

RACCOON:

A Real-time Autonomous Car Chaser Operating Optimally at Night

Rahul Sukthankar
Robotics Institute
Carnegie Mellon University
5000 Forbes Ave., Pittsburgh, PA 15213-3891
Phone: (1)412-268-3083, Internet: rahuls@ri.cmu.edu

Abstract

RACCOON is a vision-based system that tracks car taillights at night. It builds a global map in real time of the lead vehicle's position based on the location and separation of the taillights in a sequence of video images. RACCOON has been integrated into a car following experiment on the CMU Navlab II, a computer-controlled HMMWV testbed. The Navlab II safely followed a lead vehicle on a winding road in light traffic at 32 km/h.

1. Introduction

Although vision based autonomous navigation has been an active research area for a number of years [1, 2, 3, 9], car tracking [5, 8] and car following [7] have only been addressed recently. Vision based car following in low light conditions has remained unexplored. At night, an autonomous land vehicle must deal with an additional set of problems including:

- The road cannot be seen clearly at night.
- Traffic looks like a pattern of bright lights on a black background.
- Unlit landmarks cannot be detected so corners and intersections have to be negotiated based solely on the observed actions of the lead vehicle.

Some of these problems can be solved by developing a system which follows a human controlled lead vehicle. Since taillights can be easily extracted from a dark background, an intuitive approach to the car following problem is to steer the autonomous vehicle so that it heads towards the taillights of the lead vehicle. This implementation may produce satisfactory results on straight roads when both vehicles are moving at the same speed. However this naive algorithm fails in any realistic scenario since lead vehicles make turns to follow winding roads, and steering

towards taillights results in corner cutting — possibly causing an accident as the computer controlled vehicle drifts into oncoming traffic or off the road entirely.

RACCOON solves these problems by creating an intermediate map structure which records the lead vehicle's trajectory. The path is represented by points in a global reference frame, and the computer controlled vehicle is steered from point to point. The autonomous vehicle follows this trail at any desired speed, while keeping the lead vehicle's taillights in sight. Using this approach, the autonomous vehicle steers around corners and obstacles rather than through them. RACCOON has been successfully implemented on the Carnegie Mellon Navlab II testbed and follows a variety of lead vehicles at night, on real roads, in light traffic. Its only inputs consist of image sequences from a single color camera and Navlab II position information from the onboard controller.

2. System Architecture

RACCOON consists of the following subsystems: (see Figure 1)

1. Initialization
2. Image Acquisition
3. Taillight Tracking
4. Lead Vehicle Position Calculation
5. Consistency Checking & Error Recovery
6. Output

Each of these will be individually discussed below. RACCOON can be used in two modes:

1. Passive mode: stand-alone passive car tracking system.
2. Active mode: sends output to path planner module for car following.

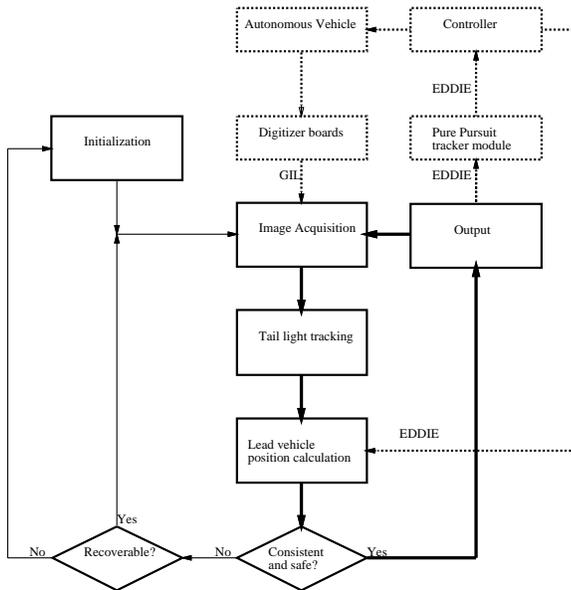


Figure 1: System Architecture

The former mode enables RACCOON to be tested in the lab on video taped image sequences, while the latter may be used either to reconstruct the path of the lead vehicle or to actively follow the lead vehicle on a real road.

