Chapter 9

Real World Driving

9.1 Data collection

The real world driving data were collected using the CMU Navlab 8 test vehicle, shown in Figure 9-1 [Pomerleau et al, 96]. A CCD camera is mounted on the windshield, underneath the rear-view mirror. This camera is used for lane tracking and vision based obstacle detection. A radar obstacle sensor is mounted behind the front license plate, and is used for detecting vehicles directly ahead and to the front-left/right. Two side sensors are mounted on the sides of the vehicles, near the rear. A single line laser range finder is mounted behind the rear bumper. It also has a Differential Global Positioning System receiver, which has a resolution of +/- 3-5m. Finally, a yaw-rate gyro is mounted in the rear, along with a tilt sensor. Hence, this vehicle allows us to take the time series records of the vehicle's states, the environmental situation, as well as the control actions. Notice that currently there is no sensor to measure the throttle of the engine in NavLab 8.

After eliminating not-important ones using our prior domain knowledge, the variables listed in Table 9-1 were used for the detection experiments. The shaded variable in the table, steering angle (ϖ), is the only output variable; the other output variable, the throttle of the gas in con-

Figure 9-1: NavLab's smart van. (Courtesy of Navlab, CMU).

junction with the brake force, is absent. All others, including the previous records of ϖ, were used as inputs. All the variables were taken record at a frequency between 14 Hz and 18 Hz.

Seven people were invited to drive the vehicle. They were selected from both genders and over a range of ages from twenty to fifty. All of them have valid U.S. driver's licenses, and have at least four years driving experience in the U.S., with no major traffic violations, accidents, or DUIs. The subjects were told only that we were interested in learning driving behaviors. Details were kept sketchy, to help avoid biasing the drivers' behaviors. They were not told how to drive, but the only instruction was to drive safely.

The operators were asked to drive from CMU to Grove City, a small town about 50 miles north of Pittsburgh, then back. "The route is primarily two lane (in each direction) highway driving, with short stretches of three lanes." Each operator drove for over two hours round trip. "One concern is that the subjects most likely have never driven a Silhouette, or even a mini-van. A mini-van is large enough that it is hard to get a good feel for the boundaries and available space, particularly on the right hand side. Due to this, most drivers initially tended to hug the left side of the road. However, this effect seems to subside within a half hour or so of driving."

Table 9-1: Real world driving variables

Variables	*Description*	*Variables*	*Description*
x_ξ	The lateral position	$1/s_F$	Inv. distance to the front obstacle
v_ξ	The lateral velocity	$1/s_{FL}$	Inv. dist. to the front-left obstacle
v_η	The longitudinal velocity	$1/s_{FR}$	Inv. dist. to the front-right obstacle
θ	Road Curvature	$1/s_B$	Inv. distance to the back obstacle
ϕ	Vehicle yaw	$1/s_{BL}$	Inv. dist. to the back-left obstacle
ϖ	Steering angle	$1/s_{BR}$	Inv. dist. to the back-right obstacle

Unlike the simulation cases discussed in last chapter, it seems to us that linear models are not appropriate for describing the real world driving behavior, because of the existence of traffic. For example, most drivers tend to take cut to the inside on a curvy road if there is no traffic, as illustrated by the dash curve in Figure 9-2. However, in case there is traffic, especially if there are other vehicles in the shortcut route, the drivers are more likely to stay in the middle of the lane. We can measure the distance from our vehicle to other vehicles in the curve, such as "d" in Figure 9-2. If there is no traffic in the curve, d goes to infinity or $1/d$ is equal to zero. To decide to take the shortcut, the crucial issue is that $1/d$ should be zero, however, it does not matter that $1/d$ is equal to 0.25 or 0.32. Therefore, it is not proper to model the relationship among the vehicle's lateral position, the road curvature and $1/d$ as a linear function.

