Rapidly Adapting Machine Vision for Automated Vehicle Steering

Dean Pomerleau and Todd Jochem, Carnegie Mellon University and AssistWare Technology Inc.

THE RALPH VISION SYSTEM HELPS AUTOMOBILE DRIVERS STEER, BY SAMPLING AN IMAGE, ASSESSING THE ROAD CURVATURE, AND DETERMINING THE LATERAL OFFSET OF THE VEHICLE RELATIVE TO THE LANE CENTER. RALPH HAS PERFORMED WELL UNDER EXTENSIVE TESTS, INCLUDING A COAST-TO-COAST, 2,850-MILE DRIVE.

NEARLY 15,000 PEOPLE DIE EACH year in the US in single-vehicle roadway-departure crashes.[1] These accidents often stem from driver inattention or impairment (for example, fatigued or intoxicated drivers). A system could eliminate many of these crashes if it warned drivers when their vehicle was departing the roadway and controlled the vehicle's lateral position to keep it in its lane. Nearly 70% of such crashes occur in rural or suburban settings on undivided, two-lane roads.[1] Because these roads will most likely not be upgraded in the foreseeable future, a crash-prevention system must rely on the existing road structure.

In this article, we describe Ralph (Rapidly Adapting Lateral Position Handler), a vision system developed jointly by Carnegie Mellon University and AssistWare Technology Inc. Ralph decomposes vehicle steering into three steps: sampling the image, determining the road curvature, and assessing the lateral offset of the vehicle relative to the lane center. Ralph combines the outputs of the latter two steps into a steering command, which it can send to the steering motor on our Navlab 5 testbed vehicle (see Figure 1) for autonomous steering control.[2] As a road-departure warning system, Ralph can also compare this command with the human driver's steering direction.

Other steering systems

Research on other steering systems has focused on machine-vision techniques that detect particular features in video images of the road in front of the vehicle and determine the desired vehicle trajectory on the basis of the relative positions of these features. Many of these systems track specific features, such as lane markings, from one image to the next.[3-6] Others detect regions of the image that represents the road on the basis of features such as color[7,8] or texture.[9]

Nevertheless, all of these systems share one characteristic: They have a detailed, a priori model of the road's appearance, and they locate specific features via hand-programmed detection algorithms. Unfortunately, roads are not always cooperative. Road markings vary dramatically, depending on the type of road and the state or country where the road is located. For example, many California freeways embed regularly spaced reflectors—not painted markings—to delineate lane boundaries. Further challenges arise from the environmental context, which can greatly affect road appearance. The road's appearance often varies dramatically because of changes in illumination due to shadows, glare, or darkness and because of obstructions by other vehicles, rain, snow, salt, or other objects. These variations often invalidate the assumptions underlying vision algorithms, leading to poor road-detection performance.

Alternative approaches combine machine-vision and machine-learning techniques for enhanced systems capable of handling variations in road appearance.[4,10,11] One such system, Alvinn (which we've developed),

0885-9000/96/$4.00 © 1996 IEEE

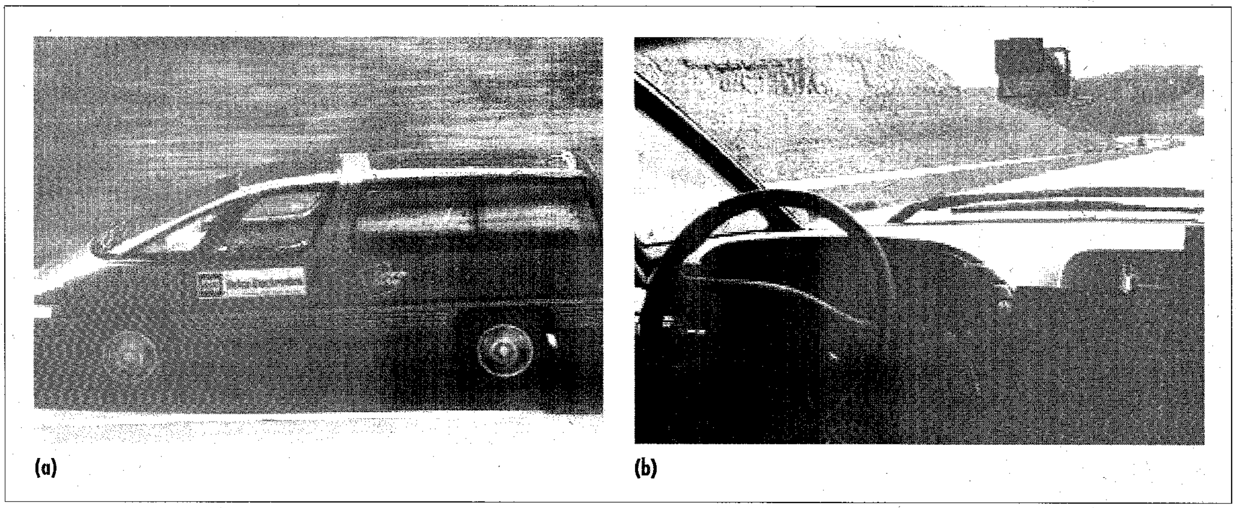

Figure 1. The Navlab 5 testbed vehicle: (a) exterior; (b) interior.

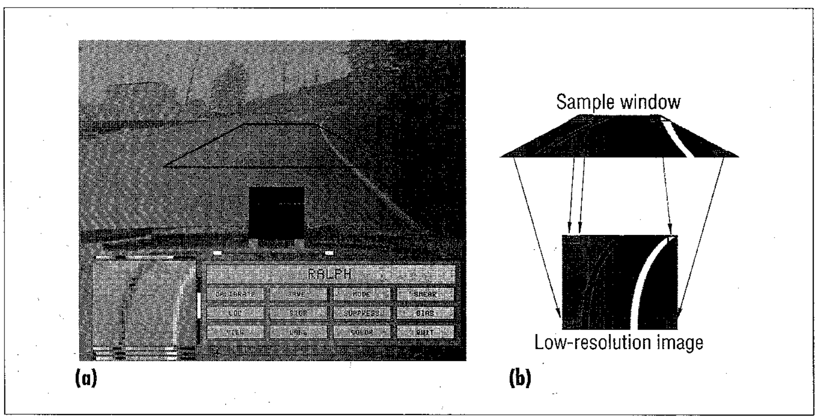

Figure 2. (a) Forward-looking image; (b) Ralph's sampling strategy.

uses an artificial neural network to learn the characteristic features of particular roads under specific conditions. Alvinn uses this learned road model to determine how to steer the vehicle to keep it in its lane.

