Vision Guided Lane Transition

Todd Jochem, Dean Pomerleau, Charles Thorpe

{tjochem, pomerlea, cet}@ri.cmu.edu
Phone: 412-268-3260 Fax: 412-268-5571
The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Abstract

Many systems have been created which can keep an autonomous vehicle within a driving lane, but little experimental work has been reported that describes methods to transition a vehicle between lanes. Three techniques to accomplish lane transition using the ALVINN lane keeping system are reported here. The most basic involves intelligently switching between two trained ALVINN networks. The other techniques use active control of virtual camera views to move the vehicle into the destination lane.

Introduction

Much progress has been made toward solving the autonomous lane keeping problem using vision-based methods. Systems have been demonstrated which can drive robot vehicles at high speeds for long distances. The current challenge for vision-based on-road navigation researchers is to create systems that maintain the performance of the existing lane-keeping systems, while adding the ability to execute tactical level driving tasks like transitioning from one lane to another.

Lane transition is assuming increasing importance because of the requirements for new advanced traffic control projects like the Automated Highway System (AHS). For the AHS, moving vehicles into and out of automated lanes is the primary situation that lane transition maneuvers are needed and one which is on the critical path toward deployment of a full scale AHS.

In order to successfully transition from one lane to another, a simple geometric model of the road is needed. Using the information contained in the model, virtual camera views [2] can be positioned which allow existing lane keeping systems to be used for lane transition maneuvers. This paper presents three techniques which autonomously transition a vehicle between lanes. These methods are based on controlling the lateral placement of virtual camera views using a simple road model and control scheme for guidance. They are not strictly tied to any specific lane keeping system, only requiring that the system take images as input and produce a point on the road to drive over and a measure of its internal confidence in this point as output. When implemented using a slightly modified version of the ALVINN [4] lane keeping system, the techniques have been able to autonomously navigate our testbed vehicle, the Navlab 5 [3], between lanes of a rural interstate highway.

Lane Model and Control Scheme

All of the techniques described in this paper use a very simple lane model which requires constant separation of the lane centers. This basically implies that the lanes are parallel and have a constant width.

Additionally, all of the techniques use a linear control scheme to determine the appropriate way to move to the next step of the lane transition maneuver. The linear scheme provides for controlled, comfortable lane transitions when used with the ALVINN lane keeping system.

Finally, in order to simplify integration of these geometric lane changing techniques, the output of the original ALVINN system has been slightly modified. Instead of being trained to produce the required turn radius to return the vehicle to the center of the driving lane, the system is trained to produce the displacement required to re-center the vehicle at a fixed lookahead distance in front of the vehicle. These formulations are identical in function, but allow for more direct manipulation of ALVINN's output.

Other Work

There has been a significant amount of research published describing how people change lanes as well as identifying theoretically optimal control strategies which could be used to autonomously control a vehicle in a lane change maneuver. Unfortunately, most of the researchers didn't have the facilities or equipment to test their results outside of the lab. The exception to this is [1] who have integrated lane transition functionality into their model based lane keeping system.

The Virtual Camera

An integral component in the lane transition techniques presented in this paper is the virtual camera. A virtual camera is simply an artificial imaging sensor which can be placed at any location and orientation in the world reference frame. It creates images using actual pixels imaged by a real camera that have been projected onto some world model. By knowing the location of both the actual and virtual camera, and by assuming a flat world model, accurate image reconstructions can be created from the virtual camera location. Virtual camera views from many orientations have been created and the images produced by these views have proven to be both accurate and usable by the ALVINN system to navigate successfully.

Lane Transition by Network Switching

Because ALVINN has the ability to re-center itself in the driving lane if it has made a significant error it is possible to transition from the current driving lane to the destination lane by simply switching to a network that was trained to drive in the destination lane. When presented with an image of the road while in the driving lane, this network would recognize that the vehicle was offset by a large amount from the proper driving position (in the destination lane) and would produce a sharp turn to re-center the vehicle. This type of lane transition is inadequate for two reasons. First, it is a very unsafe maneuver. At 25 meters/second, the lateral acceleration resulting from the sharpest turn ALVINN typically produces for highway driving is approximately 4.0 m/s^2. Even though the system cycles at 12-15 Hz, it could not recover quickly enough to prevent the vehicle from driving through the destination lane and off the other side of the road. Second, this type of maneuver would be uncomfortable and frightening to any passenger in the vehicle. A more controlled approach is desired in which the vehicle slowly transitions from one lane to the other, monitoring its performance along the way. To accomplish this goal, a modification of the network switching idea, along with two techniques based on controlled placement of virtual camera views, have been developed.

