Proceedings of the SPIE Aerosense Conference, Orlando, FL, April 21 - 25, 2003.

Using A Priori Data for Prediction and Object Recognition in an Autonomous Mobile Vehicle

Chris Scrapper, Ayako Takeuchi, Tommy Chang, Tsai Hong, Michael Shneier
Intelligent Systems Division
National Institute of Standards and Technology
100 Bureau Drive, Stop 8230
Gaithersburg, MD 20899

ABSTRACT

A robotic vehicle needs to understand the terrain and features around it if it is to be able to navigate complex environments such as road systems. By taking advantage of the fact that such vehicles also need accurate knowledge of their own location and orientation, we have developed a sensing and object recognition system based on information about the area where the vehicle is expected to operate. The information is collected through aerial surveys, from maps, and by previous traverses of the terrain by the vehicle. It takes the form of terrain elevation information, feature information (roads, road signs, trees, ponds, fences, etc.) and constraint information (e.g., one-way streets). We have implemented such an a priori database using One Semi-Automated Forces (OneSAF), a military simulation environment. Using the Inertial Navigation System and Global Positioning System (GPS) on the NIST High Mobility Multi-purpose Wheeled Vehicle (HMMWV) to provide indexing into the database, we extract all the elevation and feature information for a region surrounding the vehicle as it moves about the NIST campus. This information has also been mapped into the sensor coordinate systems. For example, processing the information from an imaging Laser Detection And Ranging (LADAR) that scans a region in front of the vehicle has been greatly simplified by generating a prediction image by scanning the corresponding region in the a priori model. This allows the system to focus the search for a particular feature in a small region around where the a priori information predicts it will appear. It also permits immediate identification of features that match the expectations. Results indicate that this processing can be performed in real time.

Keywords: robot vehicle, a priori knowledge, prediction, recognition, world model, sensory processing.

1 INTRODUCTION

Autonomous mobile vehicles typically make use only of information they can acquire through their sensors. This limits their capabilities because sensors have a limited range, and some aspects of the world, such as the names of objects, cannot be sensed directly. In addition, there are many behaviors that are applicable only in certain situations, such as the rules of the road. A vehicle that does not have access to a priori information could not apply such rules.

In order to address some of the issues that arise when a priori information is available, we have developed a means of accessing information in a database using the position of the vehicle as an index. This provides a continuous mapping from the database to the vehicle. So long as the position is known, relevant information should always be provided. It is not always possible to know exactly where the vehicle is in the world, and even if the position is known precisely, the information in the database may be incorrect or out of date, and the a priori information may not be correctly registered with the real world.

Our vehicle deals with these contingencies by maintaining confidences in each piece of information. Confidences are computed by combining information over time on how consistently the sensors report each piece of information, and how well the sensory data matches the a priori knowledge. The confidences are updated when new information comes in. The weight given to each type of information can change dynamically, depending on the current state of the vehicle. Thus, the vehicle takes advantage of a priori information when it is available, but can do without it and use only the on-board sensors when it is not. Further, if the a priori information contradicts the sensors, the sensors take precedence, although the confidence in what they report is reduced.

Our source of a priori information takes the form of a OneSAF database[1]. OneSAF is a military simulator that is commonly used for mission rehearsal and planning. It can store terrain information as well as dynamic information about objects and vehicles that move in the terrain. Our database is populated from a number of sources, the major one being an aerial survey of the grounds at the National Institute of Standards and Technology (NIST). The survey provides high resolution information about the roads, buildings, trees, lakes, and other features at NIST and part of the surrounding area. The features are annotated with names (or class membership) and properties that may help the sensors to recognize them.

The position of the vehicle is determined using an inertial navigation system coupled with data from the Global Positioning System (GPS). Typically, the position can be known to within a few meters, but the uncertainty can increase when the GPS information is lost. This may happen, for example, when the vehicle drives through the woods.

The vehicle uses two main imaging sensors, a color camera and an imaging ladar, to sample the terrain in its immediate vicinity. Both sensors are effective for a fairly short distance, partly because of inherent range limits, and partly because of the height at which they are mounted on the vehicle. The vehicle builds and maintains an internal model of the world that extends for about 30 m around it in all directions. The vehicle is always kept at the midpoint of the model, so the model is scrolled as the vehicle moves. The size of the model nicely limits the information that needs to be requested from the a priori database. We note that the vehicle maintains world models at more than one level of resolution, according to the principles of the 4D/RCS control system that provides its autonomous navigation capabilities[2, 3]. We currently make use of the a priori information described in this paper at only one level of resolution (30 meters around the vehicle, with a 20 cm by 20 cm grid).