Although systems of this type have been quite successful at handling a wide variety of road types under many different conditions, they have several shortcomings. First, the process of adapting to a new road requires a relatively extended "retraining" period (at least several minutes). While this adaptation process is relatively quick by machine-learning standards, it is unquestionably too long in domain-like autonomous driving, where the vehicle may be traveling at nearly 30 meters per second. Second, the retraining process invariably requires human intervention of one form or another. These systems employ a supervised learning technique such as back propagation, where the driver must physically demonstrate the correct steering behavior for the system to learn.

A truly flexible system should exploit whatever features are available to determine vehicle location, adapt almost instantly when the available features change, and perform this adaptation without human supervision. Ralph demonstrates these characteristics.

Ralph's sensor configuration

Figure 2a depicts a typical scene of the road ahead, as imaged by a video camera mounted next to the rearview mirror on Navlab 5. Ralph can use either black-and-white or color images via a color-based contrast-enhancement technique.[10] Obviously, many parts of this image are not relevant to the driving task (for example, the parts depicting the sky or showing the vehicle dashboard). Ralph eliminates these parts, and processes only the portions of the scene inside the red trapezoid. Although the lower and upper boundaries of this trapezoid vary with vehicle velocity (moving further ahead of the vehicle, toward the top of the image, as vehicle speed increases), they are typically about 20 and 70 meters ahead of the vehicle, respectively.

The second, and perhaps more important, aspect of the trapezoid's shape is its horizontal extent. It is configured so that its width on the ground plane is identical at each row of the image. The horizontal distance that each row of the trapezoid encompasses is approximately 7 meters, about twice the width of a typical lane. The trapezoid is selectively sampled according to the strategy depicted in Figure 2b. This sampling process creates a low-resolution (30×32 pixels) image where important features such as lane markings (which converge toward the top of the original image) appear parallel. This image resampling is a simple geometric transformation and requires no explicit feature detection.

Curvature calculation. The parallelization of road features is crucial for the second step of Ralph processing: curvature determination. To determine the curvature of the road ahead, Ralph hypothesizes a possible curvature, subtracts it from the parallelized low-resolution image, and tests to see how well the hypothesized curvature has "straightened" the image.

Figure 3 depicts the process Ralph uses to determine curvature. In this example, Ralph has hypothesized five curvatures for the original image. In each curvature, Ralph has differentially shifted the image's rows to "undo" the curve and straighten out the image features. For left-curve hypotheses Ralph shifts rows toward the right, and for right-curve hypotheses it shifts rows toward the left. For the more extreme hypothesized curvatures (on the far left and right), Ralph shifts the rows of the original image further than for the less extreme curvatures (in the middle). For all hypothesized curvatures, Ralph shifts rows near the top of the image (corresponding to regions on the ground plane further ahead of the vehicle) further horizontally than rows near the bottom of the image. This differential shifting compensates for the road's greater displacement at the top of the image (far ahead of the vehicle) than at the bottom. The exact shift distance for each row in the transformed images depends on both the geometry of the camera and the particular curvature hypothesis being tested.

In Figure 3, the second curvature hypothesis from the right (corresponding to a shallow right turn) yields a transformed image with the straightest features and is, therefore, the "winning" hypothesis. Figure 4 illustrates how Ralph scores the straightness of each hypothesis. After differentially shifting the rows of the image according to a particular hypothesis, Ralph vertically sums the columns of the resulting transformed image to create a scan-line intensity profile (shown in the two curves at the bottom of Figure 4). When the visible image features are correctly straightened, sharp discontinuities between adjacent columns occur in the image (right scan-line intensity profile in the figure). In contrast, when the hypothesized curvature has shifted the image features too much or too little, there are smooth transitions between adjacent columns of the scan-line intensity profile (left profile in the figure). By summing the maximum absolute differences between intensities of adjacent columns in the scan-line intensity profile, Ralph can quantify this property to determine the curvature hypothesis that best straightens the image features.

An important attribute of this technique for determining road curvature is that it is entirely independent of the particular features present in the image. As long as visible features run parallel to the road, this technique exploits them to determine road curvature. Those features need not be at any particular position relative to the road, and they need not have distinct boundaries (characteristics required by systems that use detailed a priori road models and edge detection).

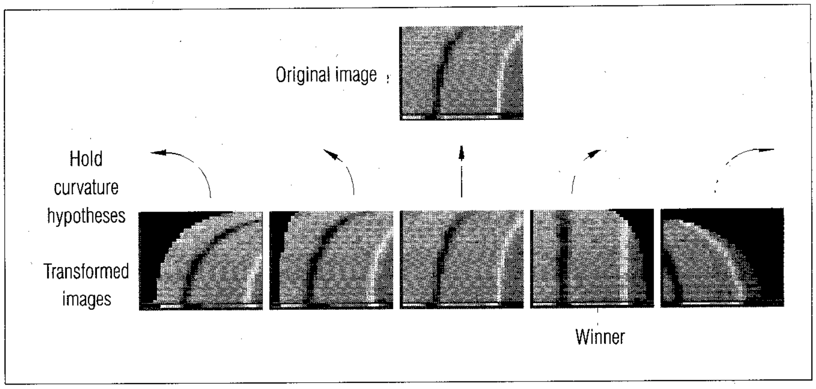

Figure 3. Ralph's curvature hypotheses.