2.1. Initialization

A video image of a night-time driving situation consists primarily of a black background punctuated by numerous bright spots and occasional specular reflections from metallic surfaces. The lead vehicle's position must be extracted from images containing a large number of taillights. RACCOON addresses this task by only considering features which are contained within a rectangular box in the image. By ignoring lights which lie outside this region, processing speed is increased and the possibility of tracking spurious features is decreased. The following assumptions are made about the lead vehicle:

- The taillights on the lead vehicle are red, and bright enough to be distinguished from the background.
- The distance between the outermost taillights is known and fixed.
- The taillights of the lead vehicle will only move small distances from image to image.

Since the patterns of taillights on vehicles can vary tremendously, RACCOON makes no assumptions about the number of lights, their relative positions, or their height from the road surface. Any red lights inside the region of interest are considered to be part

of the lead vehicle, and thus RACCOON robustly handles situations like:

- Illumination of turn signals or brake lights.
- Variations in relative camera pitch due to rough roads or slopes.
- Specular reflections from red cars.

The size of the region of interest shrinks and grows as the lead vehicle changes distance relative to the autonomous vehicle. The increase in flexibility provided by this approach is partially offset by the danger of tracking spurious features which appear inside the region of interest. The consistency checking module attempts to recognize and correct for these rare occurrences.

Choosing which vehicle to follow is in itself a difficult task, so the current implementation of RACCOON allows the user to interactively select the initial region of interest by dragging a rectangle around the taillights of the lead vehicle in the initial image. Manual re-initialization of the system is only necessary if the lead vehicle has been irretrievably lost, or if a different lead vehicle is to be followed. By default, the lead vehicle is expected to have an approximate width of 1.5 meters, but this can be changed at run time if the lead vehicle has unusual dimensions.

2.2. Image Acquisition

RACCOON receives video input as a sequence of digitized images and operates identically on live camera input, videotaped data or images stored on disk in the CMU Generalized Image Library format [6]. The rectangular region of interest is pre-processed to yield normalized red intensities at every pixel using the formula:

$$r_N = \frac{R}{R + G + B} \quad (1)$$

Thresholding is then performed to keep only the pixels which satisfy:

- Absolute red intensity $> T_A$
- Normalized red intensity $> T_B$

where T_A and T_B are thresholds specified at run-time. This selection process ensures that only pixels which are bright enough and red enough to be taillights are considered by the subsequent tracker. The first threshold eliminates light sources which are too dim to be taillights. The second threshold measures the relative amount of red content per pixel, and filters out bright lights which are not sufficiently red, such as head lights and street lights. Any part of the image which lies outside the region of interest is ignored.

2.3. Taillight Tracking

Within the given region of interest, the centroid of taillight pixels and their bounding box are calculated. The new region of interest is computed by expanding the bounding box by a fixed amount. This straightforward approach is significantly faster than more complicated algorithms, and has proven to be adequate both in the lab and on the road. A model which predicts the motion of the region of interest from image to image was also implemented, but found to be unnecessary.

The distance to the lead vehicle and its bearing from the autonomous vehicle is calculated from the horizontal position of the centroid and the bounding box size as follows:

distance: Assuming that the taillights of the lead vehicle are enclosed in the bounding box, the horizontal size of the box, S is related to the distance to the lead vehicle, D by the formula:

$$D = k/S \quad (2)$$

where the constant of proportionality k is pre-computed from the known camera parameters as shown below. The vertical size of the bounding box is expected to change drastically with the illumination of the third brake light and is therefore ignored.

bearing: The horizontal position of the centroid of the taillights in the region of interest is a good indicator of the angle between the lead vehicle's position and the autonomous vehicle's current heading. When combined with the distance information, the lateral offset of the centroid l from the camera centerline enables us to compute the vector to the lead vehicle.

The relationship between measurements in camera and image coordinates can be deduced from Figure 2. The following assumptions are implicit in the equations:

- The field of view (FOV) or the focal length of the camera is known.
- The camera is not rolled relative to the lead vehicle, so that width of the bounding box maps to the width of the car in the given image.
- The yaw of the camera is known, and relatively small so that any foreshortening in the lead vehicle's image is negligible.
- The lead vehicle is pointing approximately in the same direction as the autonomous vehicle so that foreshortening in the bounding box is negligible.