The combination of vehicle position and orientation information (pose), input from sensors describing the local terrain, and a geographic database that can be indexed by vehicle position are the essential foundation on which a priori knowledge can effectively be applied. We describe how we use a priori information to predict what should be seen in ladar images, and how the information can be used to perform low-level segmentation of the images. The a priori knowledge also enables the vehicle to focus its attention on regions predicted to be interesting, and to apply only a subset of its feature detectors to locate predicted features.

2 ARCHITECTURE

The a priori knowledge module is divided into three processes: The Geographic Database Server, the Area Feature Server and the Area Feature Client. The processes communicate with each other and with the database server using the Neutral Messaging Language (NML)[4]. NML is also used to transmit data from the Area Feature Client to the sensory processing and world modeling processes.

The Geographic Database Server provides an interface to the OneSAF simulator, which contains a map of the region over which the vehicle traverses. A number of queries have been defined that return information about the simulated world. Of most interest for our application are queries that refer to a specific area. Given a location, the database can return the topography and features within a specified area around the location. Features include woods, rivers, roads, etc., and other objects such as vehicles that may be moving. Another useful query sets up a standing request for changes in a given region. Whenever anything enters or leaves this area, a notification is sent to the requesting process.

OneSAF is configured to allow a client to query its database and obtain information such as the object type and the vertices that define the object's borders. Each vertex is represented as a point described by its GPS coordinates. A query may result in several response messages, and the time to completion of the response is proportional to the number of feature objects returned from the query. This uncertainty in response time can be a problem for a vehicle driving in real time, and leads to the separation between the Area Feature Server and the Area Feature Client. The Area Feature Server takes the responsibility for querying the database in advance and always having up to date information about the area the vehicle is traversing. These data are sent to the Area Feature Client, which interacts in real time with the sensory processing and world modeling elements.

The Area Feature Server constructs an internal representation of the terrain to be used by the Area Feature Client to recognize features in an area around the vehicle. Using vehicle pose data to index into the OneSAF database, the Area Feature Server constructs a feature map representing the terrain features in the area centered on the vehicle's current location. The Area Feature Server is responsible for keeping the features up to date, and for ensuring that the correct region is always represented in the feature map by the time the vehicle starts to traverse it.

When the vehicle first starts up, the server queries the OneSAF database and creates the initial feature map centered on the vehicle's starting position. The server fixes the dimensions of the area to be queried and defines the radius of a circular region of safe operation within the feature map. Using a simple distance equation, the server predicts when the vehicle will leave the safe operating area. When this happens, the server resets the center of the feature map to the current position and queries the OneSAF database for the new terrain information. Querying in advance enables the internal data structures and feature map used by the Area Feature Server to remain up to date and locked onto a defined area around the vehicle.

The Area Feature Client contains a 2D vector model that approximates the robotic vehicle's size and the fields of view of its sensors. Vehicle pose data is used to orient the vector model. The Area Feature Client thus contains a priori knowledge for the region around the vehicle, registered with the actual terrain. The knowledge is used by the system to focus its attention on features of interest in the terrain, and to apply appropriate sensor processing operations to detect the expected features. When the sensors are successful in locating a feature, they can label it, because they know its identity from the a priori knowledge.

The feature map produced by the Area Feature Server covers a wider area than needed at any one time by the Area Feature Client. The Client copies a portion of the map that represents the vehicle's immediate vicinity. The navigation data that provides the vehicle's position is also used to keep the a priori knowledge data centered and orientated appropriately on the vehicle. The Area Feature Client determines when the vehicle has moved one unit of resolution in any direction and updates the map by asking the Server to send it the data to fill in the new visible region. Data no longer in the area are deleted.

It is also possible to construct features in the a priori data that are not defined in the database. For example, road intersections are not a defined type in the OneSAF database, but are considered important for mobile navigation. Being able to predict the location of intersections is useful for planning, so a module was developed in the Area Feature Server that detects intersections. This algorithm creates an intersection object that contains references to the corresponding road segments. The graphs of the different road segments can then be bridged by the intersection object, which allows a vehicle to navigate through the road network. Constructed features are treated as a priori knowledge in the rest of the system.

3 USING THE A PRIORI DATA

The a priori data are used to populate the world model with features that are expected to appear within the area surrounding the vehicle. This information is merged with other information acquired from the sensors and used to provide predicted features to sensory processing modules. The sensory processing modules use the information to guide their feature extraction and recognition routines.