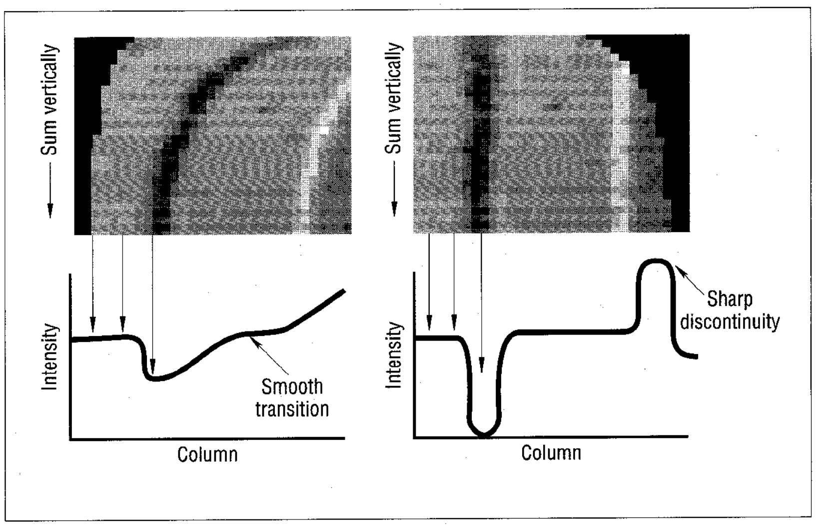

Figure 4. Ralph's curvature scoring technique.

Lateral offset calculation. Next, Ralph determines the vehicle's lateral position relative to the lane center. It uses a template-matching approach on the scan-line intensity profile generated in the curvature estimation step. The scan-line intensity profile is a one-dimensional representation of the road's appearance as seen from the vehicle's current lateral position. By comparing this current appearance with the appearance of a template created when the vehicle was centered in the lane, Ralph can estimate the vehicle's current lateral offset.

Figure 5 illustrates this lateral-offset estimation procedure in more detail. The current scan-line intensity profile is on the left, and the template scan-line intensity profile (generated when the vehicle was centered in the lane) is on the right. By iteratively shifting the current profile to the left and right, the system can determine the shift required to maximize the match between the two profiles (as measured by the correlation between the two curves). The shift distance required to achieve the best match is proportional to the vehicle's current lateral offset.

As with curvature determination, this process does not require that any particular features be present in the image. As long as the visible features produce a distinct scan-line intensity profile, the correlation-based matching procedure can determine the vehicle's lateral offset. In particular, even features without distinct edges, such as pavement discoloration due to tire wear or oil spots, generate identifiable profile variations that Ralph can exploit to determine lateral offset. This is a performance feature that edge-based road detection systems do not share.

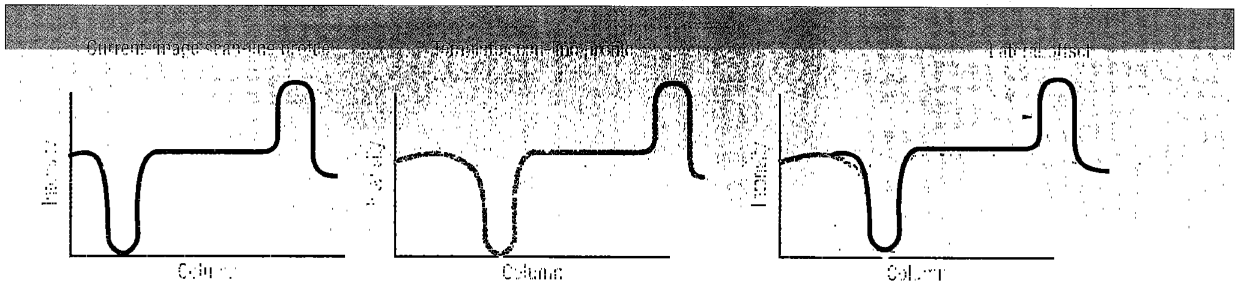

Figure 5. Ralph's technique for determining lateral offset.

Adaptation to changing conditions. Another important feature of Ralph stems from the simplicity of its scan-line intensity profile representation of road appearance. For Ralph to handle a new road type, only the 32-element template profile vector need be modified, and this is extremely easy. In the current Ralph implementation, there are four ways to adapt the template to changing conditions.

In the first method, the driver centers the vehicle in its lane and presses a button to indicate that Ralph should create a new template. In under 100 milliseconds, Ralph performs the processing steps described above to create a scan-line intensity profile for the current road, and then saves it as the template. From that point on, Ralph can either actively control the steering wheel (the driver lets go of the wheel) or warn the driver of road-departure danger (the driver continues to steer manually), using the newly created template to determine the vehicle's position relative to the lane center. (In another warning mode, the driver continues steering manually; but instead of merely warning the driver when the vehicle starts to drift off the road, Ralph can give a momentary nudge to the steering wheel to bring the vehicle back into the lane.)

In the second method, Ralph selects the template from a library of stored templates recorded previously on various roads. Ralph selects the best template for the current conditions by testing several of these previously recorded templates to determine which has the highest correlation with the scan-line intensity profile created for the current image.

The third method of template modification occurs after an appropriate template has been selected. During operation, Ralph slowly evolves the current template by adding a small percentage of the current profile to the existing template. This lets the current template adapt to gradual changes in the road's appearance, such as those caused by changes in the sun's angle.

Ralph handles more abrupt scene changes, such as changes in lane-marker configuration, via the final and most interesting template-modification strategy. In this technique, Ralph uses the appearance of the road just ahead of the vehicle (the foreground) to determine the vehicle's current lateral offset and the curvature of the road ahead. At the same time, Ralph is constantly creating a new *rapidly adapting template* based on the appearance of the road far ahead (typically 70 to 100 meters) of the vehicle. Ralph creates this rapidly adapting template by processing the top rows of the image, far ahead of the vehicle, in the same manner as described previously. Ralph assumes that the road's curvature between the foreground and the background is nearly constant, so it can project where the road will be and, hence, what the new template should look like when the vehicle is centered in its lane.