Incremental Network Switching

A method to transition lanes which is based on the simple network switching idea described above is called Incremental Network Switching (INS). Instead of immediately switching from the driving lane network to the destination lane network, this method combines the output displacements of the driving and destination lane networks to slowly transition the vehicle from the driving to the destination lane.

The INS method requires two trained networks, one for each lane. When a lane transition is initiated, road images are passed not only to the current driving lane network, but also to the network trained to drive in the destination lane. Each network responds by producing the displacement at the lookahead distance that would be required to re-center the vehicle in its respective lane.

Initially, the destination lane network's displacement indicates that a hard turn is necessary. But instead of using this displacement alone, the driving lane network's displacement is also utilized to compute a target displacement, which will determine how to steer the vehicle. This is done by incrementally adjusting how much contribution the displacement associated with the driving and destination lane networks have to the target driving displacement.

Initially, the target displacement is very similar to the driving lane displacement. This prevents the sharp transition into the destination lane which is present in the basic lane switching method. But because a small part of the destination lane displacement is used, the vehicle does begin to move toward the destination lane. In subsequent iterations, the amount of the displacement contributed by the destination lane increases and the amount contributed by the original driving lane decreases. When the amount contributed by the destination lane become significant, its output displacement is smaller because the vehicle is much closer to the center of the destination lane. This means that sudden, sharp movements are again averted. As the transition ends, the destination lane output displacement is the major contributor to the target driving displacement. Although the original driving lane network is producing a displacement which would cause the vehicle to return sharply to the original driving lane, its contribution is small and has only a limited effect. For these experiments, the lane transition was specified to be accomplished in 30 increments. This means that at each step, the contribution of the driving lane displacement was reduced by 1/30 while the contribution of the destination lane displacement was increased by 1/30.

Two constraints were used to ensure that both networks were performing correctly during the lane transition maneuver. The IRRE confidence measure [4] produced by each network was the first constraint. This measure is an indicator which is strongly correlated to satisfactory driving behavior. In order for an

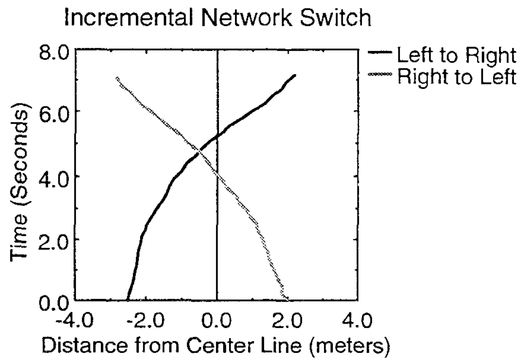

Figure 1. Average vehicle position during Incremental Network Switching Lane Transitions.

INS lane transition to proceed, the IRRE measure produced by either of the networks was required to be greater than a threshold value. Typically both networks produced satisfactory measures, with the measure increasing (or decreasing) as the vehicle moved into (or out of) the lane a network was trained to drive in. The second constraint was related to the geometric relevance of the network's output displacements. In general, the difference between the displacements should be equal to the lane separation distance, specified in the road model. The driving lane network should specify the center of the driving lane while the destination lane output specifies the center of the destination lane. In order for this constraint to be satisfied, the difference between displacements was required to be within 40% of the lane separation distance. This figure allowed for inadequacy in the network's output to represent all displacements while still prohibiting obviously false combinations of displacements.

Of the 38 lane transitions (19 in each direction) that were attempted, 37 succeeded. The results of these transitions are show in Figure 1. This figure shows the average vehicle position with respect to the center line of the road versus time. The most obvious feature in this graph is that the two trajectories do not cross at the center line of the road. Initially it was suspected that this was due to the vehicle not being properly centered during training of the left lane network. But as will be shown later, this assumption proved to be false. The exact cause of the anomaly is not known, but it is suspected that it is related to some small difference in how the left and right lane networks respond to images which produce large displacements. A somewhat subtler feature that is visible in the figure is the that the trajectories cross later in time than would be expected. If the transition was perfect, the trajectories would cross at about the 3.5 second mark, but using this method, they cross much later. This characteristic manifests itself during experiments by causing the vehicle to "dive" into the destination lane. Specifically, instead of moving smoothly into the destination lane, the vehicle's lateral velocity increases as the lane transition continues. This leads to overshoot in the destination lane after the maneuver has completed because of the somewhat long correction time required for smooth driving at high speeds.