The Area Feature Client communicates with the world modeling processes using NML channels that allow different modes of interaction. The available methods of delivery include sending the entire region in response to every request, sending only the part of the region that changes as the vehicle moves, and initially sending the entire region followed by incremental updates. The processes that call the Area Feature Client have control of the data that is sent, enabling them to prioritize their tasks and allowing both modules to operate without being blocked.

The Area Feature Client periodically broadcasts a message that provides a summary of the available a priori data. The message includes GPS coordinates of the vectors that define the area currently stored by the Area Feature Client and a set of flags that summarize the information available for the region. The flags also indicate whether or not the vehicle is on the road or in an intersection, and lists the features that are present in a rectangular region in front of the vehicle that corresponds to the view from the vehicle's sensors.

Inserting prior knowledge into the world model requires that the information in the world model be spatially registered with that in the Area Feature Client. The primary source of information in the world model is the vehicle's sensors, which provide their view of what surrounds the vehicle. The a priori information cannot simply override the sensors because it could be out of date or offset from the correct location. In our world model, information is always associated with a level of confidence. The confidence in a given feature is enhanced if successive views confirm its existence, and reduced if they don't. A priori data are treated like output from a virtual sensor that modifies the confidence in the feature. While the a priori knowledge can disagree with the sensed information, we currently have no procedure for modifying the a priori knowledge in the OneSAF database to reflect the real world more closely.

The information in the world model is used to influence future sensory processing by providing predictions of what will appear in the sensed data. In the case of the ladar range sensor, projecting the field of view of the sensor into the world model and constructing a synthetic image generates a predicted view. This image is the same size as that produced by the sensor, and can be compared pixel by pixel with the sensor data. Instead of range data, however, the predicted image contains the identities of the features at each location. Thus, when the vehicle is driving down the road, the prediction data would identify the road region segmented from the rest of the image, enabling simple confirmation algorithm to be applied rather than having to apply a range of more complex operators to the whole image.

We have constructed a OneSAF database of the NIST grounds based on an aerial survey. The NIST HMMWV (Figure 1) can drive around anywhere on the grounds and receive a priori information about the roads, buildings, trees, lakes, fences, etc., in its neighborhood. The black line behind the vehicle in Figure 2 shows the trace of the path followed by the HMMWV as it drove. The path is overlaid on a representation of the a priori data. The INS/GPS unit on the vehicle was used both to plot the path and to update the features on the display as the vehicle moved around.

Figure 3 shows a view of the vehicle's path overlaid on the a priori map. Figure 4 shows a ladar image taken when the HMMWV was at the same location in the real world, while Figure 5 shows the image resulting from projecting the ladar field of view into the world model. The prediction contains feature classes and names rather than range values. As can be seen, there is a good match between the a priori and sensed information, enabling the features to be segmented and classified without having to search the whole image. Such good registration is not always available only using INS and GPS, because the vehicle may lose the GPS signal and the inertial unit may then start to drift. We are working on using matches between a priori and sensed features to correct the drift and provide more accurate location information[5].

Figure 1. The NIST HMMWV

4 DISCUSSION AND CONCLUSIONS

Humans are constantly relating the environment around them to their store of knowledge. This is what allows us to recognize objects, situations, and events, and lets us switch our behavior to match the current environment. Many activities are greatly simplified by access to experience and by focusing attention on the important aspects of the situation. Autonomous systems typically have no access to a priori knowledge and can't adapt their behavior unless programmed explicitly to do so.

In this paper we showed how we are building the capability for an autonomous vehicle to make use of a priori knowledge to improve its understanding of the world and to target its sensory processing routines to the regions that are predicted to be most interesting. In order to ensure that the act of retrieving the relevant information does not adversely impact the real time aspects of the system, we divided the retrieval into separate parts, each of which was responsible for only part of the retrieval task. The Area Feature Server is responsible for interacting with the database server and ensuring that data are always available for the locations that the vehicle is traversing or is about to traverse. The Area Feature Client manages the interface between the a priori knowledge and the world model. The world model then projects the a priori knowledge into the sensor coordinate systems. All of these operations are made possible because the vehicle position and orientation are known at all times and the a priori knowledge is stored in a way that allows it to be retrieved based on its location in the real world.

Building the geographical knowledge base to make this project possible required obtaining an aerial survey and storing it in the OneSAF simulator. This is not a straightforward task, especially when given the high resolution data we desired. The availability of such information is growing, however, especially for road networks. Application of our techniques to autonomous vehicle that drive on roads should be feasible within the not too distant future.