In places where the road's appearance changes dramatically, such as the entrance to a tunnel or an on-ramp to a highway, Ralph uses this technique to quickly create a template appropriate for the new road appearance. When the vehicle actually reaches the new road, Ralph determines that the template it was previously using is no longer appropriate, because it does not match the scan-line intensity profile of the current image. It therefore swaps in the rapidly adapting template and continues driving. This rapid adaptation occurs in about 2 seconds without any human intervention.

Ralph's performance

We have conducted extensive laboratory, test-track, and on-road experiments with the Ralph system as part of a program by the National Highway Traffic Safety Administration (NHTSA) to assess roadway-departure collision-avoidance systems.[12] The results indicate that Ralph can accurately estimate a vehicle's lateral position on the road along with the curvature of the road ahead, under a wide variety of conditions.

Laboratory tests. An important factor in determining autonomous driving effectiveness is the sensing system's accuracy. The crucial accuracy metric for Ralph is how well it can estimate the location of the road ahead of the vehicle, because road location determines the direction in which to steer the vehicle.

To quantify Ralph's capability to accurately determine the position of the road ahead, we conducted controlled laboratory tests to accurately measure the road's actual location. To facilitate these measurements, we collected high-quality video sequences of road scenes, using a Umatic 3/4-inch VCR. We gathered these scenes with the Navlab 5 test vehicle, using the same camera mounted in the same location (next to the rearview mirror) as in the on-vehicle experiments (described later). These sequences included both day and night operation as well as images of various road types, including rural roads and multilane divided highways. The test road sequences recorded on videotape were all between 4 and 9 miles long. While recording the sequences, the driver repeatedly changed the vehicle's lateral position within the lane to obtain a wide range of images.

After replaying the video sequences in the laboratory, we used Ralph to track the road. More specifically, Ralph combined its estimates of the vehicle's lateral offset and the curvature of the road ahead into an estimate of the lane-center location at 1 second (about 25 meters) ahead of the vehicle.

In real time, we compared Ralph's lane-center position estimate with the estimate that the experimenter provided manually. The experimenter continuously indicated his estimate of the lane-center location, using a computer mouse to keep a crosshair centered over the right-lane marking at 1 second ahead of the vehicle in the image.

Table 1 summarizes the results of these

Table 1. Accuracy of Ralph's lane-location estimation.

Condition	Mean error (cm)	Error, standard deviation (cm)
Daytime highway	11.4	14.3
Daytime highway with shadows	13.8	18.9
Nighttime highway	11.1	13.8
Daytime rural road	13.7	16.2
Daytime rural road with glare	15.8	17.2
Nighttime rural road	13.8	16.8
Average	13.2	16.2

tests. For each condition tested, the table shows the mean and standard deviation of the difference between Ralph's and the experimenter's estimate of the lane-center position. In general, Ralph's performance was quite good in all the conditions tested, with a total mean disagreement between Ralph and the experimenter of 13.2 cm—just slightly larger than the width of a typical lane-edge marker.

As expected, we observed lower mean and standard-deviation errors in the conditions with the most consistent features. Figure 6 shows one such situation, depicting a daytime highway scene where the lane markers are clearly visible. Under these conditions, the mean disagreement between Ralph and the experimenter was 11.4 cm. The variance of the disagreement was 14.3 cm. However, we can attribute a substantial portion of the disagreement to inconsistency in the experimenter's estimate of the lane-center position. Using a mouse to accurately indicate the lane position 20 meters ahead is difficult. In fact, in a series of repeatability tests, we determined that the experimenter's estimate of lane position over two different trials on the same section of videotape varied by an average of 7.3 cm.

On the same stretch of highway, under conditions of heavy shadows (see Figure 7a), the mean and standard deviation of Ralph's lane-position estimation error increased to 13.8 cm and 18.9 cm, respectively. These error increases were due primarily to the camera's limited dynamic range, which caused the shadowed regions of the image to be black and the areas in the sunlight to be saturated.

Ralph's estimation of lane location on the same stretch of highway at night improved slightly. As Figure 7b shows, the lane markers were very distinct in this situation, resulting in a mean error of 11.1 cm and a standard deviation of 13.8 cm.

Ralph's performance on rural roads such as the one in Figure 7c was fairly similar to the highway results. The mean and standard deviation under favorable daytime conditions did increase slightly over the corresponding figures for favorable daytime highway images, to 13.7 cm and 16.2 cm, respectively. Two factors were primarily responsible for these increases. First, hills on the rural roads changed the perspective of the camera relative to the road. This resulted in slight additional errors in lane-position estimation, particularly at grade transition points. Second, several cross streets intersected the section of rural road tested, which occasionally resulted in momentary inaccuracy when the lane markers disappeared. Nevertheless, the error increases due to these effects were small, on the order of 2 cm.

One problem with lane-tracking systems that rely exclusively on lane markers to locate the road ahead is that they sometimes have difficulty when glare off the pavement makes the markers hard to find. This type of glare typically occurs when the pavement is wet or the sun is low on the horizon. To quantify how these conditions affect Ralph, we collected a video sequence on the same rural road during the early morning hours heading into the rising sun. Figure 7d presents a sample image from this sequence. As expected, the mean and standard deviation of Ralph's error increased to 15.8 cm and 17.2 cm. However these increases were slight, again in the range of 2 cm. Ralph still accurately located the road ahead by adapting its processing to exploit the boundary between the bright pavement and the dark shoulder. This capability to adapt to changing conditions proved particularly valuable in the on-road tests.

In summary, the laboratory tests indicate that under various conditions, Ralph can localize the position of the road ahead of the vehicle to approximately within the width of a single lane marker. To further characterize Ralph's capability to perform repeatedly and reliably, we also performed extensive test-track and on-road experiments.