Incremental View Lane Transition

The Incremental View Lane Transition (IVLT) method uses two virtual camera views to move the vehicle from the driving lane to the destination lane. It accomplishes this by intelligently controlling the placement of the views and using ALVINN's response to the images created from them. The central idea behind this technique is that virtual camera views can be used as a tool to check if the upcoming driving situation can be correctly handled by the system before actually encountering the situation. If it is determined that the system can handle the situation, virtual camera views can be used to cause ALVINN to respond in a controlled and direct manner.

The virtual camera views created for the IVLT method lead to images in which the vehicle seems to be improperly aligned with the driving lane. Because ALVINN's neural network is trained to recover from these situations, it responds with steering commands that attempt to re-center the vehicle. Once the vehicle has re-centered itself in the virtual image, the view is shifted again. As a side effect of the shifted views and resultant steering commands, the vehicle is incrementally moved through several lane transition locations and into the destination lane.

The IVLT method requires two networks be trained, one network for driving in the source lane and another for the destination lane. During normal operation before a lane transition has begun, only one virtual camera view, the Driving View, is used to keep the vehicle in its lane. When an IVLT is initiated an additional view, called the Test View, is created. The Test View is offset a small lateral distance from the Driving View.

Although the IVLT method requires two networks, only one is used to control the vehicle during the transition. While the system is using the image created from the Driving View to servo the steering wheel, it uses the Test View image to determine if the system will be able to reliably move the vehicle to and con-

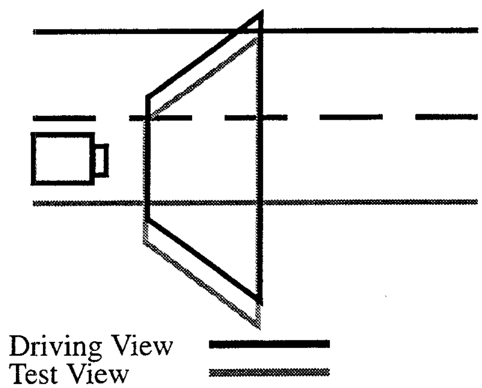

Figure 2. Driving and Test Views for a right to left lane transition.

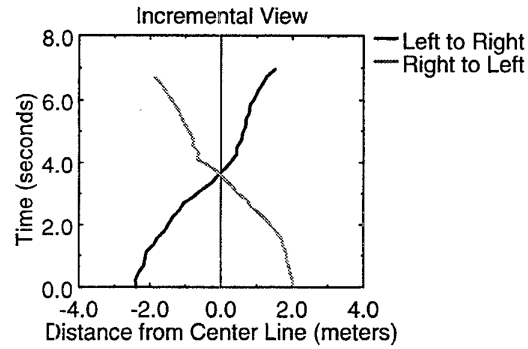

Figure 3. Average vehicle path during Incremental View Lane Transitions.

tinue control at the next transition location.

To determine if the next transition location is satisfactory, the Test View is used by the IVLT technique to compute both internal and external confidence metrics. These metrics can be used to determine if moving to the new lane location is advisable before actually executing the move. The IRRE confidence computed from ALVINN's output when presented with an image created from the Test View is used as an intrinsic measure of how well the network thinks it will be able to drive if the Test View is switched to the Driving View. Geometric knowledge about the lateral offset distance of the Test View from the Driving View, along with the output displacement of the Test and Driving Views, is used as the external metric. Because the system uses this geometric constraint in addition to the purely intrinsic IRRE measure, it is less susceptible to transitory features in the input image which can cause improper driving behavior yet not produce substantially lower IRRE values.

Before the Test View can be created, the appropriate lateral offset distance is computed. The lateral offset distance applied to the Test View is determined by the lane model as well as the desired number of intermediate lane transition positions. Because the view associated with each lane transition position must be used for a fixed period of time - while the system is determining if moving to the next location is advisable - the selection of the number of lane transition locations is proportional to the transition rate of the vehicle. Typical lane separation distances of 3.2 to 4.0 meters, along with 10-20 intermediate lane transition positions, lead to controlled and comfortable IVLT's.

If the vehicle is to transition into a lane to its left, the Test View is offset to the right of the Driving View. The opposite holds for lane transitions to the right. An example of initiating a right to left lane transition in shown in Figure 2. The network associated with the Test View is in most cases the same as the one associated with the current Driving View. Because the Test View is laterally offset, the image it creates looks as if it was taken from a vehicle that was not centered in the driving lane. Because of this, when the Test View becomes the Driving View after the internal and external reliability constraints have been met, ALVINN will compensate by steering the vehicle so that the road is centered in the image. The effect is to move the vehicle by the offset amount in the opposite direction of the offset.