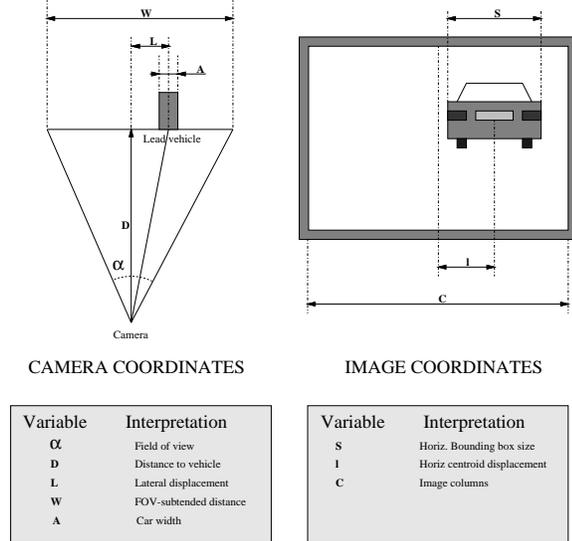


Figure 2: Recovering vehicle position vector from camera image

From our model, we thus get the following geometric relationships:

$$W = C \times \frac{A}{S} \quad (3)$$

$$L = W \times \frac{l}{C} = l \times \frac{A}{S} \quad (4)$$

$$D = \frac{W}{2 \tan \alpha/2} \quad (5)$$

where L is the lateral displacement of the lead vehicle and D is its distance from the camera. The intermediate variable W is the distance subtended by the camera's FOV at the distance D .

Given the vector (L, D) to the lead vehicle in the camera's reference frame, we can calculate its position in robot vehicle coordinates since we know the position and orientation of the camera relative to the computer controlled vehicle.

2.4. Lead Vehicle Position Calculation

This module is given the position of the lead vehicle in the coordinate system fixed to the robot vehicle's rear axle. By querying the controller, the position and orientation of the autonomous vehicle at the time of frame capture is acquired. If RACCOON is being tested in a lab, this is usually provided by a vehicle simulator; on the Navlab II testbed, the position is computed by a combination of dead-reckoning and inertial navigation systems. Using this information, the module calculates the position of the lead vehicle in global coordinates. If this position is found to be consistent with past information, the output module will pass this position on to other systems.

2.5. Error Recovery

Occasionally, the taillight tracker encounters misleading inputs and will therefore report an incorrect position for the lead vehicle. This can occur when a spurious red light is seen in the region of interest, an obstacle occludes a taillight or when the lead vehicle temporarily leaves the camera's field of view.

To detect potential problems, RACCOON employs the following heuristic: The given lead vehicle position is compared with the latest confirmed position in world coordinates, and if the Euclidean distance between them is within a threshold, the new position is classified as being *sane*. Once a number of consecutive sane positions have been seen, they are additionally tagged as *safe* and RACCOON will send these positions to the planner via the output module. A lone bad position evokes no special response, but multiple consecutive bad positions increase the region of interest to try to reacquire the lead vehicle. Before RACCOON decides to reinitialize, it tries the following strategies:

1. Searching a uniformly grown region of interest at reduced resolution.
2. Checking a large horizontal strip in the image, in case the lead vehicle turned sharply.
3. If the bounding box is adjacent to the horizontal image boundaries, RACCOON advises steering in that direction to prevent losing sight of the lead vehicle completely. This ensures that images with clipped bounding boxes are recognized early to increase the chance of a successful error recovery.

If all of these approaches fail, RACCOON is forced to reinitialize: either manually, by allowing the user to select a new car to follow, or autonomously by exhaustively searching the image for some red lights to track. The current implementation is unable to make an intelligent decision as to which set of taillights belong to the desired lead vehicle.

2.6. Output

RACCOON, in its stand-alone mode supports a graphical interface in addition to diagnostic text output. A green box is displayed surrounding the region of interest to help the user determine whether or not the lead vehicle is being tracked correctly (See Figure 3). During error recovery, the potential search regions are highlighted on the display.