Test-track experiments. We conducted additional controlled experiments on a road segment outside of Pittsburgh that is often used for testing. These tests involved repeatedly driving on the same stretch of roadway at different speeds, with no other vehicles around.

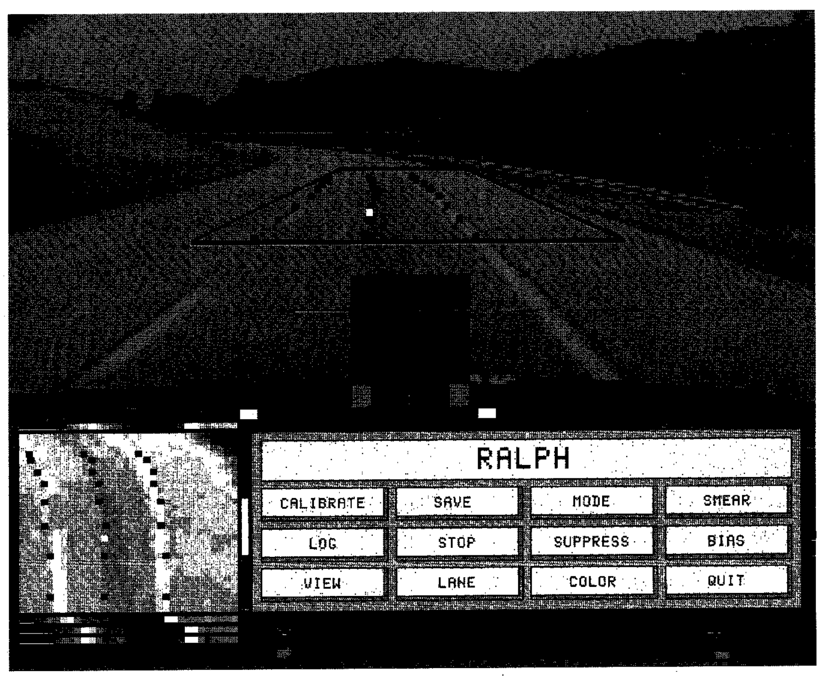

Figure 6. Ralph processing a daytime highway image.

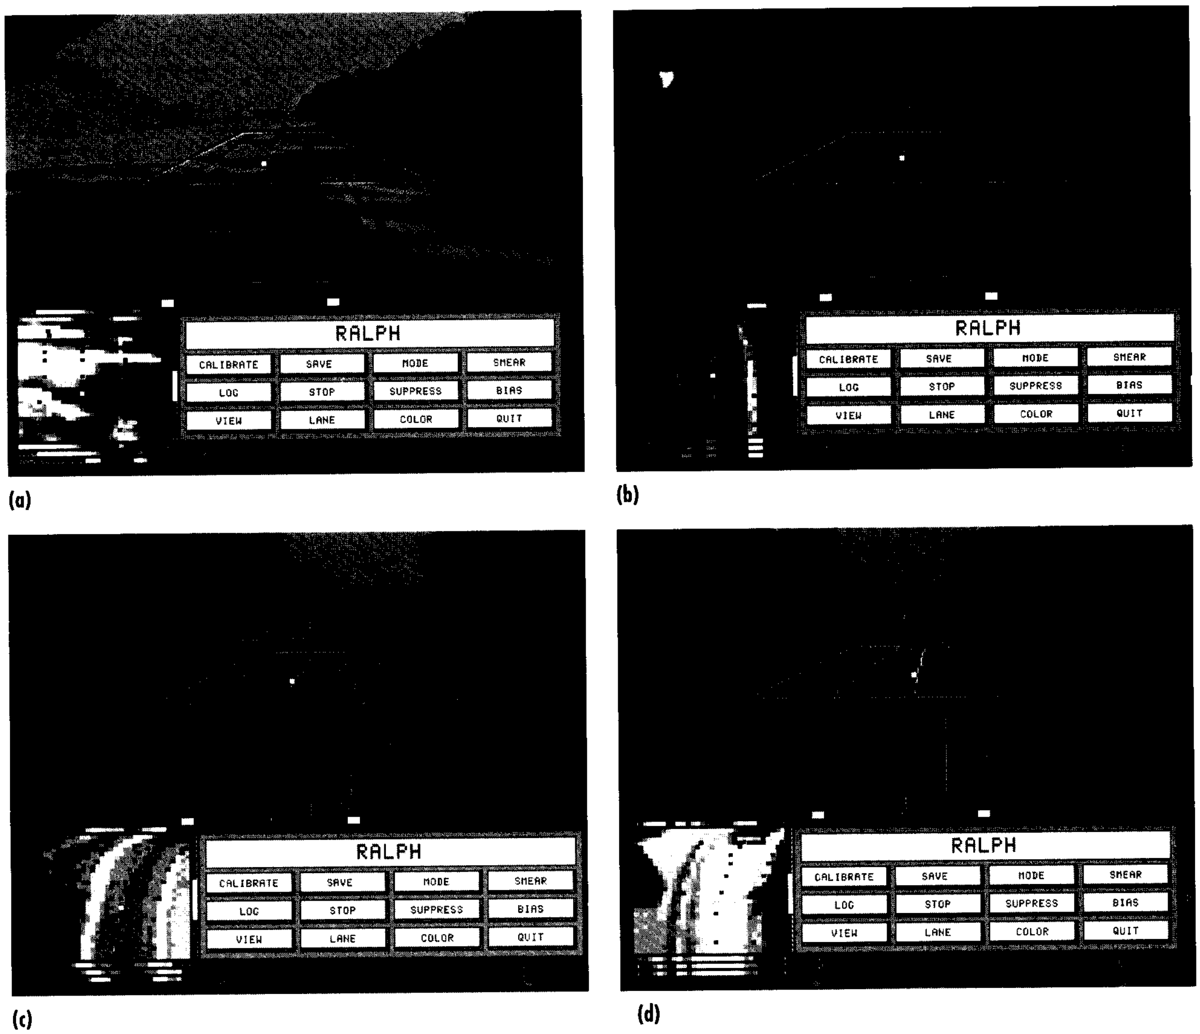

Figure 7. Ralph processing (a) a daytime highway image with heavy shadows; (b) a nighttime highway image; (c) a daytime rural-road image; (d) an early-morning, rural-road image with glare off the road.