Before actually moving from the current lane transition position, the IVLT method estimates ALVINN's ability to drive at the next location. It does this by monitoring the internal and external reliability metrics produced by ALVINN's neural network and the geometric constraints. To compute the external metric, the system computes the difference between the output displacement of the Test View network and the Driving View network. This difference is compared to the lateral offset distance of the Test View. If the difference between these figures is below some threshold (typically 25-40% of the Test View Offset), and the IRRE confidence is above another threshold (typically 0.75 out of 1.0) for a few (typically 2) Test View images, the system switches to using it as the Driving View and creates another Test View, again offset by the same distance from the new Driving View. This cycle repeats until the vehicle has reached the midway point of the lane transition. At this location, instead of using the same network as the Driving View, the Test View uses the network associated with the destination lane. Also at this step, the Test View offset distance from the Driving View is modified so that it correctly images the destination lane.

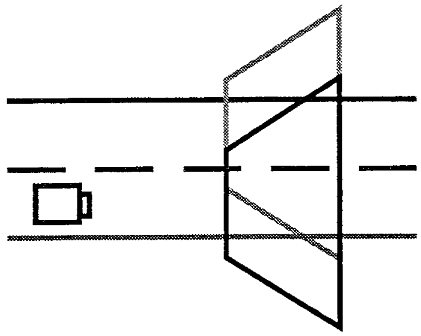

Figure 4. Initial virtual camera view placement for the Dual View Lane Transition technique.

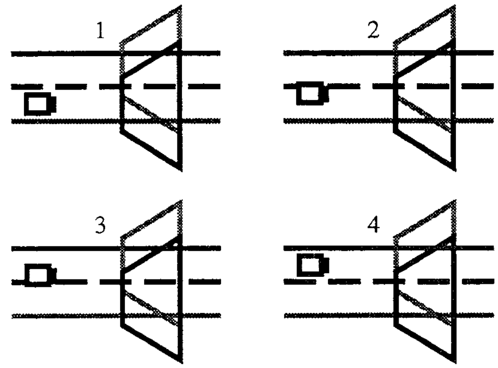

Figure 5. Progress of DVLT method. Note that as the vehicle moves, the views remain centered over the lane.

All 32 IVLT's that were attempted to verify this lane transition method were successfully completed. Like the INS method, half were left lane to right lane transitions. The same two networks that were used for the INS method experiments were used here. The results of these transitions are shown in Figure 3. There are three things to note about this graph. First, the average trajectory lines cross at the center line of the road. This is contrary to what occurred for the INS technique, which was thought to indicate improperly trained networks. Second, the trajectories are not as smooth as those in the INS method. This is a result of intermittent steering output when a Test View is being processed. The last characteristic to note is that the IVLT method relinquishes control before the vehicle is at a typical driving position in the destination lane. Initially, this was thought to be an algorithmic error, but upon closer inspection, we feel that it could be the result of a disagreement between the specified lane separation distance and the lane position which the networks were trained to drive at.

Dual View Lane Transition

Like the IVLT method, the Dual View Lane Transition method (DVLT) uses two virtual camera views to move the vehicle from the driving lane to the destination lane. While the IVLT method tracks only the original driving or destination lane to execute the lane transition, the DVLT method tracks them both. The idea behind the DVLT technique is to use a geometric road model along with accurate knowledge of the location of the center of each lane to smoothly servo the vehicle between lanes.

This technique is a bottom up approach to lane transition. Although it can use input from high level modules to initiate the lane transition, the geometry of the situation is what drives this method. The system locates the center of both lanes and then moves the vehicle based on these locations. This is in contrast to the IVLT method where ALVINN's response to the movement of virtual camera views is what causes the lane transition to proceed.

Like IVLT, the DVLT method requires networks to be trained for the driving and destination lanes. When a DVLT is initiated, a second virtual camera view is created. The road model is used to laterally offset this view, called the Destination Lane View (DLV) so that it is centered over the destination lane. The network that was trained to drive in the destination lane is associated with the DLV. Contrary to the IVLT, the DLV is not treated as a Test View, but rather as a tracking device. The driving, or Source Lane View (SLV) is treated in the same manner. See Figure 4.

After the DLV has been created, images are generated from both the DLV and the SLV and are passed to their respective network's, which produce an output displacement along with an IRRE confidence value. Both output displacements are converted to vehicle relative points. These points, called Lane Center Points (LCP's) specify where ALVINN believes the center of each lane is located. Because each network has been trained using the same lookahead distance, and because the DLV is only offset laterally from SLV, the LCP's fall on a line which is perpendicular to the direction of travel.