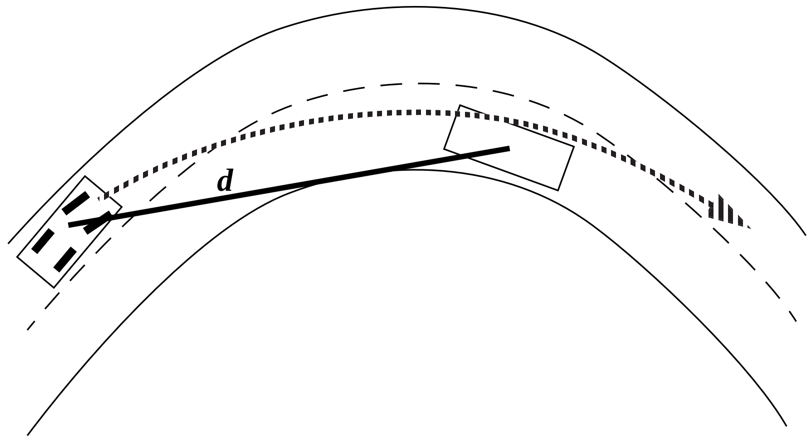

Figure 9-2: Driving in traffic may be non-linear. If there is no traffic, a driver tends to take a shortcut. Otherwise, he may stick to the same lane.

Based on empirical analysis, we found three seconds' time delay was sufficient for OMEGA to work properly. Hence, for each variable, we took its previous forty-eight $(3 \times 16(\mathrm{Hz}))$ records into account, except that for the road curvature, we took its forty-eight records ahead. For each variable, we used PCA to compress its dimensionality from forty-eight to three. Then we combined the twelve variables together, and used PCA again to reduce the dimensionality from thirty-six (12×3) to eight. The compression of the dimensionality is to make the further classification process feasible; however, as the price, we lost 17.8% information.

9.2 OMEGA result

Since there were seven drivers, and each one had two datasets, from Pittsburgh to Grove City and back, so that there were totally fourteen datasets: O_{ij}, $i = 1, ..., 7, j = 1, 2$. We can randomly select one dataset as a testing dataset, hide the real driver to OMEGA, and ask it to detect the driver to see if OMEGA is capable of detecting correctly.

To do so, OMEGA needs some training datasets. Define the datasets such that O_{ij} corresponds to journey *j* by driver *i*. If the testing dataset corresponds to a trip from Pittsburgh to Grove City, say O_{31}, where *1* refers to the route, *3* indicates the real driver. We assign the datasets collected

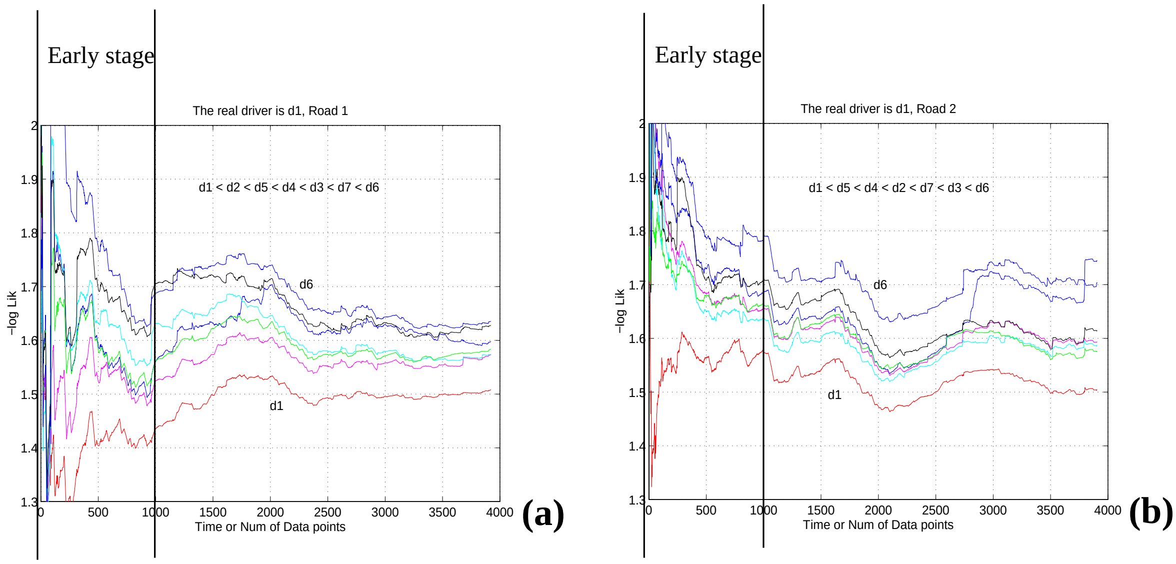

Figure 9-3: OMEGA detects the real world driving style. Two correct cases. (a) From Pittsburgh to Grove City, (b) From Grove City back to Pittsburgh.

on the way back from Grove City to Pittsburgh, as the training datasets. Thus, for each testing dataset, we have seven training datasets. For example, if the testing dataset is O_{31}, the training datasets will be O_{k2}, $k = 1, ..., 7$.