In the videotape experiments presented above, the goal was to quantify Ralph's capability to project the position of the road ahead by combining its estimate of the vehicle's lateral position and its estimate of the road's curvature. In the first set of test-track experiments, the goal was to tease apart this combination and measure Ralph's capability to estimate the curvature. In this experiment, we manually drove the Navlab 5 test vehicle through the S-curve shown in Figure 8a.

Careful measurement of the first curve indicated an average radius of curvature of approximately 343 meters. Figure 8b shows

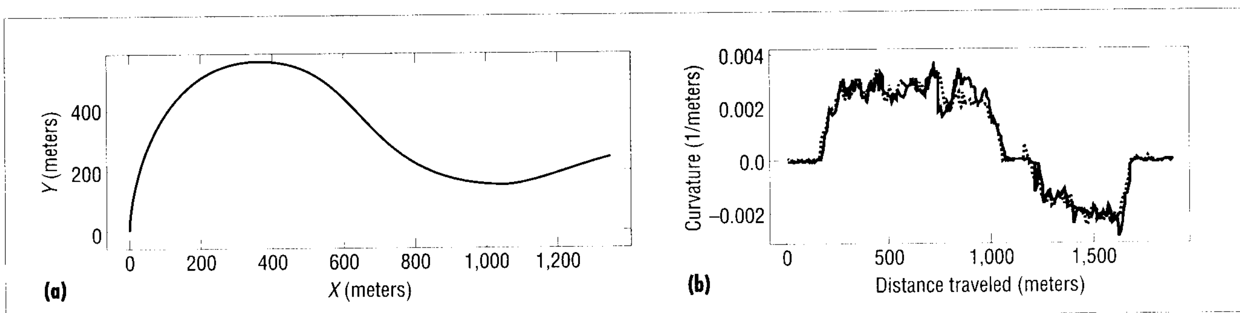

Figure 8. Experiment to assess Ralph's effectiveness in estimating curvature: (a) S-curve used for testing Ralph; (b) Ralph's curvature estimate on two traversals through the S-curve.

Ralph's estimate of the road curvature during two traversals of the entire S-curve at 55 mph.

Note the consistency in the curvature estimate between the two traversals. Ralph's mean estimate for radius of curvature during the first traversal of the first curve was 373 meters, and the mean for the second traversal was 374 meters. Not only are the two estimates extremely close, but they correspond quite closely with the measured radius of 343 meters. Moreover, we can at least partially attribute the 30-meter discrepancy between the measured curve radius and Ralph's estimate to uncertainty in the manual curvature measurement.

The next set of experiments tested whether Ralph could detect anomalous driver behavior. This information is valuable not for autonomous driving, but rather for determining Ralph's applicability as a tool for preventing roadway departure crashes. Again, the driver drove twice through the S-curve at 55 mph. The first time, the driver concentrated on accurate driving. The second time, the driver was momentarily distracted by an in-cab distractor task, which required the driver to glance toward the back of the vehicle for up to two seconds. The goal was to determine if the lane-tracking output that Ralph produces could detect the lane deviations resulting from this momentary inattention.

Figure 9 graphs Ralph's estimate of the vehicle's lateral position, both during normal driving and while the driver performed the distractor task. As the graph indicates, the relatively large magnitude of the lane deviations resulting from momentary distraction are clearly discernible compared with the driver's normal lane deviations. This characteristic is extremely useful when using Ralph as a roadway-departure warning system.

The results of these test-track experiments indicate that Ralph can repeatedly detect both the curvature of the road ahead and excessive lane deviation by the driver. However these experiments neglected two important aspects of the autonomous driving task. First, in the test-track experiments, Ralph was passively monitoring the vehicle's position on the roadway and the curvature of the road ahead. These tests did not assess Ralph's capability to combine those measurements into a command for the steering wheel that would keep the vehicle centered in its lane. Moreover, we conducted these experiments under favorable weather and lighting conditions.

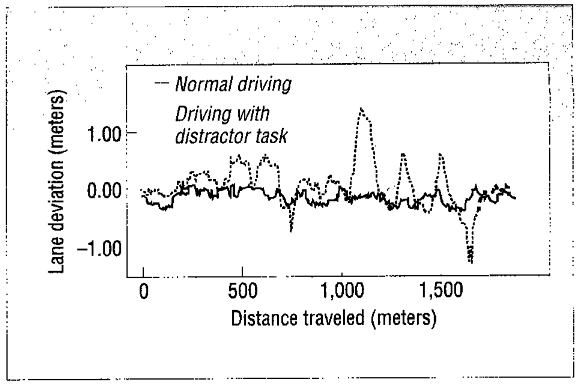

Figure 9. Lane deviations in normal driving and when the driver is distracted.

Open road tests. One of the most significant potential drawbacks of driving systems that rely on video cameras for sensor input is their susceptibility to adverse conditions. Systems that rely on visible features to determine the vehicle's position on the road can have trouble when these distinctions become difficult to detect because of adverse weather, poor lighting, or degraded pavement. To quantify this effect, we conducted a series of on-road tests of the Ralph system.

The culmination of these experiments was a 2,850-mile drive from Washington, D.C., to San Diego, in which we used Ralph's steering commands to control the Navlab 5 testbed vehicle. Except for a few detours, the trip exclusively involved highway driving. The trip included many of the difficulties typically encountered during normal driving: driving at night, during sunset when the sun is low on the horizon, during rain storms, on poorly marked roads, and through construction areas.