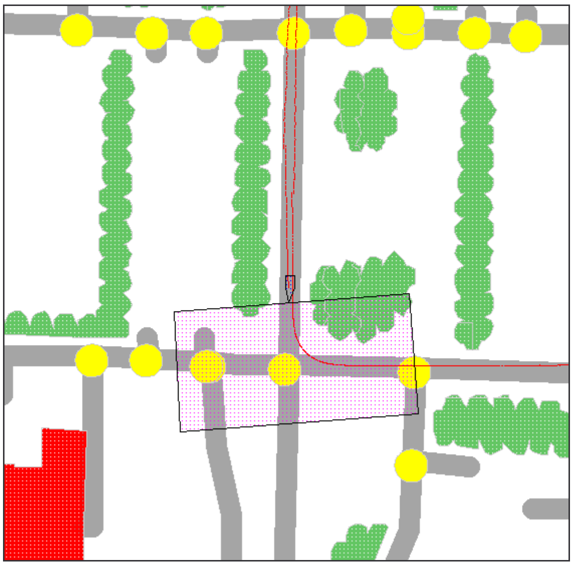

Figure 2 . HMMWV centered on the map containing a priori data obtained from OneSAF. Map shows roads, trees, building (lower left), and intersections (circles). The vehicle's path (lines on road) shows both the outbound and inbound path of the HMMWV. The rectangle indicates the projected field of view of the ladar sensor.

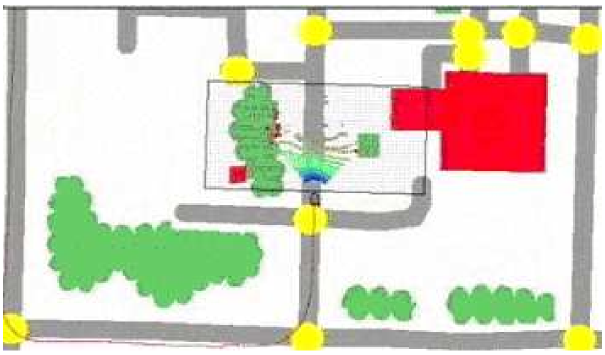

Figure 3 Feature map illustrating the a priori knowledge obtained from OneSAF. The rectangle corresponds to the field of view of the ladar sensor. Features included in the field of view are trees, roads, and part of a building.

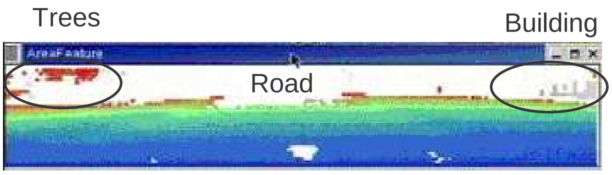

Figure 4 A ladar image taken when the NIST HMMWV was at the location shown in Figure 3. The location in the real world is based on GPS and inertial information, which is also used to define the query sent to the a priori knowledge base.

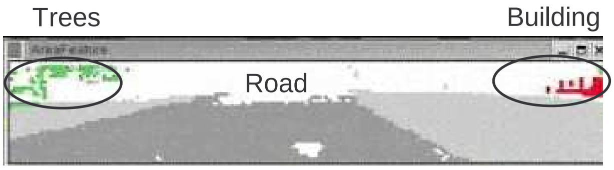

Figure 5 The predicted image generated by projecting the ladar field of view into the a priori knowledge map. It is easy to label the ladar image given the predictions in this image.

Acknowledgements

We are grateful for the support of the Army Research Laboratory and the Defense Advanced Research Projects Agency MARS program.

REFERENCES

1. OneSAF, http://www.onesaf.org/, 2003.

2. Albus, J. S. and Meystel, A., *Engineering of Mind: An Introduction to the Science of Intelligent Systems.* Wiley and Sons, Somerset, NJ, 2001.

3. Albus, J. S. et al. "4D/RCS Version 2.0: A Reference Model Architecture for Unmanned Vehicle Systems," National Institute of Standards and Technology, Gaithersburg, MD, NISTIR 6912, 2002.

4. Shackleford, W., Proctor, F, and Murphy, K, "The Neutral Message Language: A Model and Method for Message Passing in Heterogeneous Environments," *Proceedings of the 2000 World Automation Conference*, Maui, HI, 6-11-2000.

5. Madhavan, R., "Terrain Aided Localization of Autonomous Vehicles," *Proceedings of the 19th International Symposium on Automation and Robotics in Construction (ISARC 2002),* Gaithersburg, MD.