RACCOON, when activated as a module in a larger system, communicates with other processes using EDDIE [11], a message passing system developed at Carnegie Mellon for the Navlab project. The modules with which RACCOON interacts include:



Figure 3: RACCOON display while tracking lead vehicle

- Vehicle simulator: used for testing the system in the lab.
- Lead vehicle path mapping: graphical display of RACCOON output.
- Path planner: smoothly connects the given points and passes steering radius and desired velocity to the controller.
- Controller: drives the robot and updates current robot position.

3. Results

RACCOON was first developed in its passive mode and tested on videotaped images collected from the Navlab II camera during a human driven run. It consistently tracked the lead vehicle throughout the tape, detecting inconsistencies and recapturing the vehicle through whenever necessary. On a Sun-3 equipped with frame-grabber boards, a cycle rate of 3-5 Hz was achieved in debug mode. The time taken to process a given image was observed to be approximately proportional to the number of pixels in the region of interest, thus inversely proportional to the distance from the camera to the lead vehicle. The speed improvement for distant vehicles is balanced by a corresponding loss of accuracy since each pixel in the image corresponds to a larger region in the world when the vehicle is farther away.

Once satisfactory results had been obtained in the lab, RACCOON was ported to the Navlab II testbed. The goal was to follow a car-sized lead vehicle with unspecified taillight patterns, moving at a varying speed along a winding road. On a Sun Sparc-2 architecture, the processing time per

image was substantially decreased and RACCOON was able to process most images at 15 Hz. Since the Navlab II's control bandwidth is only 10 Hz, RACCOON's speed was found to be more than adequate.

The reported lead vehicle positions were connected in world coordinates by a tracker module using a pure pursuit algorithm [12]. This algorithm steered the robot so that its trajectory would intersect the target path at a pre-specified look-ahead distance. In general, smaller values of the look-ahead parameter result in increased sensitivity, possibly leading to instabilities, while larger values give smooth paths, at the expense of rounding sharp corners. For the given scenario, an empirically determined value of 15 meters resulted in stable behavior.

Initial tests in a parking lot were promising, but highlighted the potential problem of losing the lead vehicle during tight turns as it temporarily left the field of view. RACCOON was modified to correctly handle situations where the taillights of the lead vehicle are partially clipped. Difficulties with the field of view may be addressed in future versions by employing an orientable camera.

RACCOON was subsequently tested on Schenley Drive, a road which offers hills, bumps and S curves in addition to light traffic. The Navlab II safely followed the lead vehicle at approximately 20 mph (32 km/h) in light conditions ranging from dusk to complete darkness. Since RACCOON does not use a flat earth hypothesis, it is able to handle variations in the vertical position of the taillights caused by the hills in the road. RACCOON was also tested in light rain, and it was observed that the red reflections due to water droplets on the lens did not degrade performance since they fell below the brightness threshold of the taillight tracking module.

One of the major advantages of RACCOON's strategy is that the computer controlled vehicle need not maintain a constant distance to the lead vehicle — the robot can at its own speed. Furthermore, even if velocity control is not accurate, performance is not affected since the target path remains valid. RACCOON can also operate in a decoupled mode, where a human driver operates the throttle, while steering is performed under computer control.

4. Conclusions and Future Extensions

RACCOON is a system suited for a specific domain (night time car following) and it performs well in those conditions. However, RACCOON can be integrated into existing systems as a module for a number of other applications:

- Convoying operations.
- Intelligent cruise control on highways.

- Steering robot vehicle when the road is ambiguous or not visible.
- Safer driving in traffic, especially through intersections.

Planned extensions to this project are motivated chiefly by shortcomings of the existing system. The biggest problem with the current implementation is that the single, fixed camera has a restricted field of view (42.5 degrees) and this is found to be inadequate when the lead vehicle turns sharply and the Navlab II is following closely (for example, at an intersection). Losing sight of the vehicle is unavoidable and this leads to problems with reacquiring the lead vehicle on completion of the turn since there may be other cars in the scene. One solution to this problem is to mount the camera on a pan/tilt platform that can be oriented to follow the lead vehicle during tight turns.