Since we can assign any dataset to be the testing dataset, totally we can do fourteen detection experiments. OMEGA succeeded in ten cases, failed three times and was confused once[1].

Referring to Figure 9-3, at the early stage of the detections, due to the insufficient number of data points involved in the analysis, the likelihood curves are unstable. With more and more data, the curves converge eventually. However, overall the curves look bumpier than those of the simulation experiments discussed in Chapter 8, referring to Figure 8-3. There are four possible reasons: (1) The real world datasets may be noisier than the simulation datasets because of the resolutions of the sensors. (2) We lost 17.8% information when we did the PCA pre-processing. (3) One of the two output variables, the throttle of the gas/brake is absent. (4) Although

1. Again, we assigned the significance level α to be 5%, referring to Chapter 2.

the real world driving datasets are large in size, the majority of their contents consist of nothing but very routine operations which are not helpful for distinguishing different people's driving styles.

Both Figure 9-3 (a) and (b) were generated by the same driver, "d1". Figure 9-3 (a) corresponds to the trip from Pittsburgh to Grove City, and Figure 9-3(b) corresponds to the way back. Comparing the early stages of Figure 9-3 (a) and (b), we notice that the curves in (a) were more chaotic than (b)'s. As a matter of fact, we observed the same phenomena happened to almost all the drivers, in other words, all drivers' initial performance were not so well-controlled as afterwards. In Table 9-2, we compare the standard deviations of the log likelihood of each driver's performance at the early stages of the trips from Pittsburgh to Grove City, with their counterparts on the ways back. Obviously, most operator's initial performance was significantly more disordered than the latter one, except that "d5" seems more ready to drive from the very beginning. These phenomena are supported by the observation mentioned in Section 9-1: "One concern is that the subjects most likely have never driven a Silhouette, or even a mini-van. ... However, this effect seems to subside within a half hour or so of driving."

Table 9-2: Standard deviations of the likelihood at the early stages.

	D1	*D2*	*D3*	*D4*	*D5*	*D6*	*D7*
Pgh - Grove	0.162	0.232	0.215	0.161	0.175	0.182	0.192
Grove - Pgh	0.120	0.095	0.159	0.118	0.163	0.053	0.047

As usual, the curve whose tail is the lowest indicates who is the real driver. In the correct cases, as in the examples of Figure 9-3, the lowest curves are underneath the others by large margins. Figure 9-4 (a) and (b) are examples of the confused cases and the incorrect ones. In fact, all the wrong cases are similar to Figure 9-4 (b): Although the real driver's curve is not the lowest one, it is lower than most others. That is to say, although OMEGA may make mistakes, the correct one is usually within the attention scope.

