DATA FUSION FOR SURROUND SENSING

Chuck Thorpe, Christoph Mertz, David Duggins, Jay Gowdy, and Arne Suppe
Carnegie Mellon University Robotics Institute
Pittsburgh, PA 15213

One of the most important technologies that must be perfected before robotic vehicles become practical is 360 degree safeguarding. It is critical that a vehicle be able to detect and avoid objects around its perimeter, both for the safety of the vehicle itself, and, more importantly, for the safety of surrounding objects and people. If a navigation system fails occasionally, the consequences are that the vehicle may be temporarily lost; but if the safeguard system fails, the vehicle may be destroyed or people may be killed.

The problem of 360 degree safeguarding is not the large coverage area, since one can always mount more sensors on a vehicle. Instead, it is making different kinds of sensors work together to solve particular tasks which each individual sensor can not solve by itself. Sensors need to be fused at different levels to produce a robust safeguarding system.

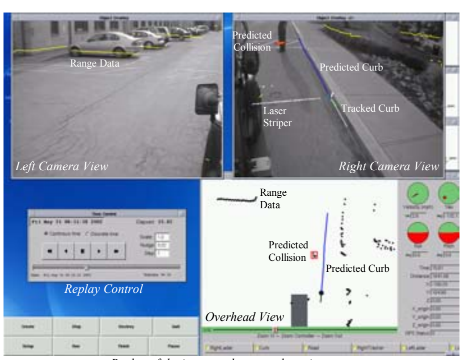

Replay of the integrated surround sensing system

The Robotics Institute at Carnegie Mellon University is building a 360 degree surround safeguard system for both robotic and human-controlled vehicles. The various sensors must be properly characterized, calibrated individually and with respect to each other, fused, and interpreted in order to build up the required understanding of the objects in the immediate robot environment. In this paper we will be discussing three different kinds of fusion: Sensor fusion over time, fusion between different sensors, and information fusion. Above is an image of our system in replay mode which illustrates the different sensors and their fusion.

Using a moving vehicle as a sensor mounting platform, the sensor reading will vary over time and reflect the changes in the environment. If the update rate of the sensor is high enough compared to the velocity of the vehicle, consecutive sensor readings can be compared with each other. This comparison can simultaneously yield information about the changes in the vehicle location and map the environment (SLAM – Simultaneous Localization and Mapping). The fusion can be extended to also Detect and Classify Moving Objects (DCMO). We monitor the output of a 75 Hz, 180^{o} laser range finder over time in order to identify, characterize, and classify stationary and moving objects in the vicinity of the vehicle. In the image above the output of the laser scanner is overlaid in yellow over the video image and the identified object is a pedestrian, marked with a red rectangle.

Fusion of different sensors has been used to detect and follow continuous structures alongside the vehicle. Output of a triangulation sensor, a vehicle state estimator, a video camera, and an obstacle detector is fused to obtain the position of a curb, a barrier, or a wall along and in front of the vehicle. The triangulation measures the position of the structure at one place. With the information of the vehicle state added, the position of the object is tracked alongside the vehicle. This track is then used to initialize a search of the structure in the video image, which gives a prediction for the curb in front of the vehicle. Finally the obstacle detector is used to eliminate false readings caused by objects (e.g. pedestrian) standing next to the structure. It should be noted, that the sensors fused have very different modalities. The triangulation sensor gives very good distance information, but only for one slice through the environment. The video on the other hand has a comparatively large coverage area, but no direct distance information. The distance information can only be inferred from the video by assuming that it views the relatively flat surface of the road. This assumption breaks down if an obstacle is on the road, a fact that can only be tested with a distance sensor like the laser scanner. The above illustration shows a detected curb alongside and in front of the vehicle in the video image as well as in the bird's-eye-view.

Once the sensors give the information about the environment, objects, and their properties, it is combined with other information. Such information can be preloaded maps or kinematic properties of certain classes of objects, models of behavior, or information from other source. One combination is the prediction of the movement of the vehicle and the objects in order to avoid collisions, plan paths, or anticipate the location of friends or foes. We use a Monte Carlo approach to this prediction. For the vehicle and each object many random paths are chosen consistent with the initial conditions measured by the sensors and also consistent with the respective kinematic properties of the vehicle or the objects. The behavioral models are implemented by weighting each path with an appropriate weight, e.g. paths where a vehicle stays on a road have a higher weight than paths where the vehicle leaves the road. The Monte Carlo method results in a probability distribution of positions of the vehicle and the objects represented by weighted points. In our example we were interested if at any time in the next few seconds the probability of a collision with an object exceeds a certain threshold. This was the case for the pedestrian in the right front of the vehicle; the rectangle representing the pedestrian was therefore marked red.

We have illustrated above, how we used fusion over time, fusion of different kind of sensors, and combined the sensor information with other information by implementing a Monte Carlo method. As an example we have shown how to find the probability of collision with another object which is an important aspect of 360 degree surround sensing.