RACCOON could easily be modified to actively servo of the camera to maintain the lead vehicle in the centre of its field of view whenever possible. RACCOON could return the distance and bearing in sensor coordinates and the result would be transformed using Panacea [10] to recover the position of the lead vehicle in the robot's frame of reference.

A second issue involves car following in daytime situations. Since RACCOON tracks illuminated taillights, the system can only function reliably in low light conditions such as dawn, dusk or at night. Alternate approaches such as optical flow or template matching could be used to extract the position of the lead car in daylight conditions without modifying the strategy used in following. Although these algorithms are unacceptably slow on the existing hardware, they may become more feasible on parallel architectures in the future.

Real-time perception systems are forced to make certain concessions in robustness in order to achieve desired speeds. Currently RACCOON relies on other obstacle avoidance systems (or user intervention) in order to avoid collisions with unlit obstacles. Also, due to hardware limitations, RACCOON only tracks a single vehicle at any one time and this is insufficient in heavy traffic situations. RACCOON has been conceived with these issues in mind however, and future versions will reason intelligently about multiple moving objects in the robot's environment.

Acknowledgements

The author wishes to thank the following people for their help with the project:

- Dr. Dean Pomerleau: who gave invaluable assistance, both in the lab and on the road.

- Dr. Charles Thorpe: for his guidance and help, especially with the initial algorithm.
- Jay Gowdy: for writing EDDIE and the pure-pursuit tracker.
- Ian Davis, Todd Jochem, Conrad Poelman and Subhash Sukthankar: for driving the lead vehicles during testing.
- The CMU UGV group: for hardware and software support.

This research was partly sponsored by DARPA, under contracts "Perception for Outdoor Navigation" (contract number DACA76-89-C-0014, monitored by the US Army Topographic Engineering Center) and "Unmanned Ground Vehicle System" (contract number DAAE07-90-C-R059, monitored by TACOM).

References

- [1] Aubert, D., Kluge, K., and Thorpe, C. (1990): Autonomous Navigation of Structured City Roads *Proceedings SPIE Mobile Robots V*
- [2] Crisman, J. and Thorpe, C. (1990): Color Vision for Road Following. *Vision and Navigation: The CMU Navlab*, C. Thorpe (Ed.), Kluwer Academic Publishers, Boston.
- [3] Dickmanns, E. and Zapp, A. (1987): Autonomous high speed road vehicle guidance by computer vision. *Proceedings of the 10th World Congress on Automatic Control, Vol. 4*, Munich, West Germany.
- [4] Gage, D. W.; Pletta, J. B. (1987): Ground Vehicle Convoying. *SPIE Vol. 852 Mobile Robots II*, pp 319-327
- [5] Graefe, V. and Jacobs, U. (1991): Detection of Passing Vehicles by a Robot Car Driver. *IROS '91*, Osaka, pp 391-396.
- [6] Hamey, L., Printz, H., Reece, D., Hyre R., and Shafer, S. (1990): A Programmer's Guide to the Generalized Image Library. Computer Science Department documentation, Carnegie Mellon University.
- [7] Kehtarnavaz, N., Griswold, N. and Lee, S. (1991): Visual Control of an Autonomous Vehicle (BART) — The Vehicle-Following Problem. *IEEE Transactions on Vehicular Technology, Vol. 40, No. 3*
- [8] Kuehnle, A (1991): Symmetry-based Recognition of Vehicle Rears *Pattern Recognition Letters 12*, pp 249-258
- [9] Pomerleau, D. (1990): Neural Network Based Autonomous Navigation. *Vision and Navigation: The CMU Navlab*, C. Thorpe (Ed.), Kluwer Academic Publishers, Boston.
- [10] Sukthankar, R., Pomerleau, D. & Thorpe, C. (1993): Panacea: An Active Sensor Controller for the ALVINN Autonomous Driving System. CMU RI Technical Report in preparation.
- [11] Thorpe, C. and Gowdy, J. (1991): Annotated Maps for Autonomous Land Vehicles *Proceedings of the Second International Workshop on High Precision Navigation*, Stuttgart, Germany.
- [12] Wallace, R., Stentz, A., Thorpe, C., Moravec, H., Whittaker, W., and Kanade, T. (1985): First Results in Robot Road-Following. *Proc. IJCAI-85*