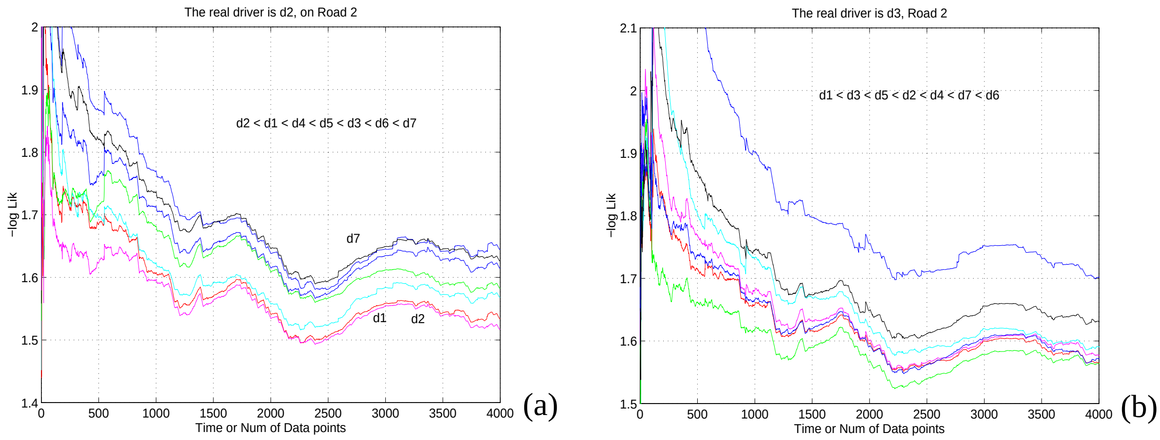

Figure 9-4: (a) A confused case. (b) A wrong case. Even as a wrong case, the real driver's curve is close to the lowest one.

9.3 Comparison with other methods

Table 9-3 is the comparison of OMEGA with other methods.

Table 9-3: Comparison of OMEGA with other methods

	Correct	*Wrong*	*Confused*
Global linear	4	7	3
HMM[a]	4	0	3
OMEGA	10	3	1

a. Nechyba has done only half of the experiments.

As we expected, the linear approach does not work properly due to the reason we discussed at the end of Section 9.1.

[Nechyba, 98, (a)]'s method did work in this domain. However, unlike OMEGA which separated the real driver from the others with a salient margin in log likelihood, [Nechyba, 98, (a)] could not make a decisive detection between two or more candidates.

Table 9-4: Cross-validation of OMEGA.[a]

	D1	*D2*	*D3*	*D4*	*D5*	*D6*	*D7*
D1	0.346	0.142	0.031	0.150	0.169	0.047	0.115
D2	0.213	0.247	0.096	0.155	0.129	0.074	0.086
D3	0.182	0.167	0.176	0.144	0.182	0.048	0.102
D4	0.159	0.119	0.059	0.337	0.126	0.087	0.113
D5	0.156	0.105	0.077	0.098	0.435	0.034	0.094
D6	0.124	0.161	0.123	0.124	0.185	0.174	0.108
D7	0.154	0.120	0.065	0.100	0.199	0.050	0.312

a. Using the datasets collected on the way from Pgh to Grove city as the training dataset, and using the datasets collected on the way back as the testing datasets.

While Table 9-3 gives a top-level comparison, Table 9-4 and 9-5 view the precision in depth. Each number in the tables is a probability of a testing dataset being generated by a certain operator. Each row corresponds to a specific testing data set, and the real operator is in the leftmost column. The other columns represent the seven candidate drivers. The testing datasets of Table 9-5 were collected on the way from Pittsburgh to Grove City, and the training datasets were collected on the way back. To be fair, so did those for Table 9-4. The number in the (2,3)'th cell is the probability that a testing dataset, which was secretly generated by the second driver, would be detected as the performance of the third driver. Thus, the sum of the seven probability values in each row is always 1.0. The number on the shaded diagonal is expected to be bigger than the others. And the bigger the diagonal number is, the better the detection system performs. Otherwise, the detection fails. We used 0.030 as a threshold to judge if the probabilities on the diagonal are significantly bigger than all the other six probabilities in the row. We notice in Table 9-4, OMEGA made wrong decisions twice. But when OMEGA made correct decisions, it was quite decisive. Conversely, HMM did not make any wrong decision, but when it came to the correct conclusion, for three times, the numbers on the diagonal could not be separated from the other six numbers in the rows by the 0.030 threshold.

Table 9-5: Cross-validation of HMM, [a]

	D1	*D2*	*D3*	*D4*	*D5*	*D6*	*D7*
D1	0.359	0.309	0.066	0.113	0.040	0.037	0.076
D2	0.108	0.226	0.123	0.193	0.098	0.090	0.162
D3	0.055	0.159	0.243	0.126	0.202	0.124	0.092
D4	0.106	0.196	0.102	0.216	0.097	0.123	0.160
D5	0.180	0.164	0.174	0.089	0.207	0.134	0.052
D6	0.053	0.127	0.087	0.208	0.105	0.232	0.188
D7	0.041	0.149	0.056	0.244	0.058	0.161	0.291

a. The same as the footnote of Table 9-4.

9.4 Summary

This chapter demonstrated that OMEGA is capable of detecting different systems accurately even in a complicated domain, where the conventional linear system identification approach, is not functional any more.