At this point, the LCP's are used to calculate the Modified Lookahead Point (MLP). The MLP is the point that the vehicle will actually drive toward. The MLP is between the two LCP's, along the line connecting them. The distance between the MLP and either of the LCP's is related to the step in the DVLT process. For

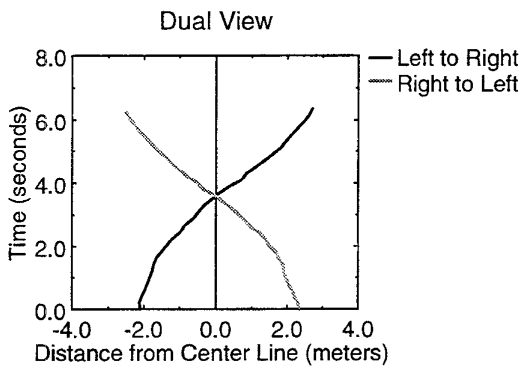

Figure 6. Average vehicle path during Dual View Lane Transitions.

example, in DVLT experiments which use 16 iterations to transition between lanes, the first MLP would be 1/16 of the total distance (along the line connecting the LCP's) away from the driving lane LCP and 15/16 from the destination LCP. The second MLP would be 1/8 and 7/8, respectively. Because the vehicle is instructed to move toward the MLP, the virtual views become misaligned with their respective lanes and must be updated. This process is similar to the one used for the IVLT method except that both views are changed. If the vehicle is moving to the left, both views are updated by moving them to the right. This process is continued until the vehicle has transitioned completely into the destination lane. See Figure 5. The combination of moving toward a MLP which is continually closer to the destination lane center and of updating the location of the virtual camera views results in a smooth, controlled transition.

During the lane transition, the IRRE confidence metric and the constant lane separation constraint are used to determine if the system is confident in its current driving ability. This is done by checking the IRRE confidence value associated with the images created by the views and by comparing the distance between the two LCP's. For the transition to continue, the IRRE confidence measure is required to be above a threshold value, typically 0.75, and the distance between the LCP's is required to be within 40% of the lane separation distance, specified by the lane model. Unlike the IVLT, this method does not check to see if the system will be able to drive at the next lane transition position.

All 42 DVLT's that were attempted to verify this lane transition method were successfully completed. Like the other methods, half were left lane to right lane transitions while the rest were right to left. The results of the experiments are shown in Figure 6.

There are two characteristics to notice in this figure. First, the lane transitions are symmetric across the center line of the road as well as with respect to the time in the lane transition. Second, the DVLT method doesn't relinquish control until the vehicle is at the proper driving position in the destination lane. Empirically, this equates to a smoother shift into the destination lane.

Conclusions and Future Work

These experiments have confirmed our belief that by intelligently controlling both the input to and output of ALVINN, tactical level driving tasks can be achieved.

Future work will include providing a high level interface to the lane changing functionality for tactical level driving modules and integrating the methods with an obstacle avoidance system, so that avoidance maneuvers, such as changing lanes and swerving, can be automated.

Acknowledgments

The authors would like to thank Delco Electronics for donating the Navlab 5 vehicle and the Pennsylvania Turnpike Commission for making its STAR Site available to us as a test facility. The United States Department of Transportation, though National Highway Traffic Safety Administration (NHTSA) contract DTNH22-93-C-07023 and National Automated Highway System Consortium, (NAHSC) contract DTFH61-94-X-00001, has provided funding for this research.

References

[1] Dickmanns, E. D., Behringer, R., Dickmanns D., Hildebrandt, T., Maurer, M., Thomanek, F., and Schielen, J., "The seeing passenger car 'VaMoRs-P,'" *1994 IEEE Symposium on Intelligent Vehicles*, pp. 68-73.

[2] Jochem, Todd M. "Using Virtual Active Vision Tools to Improve Autonomous Driving Tasks," CMU Technical Report CMU-RI-TR-94-39, October, 1994.

[3] Jochem, Todd M., Pomerleau, Dean A., Kumar, Bala and Armstrong, Jeremy. "PANS: A Portable Navigation Platform," *1995 IEEE Symposium on Intelligent Vehicles*, September 25-26, 1995, Detroit, Michigan, USA.

[4] Pomerleau, D. A. Neural Network Perception for Mobile Robot Guidance, Kluwer Academic Publishing.