During the trip, we collected statistics on Ralph's driving performance. The primary metric was the percent of the total trip distance for which Ralph was controlling the steering wheel. To measure this value, we assumed that if the steering wheel position disagreed significantly from Ralph's commanded position, the safety driver had taken control of the wheel. Specifically, when the steering direction suggested by Ralph differed from the actual steering wheel position such that following Ralph's steering arc at the current speed would result in a difference in lateral acceleration of 0.04 g (where 1 g = 9.8 meters/second^2) or greater, then we judged that the safety driver had overridden Ralph.

Overall, the results were quite encouraging. Using the above metric, Ralph was able to steer the vehicle autonomously for 98.1 percent (2,796/2,850 miles) of the trip. Because the system adapts to changing conditions, Ralph was able to drive in situations that would be difficult for other lane-keeping systems, particularly those that rely on finding distinct lane markers. Figures 10 and 11 illustrate some of the different situations that Ralph handled.

Some of the roads were just as you might expect for a major highway: nice pavement and good lane markings. Even when the lane markers were missing, Ralph continued driving by exploiting the boundary between the pavement and the off-road area. However, stretches of road without lane markers proved quite difficult at night, when the edge formed by the pavement boundary was no longer visible. In particular, on the third night of the trip, a 10-mile stretch of new, unpainted highway accounted for a significant portion of the 1.9% of the distance that Ralph couldn't drive.

Rain proved to be less of a problem. Even when the specular reflection off the wet pavement obscured the lane markings, Ralph keyed off other, more subtle variations in the road's appearance to determine how it should steer. These additional features typically formed from water pooling in ruts on the road and by the tires of vehicles ahead leaving tracks on the wet pavement.

West of the Rocky Mountains, there were stretches of very poor roads, two of which are shown in Figure 10. Often the lane markers were nearly invisible because of wear (Figure 10a). And there were several long stretches of construction where the road was composed of very fine, packed gravel, without any lane markings (Figure 10b). During these stretches, Ralph continued driving by exploiting the differences in appearance between the packed and loose gravel.

The freeways in California posed an interesting challenge. Instead of having painted markings to delineate lanes, they have reflectors that are nearly invisible during the day. In these situations, Ralph followed the diffuse discoloration from the oil spot down the center of the lane. Ralph also performed well on the Interstate-15 High Occupancy Vehicle (carpool) lane into San Diego, which has no visible lane markings but a strong

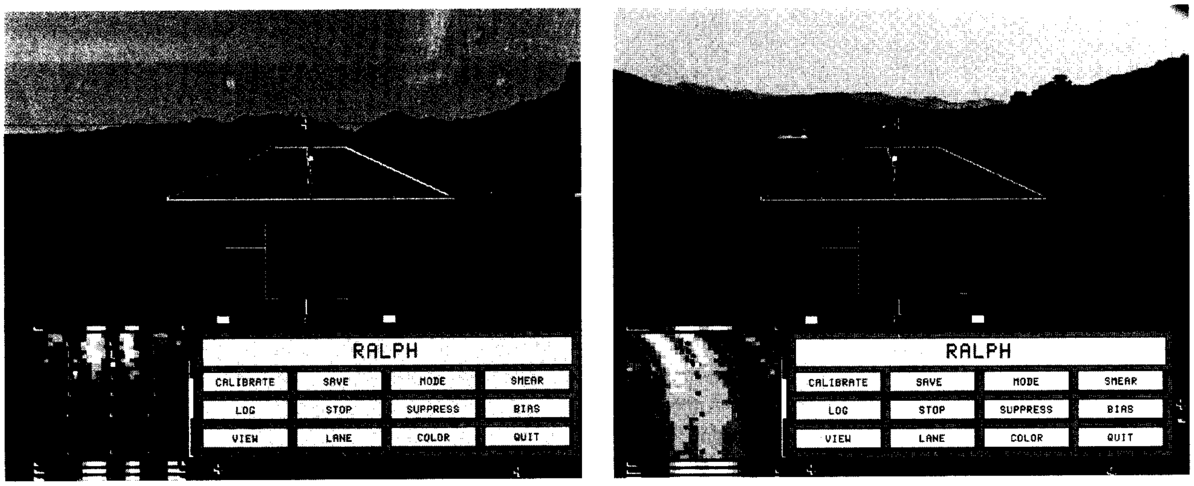

(a) (b)

Figure 10. (a) Road with severely worn markings; (b) unpaved road.

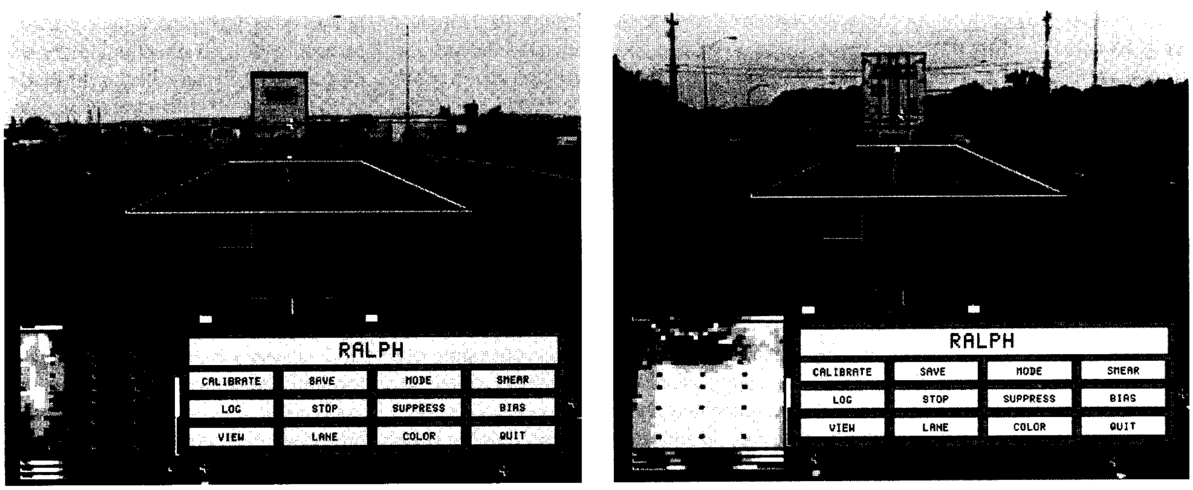

(a) (b)

Figure 11. Challenging images from city driving.

boundary between the cement road surface and the asphalt shoulder.

The situation that most challenged the system was in city traffic, when the road markings were either missing or obscured by other vehicles (see Figure 11). In such cases, Ralph was often able to lock onto the vehicle ahead and steer to follow it. In fact, Ralph used this technique to follow two trucks around nearly the entire Denver beltway. In the rare situations when Ralph became truly confused, it typically recognized that it could not correctly steer and signaled to the safety driver to take control.

EXTENSIVE TESTS OF THE RALPH vision-based lane-position estimation system indicate that it can accurately detect a vehicle's position and orientation relative to the roadway in a wide variety of situations and can use this information to steer our testbed vehicle. Current work on Ralph focuses on minimizing those few remaining conditions that are difficult for Ralph. To help Ralph overcome those difficulties, we are exploring techniques such as active camera control to focus the system's attention on important aspects of the scene. In addition, work is under way to develop techniques that will let Ralph reliably determine error bounds as it estimates the location of the road ahead.

The simplicity of the Ralph algorithm suggests that a custom hardware implementation is feasible. This has the potential to dramatically reduce both the size and cost of subsequent versions of Ralph. Our eventual goal is to build a system that is small enough to fit behind the rearview mirror and inexpensive enough to sell as an option on passenger cars. Initially, such a system would simply warn the driver when he is drifting off the road. But, in time, a system might assume at least partial control, relieving the driver of the monotonous task of steering—just as standard cruise control has done for maintaining vehicle speed.

References

1. J.S. Want and R.R. Knipling, "Single-Vehicle Roadway Departure Crashes: Problem Size Assessment and Statistical Description," Tech. Report DTNH-22-91-C-03121, National Highway Traffic Safety Administration, Washington, D.C., 1993.
2. T. Jochem et al., "PANS: A Portable Navigation Platform," *Proc. 1995 IEEE Symp. Intelligent Vehicles*, IEEE Press, Piscataway, N.J., 1995, pp. 107–112.
3. E.D. Dickmanns et al., "The Seeing Passenger Car 'VaMoRs-P,'" *Proc. 1994 IEEE Symp. Intelligent Vehicles*, IEEE Press, Piscataway, N.J., 1994, pp. 68–73.
4. K.I. Kim et al., "An Autonomous Land Vehicle: Design Concept and Preliminary Road-Test Results," *Proc. 1993 IEEE Symp. Intelligent Vehicles*, IEEE Press, Piscataway, N.J., 1993, pp. 146–151.
5. K. Kluge and C. Thorpe, "Representation and Recovery of Road Geometry in YARF," *Proc. 1992 IEEE Symp. Intelligent Vehicles*, IEEE Press, Piscataway, N.J., 1992, pp. 114–119.
6. M. Nashman and H. Schneiderman, "Real-Time Visual Processing for Autonomous Driving," *Proc. 1993 IEEE Symp. Intelligent Vehicles*, 1993, pp. 373–378.
7. J. Crisman and C. Thorpe, "Color Vision for Road Following," *Vision and Navigation: The CMU Navlab*, C. Thorpe, ed., Kluwer Academic, Boston, 1990.
8. M. Marra, T.R. Dunlay, and D. Mathis, "Terrain Classification Using Texture for the ALV," Tech. Report 1007-10, Martin Marietta Information and Communications Systems, Denver, 1988.
9. J. Zhang and H. Nagel, "Texture-Based Segmentation of Road Images, *Proc. 1994 IEEE Symp. Intelligent Vehicles*, IEEE Press, Piscataway, N.J., 1994, pp. 260–265.
10. D.A. Pomerleau, *Neural Network Perception for Mobile Robot Guidance*, Kluwer Academic, Boston, 1994.
11. M. Rosenblum and L.S. Davis, "The Use of a Radial Basis Function Network for Visual Autonomous Road Following," *Proc. 1993 IEEE Symp. Intelligent Vehicles*, IEEE Press, Piscataway, N.J., 1993, pp. 432–439.
12. D. Pomerleau et al., "Run-Off-Road Collision Avoidance Using IVHS Countermeasures: Task 2, Vol. II," Tech. Report DTNH22-93-C-07023, National Highway Traffic Safety Administration, Washington, D.C., 1996.

Dean Pomerleau is a research scientist at Carnegie Mellon University with a joint appointment in the Robotics Institute and the School of Computer Science. His research focuses on perception and learning for intelligent vehicles. He received a PhD from CMU in 1992. Pomerleau directs the Single Vehicle Roadway Departure Prevention Project, a partnership of universities and companies, sponsored by the US Department of Transportation, that develops and tests warning systems to prevent highway accidents. He is also a principal investigator on the Automated Highway System Program, a consortium of companies, universities, and the US DOT that aims to improve the safety, efficiency, and convenience of highway travel via smart-vehicle and smart-highway technology. In addition, Pomerleau is leading an effort to commercialize intelligent vehicle technology as president of AssistWare Technology Inc. Readers can contact Pomerleau at the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213; pomerlea@cs.cmu.edu or deanp@assistware.com; http://www.cs.cmu.edu/~pomerlea.

Todd Jochem is a postdoctoral research associate at Carnegie Mellon University's Robotics Institute. He received a PhD from CMU in 1996. His thesis dissertation explored using vision-based focus-of-attention techniques to let an autonomous vehicle execute driving tasks such as lane transitions and intersection navigation. He is part of CMU's Automated Highway System Team and is leading CMU's development effort for the 1997 AHS Technical Feasibility Demonstration. Readers can contact Jochem at the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213; tjochem@cs.cmu.edu; http://www.cs.cmu.edu/~tjochem.