CHAPTER IX

SUMMARY AND CONCLUSIONS

§1. Summary of Work

This dissertation has attempted to accomplish many things in autonomous robot vehicle navigation. These include:

- Global planning in the context of *a priori* knowledge
- Reactive/reflexive navigation in response to sensor information
- Perceptual strategies to provide navigational guidance
- Spatial uncertainty management
- A framework to tie these components together

This chapter first reviews the work performed for each of the above items. Future research areas are then described. A discussion of the general contributions resulting from the work in this dissertation is then presented. A brief commentary on what the future holds for mobile robotics concludes the chapter.

§*1.1 Global planning*

Global planning is accomplished through the use of digitized maps or blueprints, embedded with additional knowledge regarding terrain type, landmarks, etc. A convex decomposition algorithm operates on the original map to produce a meadow map, a connectivity graph tying together the individual convex regions. Two variations of an A* search algorithm are available to search for a "reasonable" global path through the modeled world. Chapter 4 describes the algorithms, rationale, and many advantages of the meadow map representation. These include: the flexibility to produce short, safe, and other types of paths; the ease of embedding information to guide perceptual processing, the reasonable storage requirements; and the ability to conduct multi-terrain navigation in a straightforward manner.

The initial coarse path is improved through the use of strategies which tighten and straighten the unrefined path. Both this and the extension of this method to incorporate multiple terrain types are significant advances over previous work in this domain. The navigator's path planning techniques, both for single and multi-terrain navigation are described in detail in Chapter 4. The A^* search strategies provide rapid initial path determination. The path strategies developed, for both tightening and straightening the raw path, produce high quality paths which previous methods, using similar representations, were unable to accomplish.

The global path planning component of the AuRA architecture is complete. The navigator, mapbuilder, feature editor, and long-term memory (LTM) component are fully implemented. The mission planner exists in a rudimentary version and is an area that holds great promise for future research (see below).

Good sources for navigational knowledge include building blueprints, aerial maps, and the various mapping agencies (e.g. Defense Mapping Agency). The data used to construct the landmark models for guiding perceptual processing can come from engineering drawings (e.g. for lampposts or fire hydrants). Spectral data pertinent to landmarks can be dynamically acquired from histogram processes running on images of the landmarks. In general, however, knowledge acquisition is a difficult and laborious task. Much additional research is needed to make this work easier (see Section 2).

Even if portions of the data within the *a priori* knowledge base are incorrect or inaccurate, in many instances the system performance will still be robust. Actually, navigation with little if any *a priori* knowledge is quite feasible, simply by having a means of goal recognition, a rough idea of the goal's location, and avoiding obstacles as the robot moves. In this manner, all obstacles are treated as unmodeled and high-level path planning is circumvented. Localization is of no concern as there is no world map to correlate against and schema-based navigation is the fundamental navigational principle. Efficiency, reliability, and safety, however, are sacrificed somewhat when the *a priori* knowledge is absent.

§1.2 *Reactive/reflexive navigation*

The true role of the global planner is to provide a context in which the pilot can select appropriate motor behaviors (schemas) to satisfy the navigational subgoals. These motor behaviors are exemplified by **move-ahead**, **move-to-goal**, **avoid-obstacle**, **stay-on-path** and **follow-the-leader** schemas. Each of these motor schemas is associated with one or more perceptual schemas/strategies that provide the sensor data essential for the successful completion of the particular motor task. Each individual motor schema, ideally operating in a concurrent manner, contributes a velocity vector germane to its particular motor task. These are then summed by the **move-robot** motor schema and transmitted to the motor control subsystem to update the robot's current direction and speed. Chapter 5 describes the motivation for schema usage, the potential field methodology used, and simulation results. Chapter 8 demonstrates, via experimentation, the utility of motor schema based navigation.

This control regime maps well to a distributed processing system and affords a high level of modularity that simplifies the addition and debugging of new perceptual or motor schemas. Complex behaviors are constructed through the combination of multiple simple behaviors. When suitable tools become available on the Sequent multiprocessor (i.e. the Schema Shell [41,42]) an implementation of the motor schema manager should follow. The experimental system described in Chapter 8 uses sequential evaluation of motor schemas to emulate concurrent schema activity.

Summing the output of the individual velocity vectors produced by each schema instantiation (SI) is a simplistic approach. The functions producing the velocity vectors themselves, in their linear implementation, are also simple. Nonetheless, good results for navigation have been obtained using these methods, both in simulation (Chapter 5) and in experimentation (Chapter 8). It is advisable to look at other methods for combining the outputs of the SIs and to consider other possibilities for the velocity formulations of the SIs themselves, to study just how much can be gained through the use of more complex functions and combination mechanisms.

The potential field methodology is not failsafe and reliance on a sensor-built world model constructed by the cartographer is important for recovery in instances of local potential minima or cyclic behavior. This intermediate planning level is invoked by the pilot when anomalous behavior is detected, typically by exceeding a hard real-time deadline for path completion (cycle detection) or recognition that the robot's velocity has dropped to unacceptably low levels (local minima). Some of the navigational problems can be circumvented by the addition of small amounts of noise to the navigational process, moving the robot away from potential plateaus. Chapter 4 describes the roles of the pilot for both schema selection and local navigation. A simulation illustrating this point is presented in Chapter 5.

§*1.3* *Perceptual strategies/schemas*

Action-oriented perception is the fundamental premise on which AuRA's perceptual strategies are based. By designing specific perceptual algorithms to solve particular navigational tasks, the amounts of computation required for intelligent navigation becomes tractable.

Vision is viewed within AuRA as a primary sensor modality. Chapter 6 details the different vision algorithms used to provide HARV with navigational capabilities. