a semi-autonomous or fully-manual driving scenario.

In this paper, we describe RASCL, the Robotic Assistance System for Changing Lanes, an assistive lane changing system built on Junior’s platform. It is unique in a number of ways. It is capable of communicating both to the driver, through visual and auditory feedback, and to the drivers of other cars, through autonomous control of the turn signal. To do so, the system takes advantage of a wealth of sensory data from laser scanners, inertial sensors, and GPS, both for pose estimation and localization and for tracking surrounding vehicles. The system performs these tasks using a much more complete model of the world than existing systems, which we believe ultimately results in more robust recommendations.

To assess and evaluate the effectiveness of this novel technology, an exploratory user study was designed. The user study aimed to evaluate driver reactions to the implementation of the system in their vehicle in a real-world setting.

The remainder of this paper is as follows. After a description of the system implementation and its interface, we detail the user study design and present preliminary test results. We also discuss the technological issues that remain to be addressed, the implication of the test results, and open questions that warrant further investigation.

2. Related Work

Hella manufactures a lane change warning system that uses radar to determine when another car is next to the system's car. If the driver wants to change lanes and uses the turn signal to indicate this, the Hella system issues a dual warning causing lights in the side mirror to blink and the steering wheel to vibrate. Yet the Hella warning system is not without its flaws. Indeed, a severe one is that if the system is used on a non-divided road, such as a state or rural highway, and the opposing lane is clear, the system will actually inform the user that it is safe to merge into the lane of oncoming traffic [4].

A related technology is the lane change departure system manufactured by Iteris, called AutoVue. This system utilizes video cameras to follow lane markers and detect when the car begins to drift out of its lane. If the driver has not indicated a lane change by engaging the turn signal, the system issues a warning [5].

Mobileye manufactures systems that use side mirror-mounted cameras to detect vehicles in blind spots and warn the driver if those vehicles make a lane change unsafe. It reliably detects vehicles within 50m at highway speeds, and performs analysis to determine which lane the detected vehicle is in, a much needed feature for curved roads [6].

All of these systems suffer from limited knowledge of road structure. The Hella system has no notion of lanes whatsoever beyond an adjacent occupancy region it is checking, leading to poor recommendations under certain conditions. Both the AutoVue system and Mobileye system rely on computer vision-based lane marking detection, subject to poor lighting conditions and occlusion. We believe that global knowledge of road structure using a road map (the RNDF introduced in Section 3.1) combined with precision road-based localization, enable more useful and reliable lane change assistance to the driver and overcome some of the limitations of these existing systems. In addition, global

knowledge of road structure and position allows for more sophisticated behavior, including but not limited to the autonomous turn signal control we implement.

3. System

The RASCL system has three primary components: the car tracker, the Situation Model (SM), and the user interface. These three components form a pipeline where the car tracker provides a list of cars to the Situation Model, which provides safety statuses and speed change recommendations to the user interface, which communicates the information to the driver and other drivers in a variety of ways. In order to perform these tasks, our system relies on many libraries and other support modules for localization, navigating the local road structure, and perceiving obstacles.

3.1 RNDFs, Coordinate Systems, and Localization

To accurately determine the lane change safety at highway speeds our system relies on a detailed map of the highway and precise localization. Two coordinate systems are used to describe locations in the system, local coordinates that are relative to the car and Universal Transverse Mercator (UTM) coordinates, a type of global coordinates. The local coordinate frame has the x-axis pointing forward out of Junior's front windshield and y-axis to the left (out the driver's side door) whereas in UTM coordinates, the x-axis is East and y-axis is North. The origin of the local frame is in the center of Junior width-wise and near the back bumper length-wise. The exact origin is on the roof of the car, where all of the lines converge in Figure 3.4.4.3.

We use a Route Network Definition File (RNDF) as our knowledge base for where roads are located in the world and how they are connected. The RNDF format is a specific data format defined by the Defense Advanced Research Projects Agency (DARPA) for the Urban Grand Challenge containing geometric information on the highway [8]. The file consists of a set of waypoints in UTM coordinates aligned along the center of each lane in the highway. Each waypoint is assigned a unique ID that specifies the road segment and the lane it is on.

As described in [7], the inaccuracies in GPS position estimates from Junior's GPS receivers are supplemented by continuously estimating the local alignment between the RNDF and the current position using local sensor measurements. This allows for sufficiently accurate localization within specific lanes on the highway.

To facilitate tasks such as detecting the number of lanes to the right or left of Junior or estimating clearances in near-

by lanes, we created a library that combines the localization estimate and the information in the RNDF. Specifically, the library performs the following tasks:

- 1) Detecting whether a given point is close to the current highway. This is essential for filtering the sensor inputs and thus reducing computation.
- 2) Finding the lane heading at a given coordinate point.
- 3) Finding the closest waypoint within the RNDF file and the lane number of the closest lane.
- 4) Projecting a given point to the center of the closest lane or a certain number of lanes to the right or left relative to the current closest lane.
- 5) Projecting a given point a certain distance along the current highway lane.
- 6) Finding the distance between two points in the same lane.

Additionally, we created methods that allowed for easy conversion between the different coordinate systems. All these methods play a central role throughout the system.

3.2 Perception: Aggregating Readings from Multiple Sensors

The car tracking system makes use of sensory information from various laser scanners mounted on Junior. The primary lasers for perception are the Velodyne laser and the LD-LRS laser. As described in [7], the perceptual routines in Junior provide obstacle detection from these sensors. The output from these routines is the primary input to our car tracker. In order to reduce the number of obstacles sent to the car tracker, we prune the detected obstacles to only those that occur within 3 meters of a lane’s center using our RNDF library.

3.3 User Parameters

Both the determination of what constitutes a safe lane change and how that safety information is conveyed to the driver involve a number of variables for which the optimal value is unknown and possibly varies from one driver to the next. Specifically, the Situation Model in our system relies on the notion of clearance zones surrounding the vehicle. As mentioned previously, a lane change is considered safe if no vehicle will enter a zone for the duration of the change. To make this determination, the SM needs to know the dimensions of the clearance zones as well as the expected time to complete the lane change. Thus, we expose five user parameters:


- 1) Front-left clearance
- 2) Front-right clearance
- 3) Back-left clearance

- 4) Back-right clearance
- 5) Expected lane change time

The user can set each parameter to a value between zero and one, as described in Section 3.6.3. The actual clearance or time used is computed as $(\text{min} + \text{value} * (\text{max} - \text{min}))$ where min and max are the minimum and maximum values for the parameter and value is the user-chosen value between zero and one. The allowed ranges for the front clearances, back clearances, and lane change time are [8 meters, 38 meters], [4 meters, 34 meters], and [2 seconds, 11 seconds] respectively. The front clearances have larger mins/maxes because the origin of Junior’s local coordinate frame is located near the rear of the vehicle. Figure 3.3 shows the clearances when both front clearances are set to zero and both back clearances are set to 0.5, which corresponds to 8 meters of front clearance and 19 meters of back clearance.

Figure 3.3: Clearance Parameters to Actual Clearances

The front clearance parameters are set to zero while the back clearance parameters are set to 0.5. This results in 8 meters of front clearance and 19 meters of back clearance.



Additionally, there are a number of parameters that control the properties of the LED interface. We allow for adjustments to brightness, changes to the flash rate, and the enabling or disabling of flashing for either the green (safe) or red (unsafe) LEDs.

3.4 Car Tracking

Car tracking is essential to our highway lane changing system. Our car tracking module uses the obstacles generated from our perception module to detect potential cars in each frame. Each potential car is tracked over time using a linear Kalman filter, which gives a reliable estimate of the car’s position and velocity. Our car tracking algorithm has four primary stages: detecting potential cars in the current frame, projecting the previous frame’s cars forward in time, syncing the current frame’s cars with the previous frame’s projected cars, and recovering unsynced cars.

3.4.1 Detecting Cars in a Frame. The input to this phase is the set of obstacles from our perception module. Each obstacle has a width and length of 35 centimeters, and a height that is ignored. We treat each obstacle as a box on a 2D plane where the length is along the x-axis in our local coordinate system and the width is along the y-axis. Thus the boxes are axis-aligned and computing overlaps and distances between boxes becomes a simple computation.

We begin by merging boxes that overlap each other. When merging two boxes, the resulting box is the smallest box that bounds the two merged boxes. After merging overlapping boxes, we merge boxes that are within one meter of each other length-wise.


After this length-wise merge, we discard all boxes longer than 15 meters. This eliminates boxes that are from obstacles such as median dividers or bushes on the side of the road that were not pruned during our RNDP-pruning stage as described in Section 3.2.

We then merge boxes that are within half a meter width-wise and remove any resulting boxes that have a length and width of less than half a meter. The width-wise merging distance is smaller than the length-wise distance due to the general shape of cars and our desire to not merge two cars in adjacent lanes together. Removing small boxes with a width/length of less than half a meter eliminates unwanted noise in, for example, rainy environments.

Next, we resize any small boxes to meet our minimum car length and width, which are 4 meters and 2 meters respectively. If a box's width is less than 2 meters, then we increase the box's width going away from Junior. Thus if the box is on the left side of Junior (the box's center Y value in local coordinates is greater than zero), then we add the extra width to the left side of the box and leave the box's two corners closest to the car fixed. Similarly for length, we add extra length to the front of the box if the box is in front of Junior, and to the back of the box if the box is behind Junior. This process is shown in Figure 3.4.1a.

Figure 3.4.1a: Resizing Boxes

Perception obstacles are yellow. The white bounding boxes are the original boxes after merging the obstacles. The green rectangles are the extensions to the bounding box grown away from Junior.




After resizing the boxes, we again merge overlapping boxes. This generally consolidates any remaining pieces of an actual car into the same box, but the resize/merge step can also cause two actual cars that are close together width-wise to be merged into one box. To fix this issue, we then split boxes that are 4.1 meters wide or wider into as many 2 meter wide boxes as possible with 0.1 meters of space width-wise between boxes.

Now most boxes encompass all of the perception obstacles for a single actual car, but the boxes may be excessively long because of the resize/merge step. To reach our final box sizes, we group each perception obstacle with the box that encompasses it. Once each box has all of its enclosed perception obstacles, we resize the box to be the smallest box that bounds all of its perception obstacles and then again resize the box as shown in Figure 3.4.1a.

Figure 3.4.1b depicts a situation where this resize/merge - fit to perception obstacles - resize process yields ideal results. Step A shows the state of the algorithm after performing the initial overlap, length-wise, and width-wise box merges. Clearly the perception obstacles (in yellow) are part of one actual car. Step B shows the state of the two boxes after performing the initial box resizing from Figure 3.4.1a. The two resized boxes overlap, so we then merge them together to get the state in step C. Now all of the perception obstacles for the car are enclosed by a single box, but the box is excessively long. To remedy this, we fit the box to be the smallest box capable of encompassing all of its perception obstacles, which leads us to the state in step D. In this example the final box dimensions are larger than the minimums, so the remaining resizing step has no effect. By the end of this algorithm, the remaining boxes are deemed to be cars in the current frame.

Figure 3.4.1b: Overall Resizing Process



3.4.2 Projecting Cars Forward and Syncing Frames. We project each car in the previous frame forward in time using the car's Kalman filter's prediction function. After creating this projected previous frame, we sync the two frames by looking for overlapping cars. We will use "old car" to refer

to a car in the projected previous frame and “new car” to refer to a car in the current frame.

For a given new car, we find all old cars that overlap with the new car. The new car assumes the identity of the overlapping old car that has existed for the longest amount of time. For the case where the new car overlaps with more than one old car, this effectively corrects the error where the actual car in the previous frame was broken into two cars but is correctly merged into one car in the current frame. Similarly, if more than one new car overlaps with the same old car, it is assumed that the multiple new cars are pieces of the actual car that need to be merged together.

3.4.3 Recovering Unsynced Cars. After checking all overlaps, new cars may exist that are not matched to old cars and vice versa. The new cars are assumed to be cars that have never been seen before. The old cars are kept around until their confidence level drops below zero.

For new cars with no old car match, their confidence is set to zero. For each new car that has at least one old car match, we increase the car’s confidence by 0.33 (with a cap at 1.0). For old cars with no new car match, we decrease the car’s confidence by 0.1. The car tracking module outputs a car as a “real” car when its confidence is greater than or equal to 0.5. Thus new “real” cars appear after seeing them in 3 subsequent frames and cars with full confidence (1.0) disappear if they have not been seen for half a second.

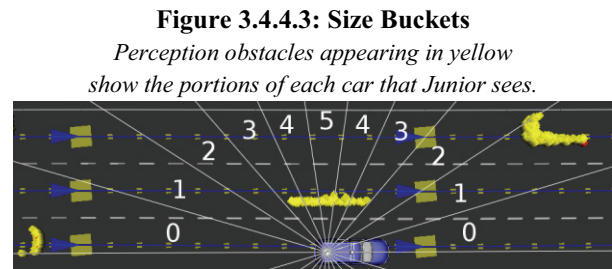
3.4.4 Car Spec Estimation. We have yet to describe in detail how our car module estimates position, velocity, direction, and size.

3.4.4.1 Estimating Position and Velocity. Our linear Kalman filter uses four state variables. We used $r^2=0.7$ and $q^2=0.1$ for the process and measurement noise respectively. As mentioned earlier, the Kalman filter provides reliable estimates of position and velocity from the noisy measurements the car’s position. Readers interested in Kalman filters should review [9]. For new cars that have never been seen before, we initialize the velocity to that of Junior.

3.4.4.2 Estimating Direction. For each car we maintain the estimated position of the car over the last 10 frames. With our 100ms frame rate, this equates to remembering the car’s position over the last one second. If the car is less than four frames old, we set the car’s direction to the heading for the lane that the car is in (which we know from the RPDF). Otherwise, we calculate the car’s direction by fitting a straight line to the history of positions.

The car’s direction can also be estimated using the Kalman filter’s output since we obtain x and y component velocities in the global UTM frame. Both methods yield similar results, with an average absolute difference between each method’s estimate of only 0.5 degrees and a rare maximum difference of 1.5 degrees.

3.4.4.3 Estimating Size. We estimate a car’s width and length by maintaining six exponentially decaying averages of the car’s dimensions with a decay rate of 0.985. We split the region around Junior into 22 sections, where each section spans approximately 16 degrees of the 360 degrees around Junior. Each section is marked as being a part of one of six buckets, as shown in Figure 3.4.4.3. The 180 degrees on the right side of Junior mirrors the left side, which is shown.



For a given car in the current frame, we place the car’s width and length estimate in the bucket that corresponds to the section that the car’s center is located in. Thus the car in the top right of Figure 3.4.4.3 would have the current frame’s width and length estimate added to bucket 1, while the car directly to the left of Junior would have the estimates added to bucket 5. Adding a width or length to a bucket involves adjusting the bucket’s numerator and denominator as follows:

$$\begin{aligned} \text{numerator} &= \text{numerator} \bullet 0.985 + \text{new_value} \\ \text{denominator} &= \text{denominator} \bullet 0.985 + 1.0 \end{aligned}$$

After adding the new measurement to the appropriate bucket, the car’s width and length are determined via a weighted average of the buckets. For width, as is evident in Figure 3.4.4.3, our estimates are more accurate when the car is closer towards the front or back of Junior. For length, our estimates are better when the car is closer to Junior’s side. To take advantage of this, for the width calculation we give the buckets 0, 1, 2, 3, 4, 5 weights of 6, 5, 4, 3, 2, 1 respectively. Length calculations use the opposite: for buckets 0, 1, 2, 3, 4, 5 we give weights of 1, 2, 3, 4, 5, 6.

After calculating the new width and length we adjust the car’s size using the same method as in Section 3.4.1. For the top-right car in Figure 3.4.4.3, this would involve leaving the corner made by the car’s right and rear walls fixed,

and growing the car's box away from the fixed corner to the appropriate size. For the car immediately to the left of Junior, the car's right wall would be fixed, its box would be grown to the left to the given width, its length would be set to the given length, and the center of the box length-wise would be centered on its old length-wise center. The car behind Junior would have its front wall fixed and its length grown away from Junior.

For new cars with no history we initialize the width of each bucket to 2 meters and the length of each bucket to 4 meters with the weights of these prior estimates set to 0.25.

3.5 Situation Model

The Situation Model receives information about tracked cars, combines it with information about the road (from the RNDF) and information about Junior's pose, and produces an estimate of whether or not it is safe to change lanes to either side. If performing a lane change is not safe, the system will give rudimentary speed change recommendations for making it possible to safely execute the desired lane change. In addition to the above, the SM also keeps track of Junior's forward clearance, so that the driver can be alerted if someone cuts him off.

3.5.1 Safety. To make a safety estimate, the SM plots the position of Junior over the lane change duration (using information about the shape of the road, from the RNDF) assuming the driver was to execute a lane change immediately. The module then plots the estimated positions of each tracked car, given its position and velocity. (In the standard mode, it is assumed that other cars will not change lanes. See Appendix A for information about experimental work on lane change prediction.) If any car enters Junior's safety zone as defined by the user parameters discussed in Section 3.3 during the duration of the lane change, the lane change is estimated to be unsafe. Otherwise, the lane change is considered safe.

3.5.2 Lane Adjacency. When estimating lane change safety, the simplest approach to determining which lanes to check is to check the lanes adjacent to the lane occupied by Junior's center. However, this becomes problematic when changing lanes. Roughly halfway through a lane change Junior's center will cross from one lane to the next, causing the new adjacent lanes to be monitored. If a car is in one of these new adjacent lanes (two lanes over from the lane Junior was originally in), the system would start alerting of an unsafe lane change because the driver is still in the process of mak-

ing the current lane change and most likely still has the turn signal on.

Think of a road with five lanes, labeled from left to right as L1 to L5. Imagine that Junior begins in the center lane, L3 (and, therefore, his left lane is L2 and his right lane is L4). If Junior makes a left lane change, at some point during the lane change, Junior's center passes from L3 to L2. At that moment, Junior's new left and right lanes would be L1 and L3, respectively.

Now, say that there is a car in L1 for the duration of the lane change, parallel to and keeping pace with Junior. When Junior's center passes into L2, the SM would begin checking L1 for lane changes; because L1 is occupied, the left lane change in progress would suddenly be deemed unsafe. Clearly, this is incorrect behavior – L1 should not be checked until the lane change has completed. To mitigate this issue, the SM defines Junior's left lane as the lane to the left of the lane occupied by Junior's right wheels, and Junior's right lane as the lane to the right of the lane occupied by Junior's left wheels. As shown in Figure 3.5.2, the effect of this is that when Junior is straddling two lanes, those two lanes are the lanes checked.


Often when Junior has fully changed lanes (both sets of wheels are now in the new lane), the driver still has the turn signal on. This is reasonable, because the driver is still not near the center of the new lane. If the SM were to begin checking the new adjacent lane immediately, the driver could be frightened. At the same time, the driver may intend to change two lanes at once.

To strike a compromise between these two use cases, after the final set of wheels crosses into the new lane we start a one second timer and turn off the system's recommendations on the appropriate side (the left side in our example) for the duration of the timer. Turning off the system's recommendations for a given side results in no sounds being played if the turn signal is set to that side and the given side's green and red LEDs turning off.

3.5.3 Speed Change Recommendations. Under certain conditions, if a lane change is unsafe the SM is able to make basic speed change recommendations that it estimates will create a safe situation for performing the lane change. This is done simply by extending the simulation, and delaying the safety checks, by the estimated time required for Junior to accelerate to the new desired speed.

Figure 3.5.2: Straddling Lanes

When straddling lanes, the two lanes that are monitored are the two lanes that Junior is simultaneously in.



3.6 User Interface

Our system incorporates a multitude of user interface mechanisms into a coherent user interface that communicates valuable information to the driver of Junior and to other drivers on the road. The parts of this user interface include LEDs for indicating the safety of changing lanes to either side, auditory signals for conveying the safety or danger level of making a specific lane change when using the turn signal, an LCD screen for adjusting user parameters and providing speed change recommendations when using the turn signal, and automatic turn signal control for when the driver of Junior forgets to use the turn signal during a lane change.

3.6.1 LEDs. Two sets of LEDs are attached to the dashboard to convey the relative safety of a lane change for each side of the car (Figure 3.6.1). Each set contains a red LED and a green LED. The red intuitively represents that it is not safe to change lanes on that side, and conversely, green represents that a lane change can be made safely. If it is neither safe nor unsafe (say, for example, there is no lane on that side or the car is not on a road in the RNDF), the LEDs for that side will remain off.

Figure 3.6.1: LED Interface



To further enhance the vehicle-driver interaction and to express more information to the user, an unsafe blinking algorithm is implemented. If it is unsafe to change lanes on a given side, the red LED not only turns on, but also blinks at

a rate that depends on how unsafe it is. The less safe it is to make the lane change, the faster the red LED will blink, expressing the danger level of changing lanes. The equation used to calculate the period between LED blinks is:

$$LEDPeriod = \max\left(0.05, 1 - \left(\frac{\max(safetyZoneSize - distanceToCar, 0.0)}{safetyZoneSize}\right)\right)$$

Thus, the closer a car is to the driver, the faster the red LED will blink to convey this. However, if there is no car in the driver's safety zone for a given side, then the green light will be turned on, indicating that it is safe to change lanes.

3.6.2 Sounds. The second method to convey the safety of a lane change is through sounds. Sounds are only used if the turn signal is engaged, and only on the side that the turn signal is engaged on. For example, if the driver turns on the left turn signal, sound will only be played on the left side of the car, and only evaluate the safety of the left side. There are three sounds incorporated into the system: a safe sound, warning sound, and unsafe-safe transition sound. The safe sound is the familiar sound of the normal turn signal.

The warning sound will only play if it is not safe to make the desired lane change, and will change repeat rate depending on the relative safety. For example, if it is unsafe to make a lane change, but the car posing the danger is at the edge of the safety zone, say 20 meters away, then the sound will play very slowly. However, if the car is right next to the driver, the sound will be repeated quickly.

The unsafe-safe transition sound will play if it was previously unsafe to change lanes but then became safe. If the driver is waiting for the status to become safe before executing the lane change, this sound will accurately tell the user that now the situation has changed, and to go ahead and make the lane change.

3.6.3 Screen. Our system includes a 9" LCD display mounted in the center of the dashboard as shown in Figure 3.6.3a. It serves two purposes. In normal operation, when the driver has not indicated his intention to change lanes by engaging the turn signal, it functions as an input for controlling the previously described user parameters. Figure 3.6.3b shows an example control screen, with multiple bars indicating the current setting for a number of parameters. Bars can either be represented as several discrete blocks or a continuous percentage scale, depending on the desired level of control. In addition several modes are available to adjust each clearance zone independently or in several grouping

levels. The user adjusts the parameters using a control wheel.

In its other mode of operation, when the driver has engaged the turn signal and a lane exists to change into, the screen conveys lane change recommendations. When the system deems it safe to change, a green check mark is displayed. When unsafe, a “speed up” visual, “slow down” visual, or “wait” visual is displayed. These screens are shown in Figure 3.6.3c.

Figure 3.6.3a: LCD Screen Interface



Figure 3.6.3b: Parameter Adjustment Screen

The interface allows for values to be chosen in the continuous percent format (shown) or as 5 discrete blocks. The interface also contains several modes for specifying each parameter separately, linking the two front parameters and two back parameters (shown), and linking all clearance parameters and the lane change time into one parameter.

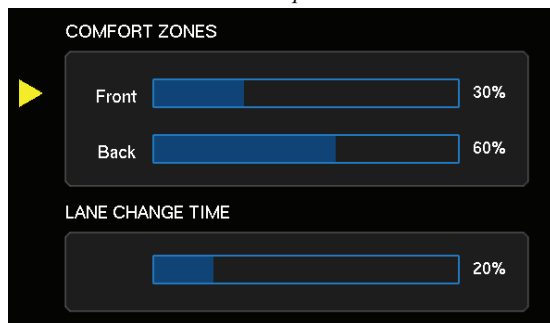



Figure 3.6.3c: Lane Change Recommendations



3.6.4 Automatic Turn Signal Control. Previous sections have explained how the system assists the driver of the car in making a safe lane change. Another interesting aspect however is how the system can communicate information to other drivers close by. A natural way to do so is by automatically controlling the turn signal if the driver forgets to do so himself.

Our automatic turn signal system works in exactly this way. It monitors the location of the car within the current lane along with the status of the turn signal. If the wheels of the car straddle a lane line without the driver engaging the turn signal, the system will turn the turn signal on. Once the system has automatically engaged the turn signal, two use cases are possible. If the driver proceeds with making a lane change the system will turn off the turn signal once the car is completely within the new lane. If the driver instead does not proceed with the lane change and moves back towards the center of the current lane, the system will turn off the signal automatically after 1 second.

The purpose of controlling the turn signal automatically in this way is two-fold. On one hand, it warns other drivers if the driver of this car is likely to make a lane change soon and thus the system communicates information to surrounding cars. On the other hand, it encourages the driver of the car to stay within the middle of the lane instead of drifting close to the lane lines.

4. User Study

Twelve Stanford University students were run as subjects in a preliminary user study. Each participant began the experiment from a central base camp and performed a trial consisting of two drives, accompanied by an experimenter. Both legs of the drive entailed driving the vehicle onto the highway and across a stretch of road. For the first drive, the participant drove across the highway for approximately 5 minutes until he reached the established endpoint. Once he exited, the experimenter had the participant fill out a paper questionnaire.

After completion, the participant was instructed to get back onto the highway, where he drove the exact reverse of the first route he had driven, returning to the base camp. Upon returning, the participant filled out the same questionnaire as previously, this time responding to the second drive (return back). Following this, the participant was given another questionnaire, which asked for evaluations about RASCL and the participant’s experience driving with it, in addition to open-ended questions about safety and potential improvements to the RASCL.

The order that the participants received the conditions was randomized to prevent any confounding systematic effects that could arise if the order was kept constant. Thus some participants drove the first leg with RASCL and the second leg without RASCL, while other participants drove the first leg without and the second leg with.

It should be noted that the user study was done before the automatic turn signal and the variable repeat rate of LEDs and sounds had been implemented.

4.1 Measures

We assessed the participant's experience of driving with RASCL by their responses to the questionnaire items. To measure feelings about driving experience, we asked "How well do the following words describe how you felt while driving?" Participants responded by rating a series of adjectives on a scale of one to ten, with ten marked by "Describes Very Well" and one marked by "Describes Very Poorly." Attitudinal measures reflecting participants' feelings toward the car and system were obtained by asking participants to rate a second set of adjectives, this time referring to "How well the following words describe the car." The same ten-point scales were used.

System aptitude was a measure of the degree to which participants felt that RASCL was intelligent and was determined by participant responses to three of the items on the questionnaire relating to the car: smart, intelligent, and unintelligent.

Overpowering measured the extent to which participants felt unpleasantly overpowered by the system and was determined by participant responses to one of the items on the questionnaire relating to the driving experience (comfortable) and three of the items relating to the car: authoritative, annoying, and dominant.

Perplexity measured confusion of the participant while driving with the system and was determined by participant responses to three of the items on the questionnaire relating to the driving experience (puzzled and confused) and one of the items relating to the car: confusing.

Efficacy measured the helpfulness and effectiveness of the system and was determined by participant responses to two of the items on the questionnaire relating to the car: helpful and inefficient.

For each of the above measures, the factors comprising it loaded on each other with a Cronbach's Alpha correlation coefficient of over 0.5, showing that the indices were reliable.

Other measures we attempted to analyze in our questions included indices of engagement (engaged, attentive), fear (fearful, afraid), enjoyment (enjoyable, fun to drive) and perceptions of vehicle safety and reliability (safe, trustworthy, dependable, reliable).

4.2 Results

As each participant experienced both conditions, we were able to compare each participant to him or herself. Consequently, we used a paired two-tailed t-test to determine if

there was a significant difference at the .05 alpha level between conditions (with/without RASCL) for each of our indices.

Significant main effects of driving with RASCL with regards to the indices *system aptitude*, *a sense of being overpowered*, *perplexity*, and *system efficacy* were found.

There was a statistically significant difference in how participants rated the vehicle with and without RASCL in aptitude, $M = 6.16$, $SE = 2.26$, with $t = 2.73$, $p = .02$. (Where M is the mean difference between with and without RASCL with regard to this index, SE is the standard error, t is the resulting t-statistic, and p is the probability that such a result would occur by chance if there were no difference between how the users felt about the two systems with regard to the index).

There was also a significant difference in how participants rated the sense of being overpowered with $M = 5.07$, $SE = 1.68$, $t = 3.02$, $p = .01$, the degree of perplexity experiences with $M = 2.48$, $SE = 1.01$, $t = 2.46$, $p = .03$, and the perceived efficacy with $M = 2.98$, $SE = .69$, $t = 4.32$, $p = .0012$. All statistically significant differences were positive in the direction of RASCL. For all other measures, none of the indices yielded a statistically significant difference.

4.3 Driver Behavioral Analysis

As part of the user study we logged the position of Junior, location of other cars, the driver's use of the turn signal, and other information. From these logs we were able to compute interesting statistics, such as the average lane change time and the average front and back clearance in the new lane during the duration of a lane change. The average lane change duration was three seconds with a standard deviation of one second.

Unfortunately, the data were too sparse to yield statistically significant results in the difference between values with using RASCL and not using RASCL. For example, amongst the 12 users only 3 made lane changes both with and without RASCL where a car existed within 40 meters behind Junior in the lane that the user was changing in to. Thus the sample size for the correlated t-test of the statistic "difference with and without RASCL on the average distance to the car behind Junior in the new lane at beginning of the lane change" was 3. Sample sizes for all statistics except time were 5 or less.

5. Discussion

5.1 Usability

The results of the user study indicate that participants found the vehicle to demonstrate a higher degree of aptitude and intelligence when they were driving with RASCL. Participants also felt that the vehicle was more helpful and greater in efficacy when driving with RASCL. Yet at the same time they also found the system to be more overpowering, to the extent that they felt vexed and unsettled. In addition, participants found the driving experience to be more confounding and perplexing when the system was in operation.

This seeming contradiction indicates that a lane changing assistance system is indeed useful, but that the specific interface of overt, sometimes-on sounds, and subtle, always-on lights, needs to be refined to make it less agitating, dominant, and confusing. A limitation of this study is that we cannot focus in on what specifically makes this warning system overpowering and perplexing without comparing it to other systems. Fortunately, this study paves the way for further studies which will compare the different types of systems from the matrix of possible combinations of overt and subtle signals, and ideally these studies will tell us exactly what elements of a system render its different qualities. Ultimately, we aim to determine what desirable elements make the system more user friendly, and what elements detract from its effectiveness and ease of use.

5.2 Limitations

From an engineering perspective, the system ultimately performed as designed, although we encountered a number of important limitations with the current implementation. One of the most significant is the system's reliance on near-perfect localization for proper functionality. Under most conditions, the RNDF-based localizer provided sufficient accuracy, but we observed certain conditions under which localization performed poorly, including the first few seconds of driving on the RNDF after entering the freeway and sections of road close to on ramps and off ramps.

There are two ways to look at this problem. On the one hand, improvements to the localizer are probably sufficient to address the observed failures. On the other hand, such failures really highlight an overreliance on the localizer in general. A truly robust system probably requires a way to more gracefully handle localization failure, especially if it is ever to operate in a semi- or fully-autonomous mode. Indeed, the ability to provide reliable recommendations based

only on local information could also make for better overall recommendations even when localization is good.

Another shortcoming we observed is the poor performance of the Velodyne laser scanner when it is raining. The water kicked up by Junior was detected and tracked as a car immediately behind us. In addition, the density of sensor readings for surrounding cars diminished substantially, likely due to water droplets scattering the lasers. This led to the intermittent loss of visibility of tracked cars.

5.3 Future Work

We believe the success of our system warrants further investigation. One area that is particularly interesting is better modeling of driver behavior. For instance, our system currently makes a number of simplistic assumptions about driver behavior such as a static, user-defined time to complete a lane change. Better approaches to the problem might involve a model based on experimentally-determined lane change statistics or an adaptive model. Additionally, it is important to extend our approach to harder, more complex lane change situations such as denser traffic and on ramp merging.

Future systems might also be able to utilize Junior's trajectory planning module – which already takes into account the dynamics of the vehicle and obstacle avoidance – as a way to test possible lane changes and predict whether it could be accomplished safely.

Finally, our exploratory user study provided valuable feedback on the effectiveness of our user interface, but clearly longer tests that allow drivers to get used to the system and allow us to gather more meaningful data are needed to truly understand how drivers view the system.

6. Acknowledgements

We would like to thank Professor Sebastian Thrun and Professor Clifford Nass for their guidance and support. We would like to extend a special thanks to David Stavens for maintaining the car throughout our tests and user study, for providing guidance and overall direction on the project, and for keeping current and participating in our frequent newsgroup discussions.

7. References

- [1] Centers for Disease Control and Prevention. 2003. *10 Leading Causes of Injury Death by Age Group Highlighting Unintentional Injury Deaths*. <ftp://ftp.cdc.gov/pub/ncipc/10LC-2003/PDF/10lc-2003.pdf>
- [2] Fatality Analysis Reporting System. *National Statistics*. <http://www-fars.nhtsa.dot.gov/Main/index.aspx>

- [3] *Improper Lane Change: A Collision Countermeasures Presentation.*
<http://my.safetyfirst.com/presentations/ImproperLaneChange.ppt>
- [4] Automotive Design Europe. 2008. *Hella provides lane change warning for BMW.* Automotive DesignLine-Europe.
<http://www.automotivedesign-europe.com/210200189>
- [5] Reed, P. *AutoVue, Lane Departure Warning System.*
<http://www.emunds.com/ownership/safety/articles/100176/article.html>
- [6] I.Gat, M.Benady and A.Shashua. 2005. *Monocular Vision Advance Warning System.* Detroit, MI: Automotive Aftermarket SAE World Congress & Exhibition.
- [7] Montemerlo, M. et al. 2008. *Junior: The Stanford Entry in the Urban Challenge.* Journal of Field Robotics.
- [8] Defense Advanced Research Projects Agency (DARPA). 2007. *Route Network Definition File (RNDF) and Mission Data File (MDF) Formats.*
http://www.darpa.mil/grandchallenge/docs/RNDF_MDF_Formats_031407.pdf
- [9] Kalman, R.E. 1960. *A new approach to linear filtering and prediction problems.* Journal of Basic Engineering 82 (1): 35-45.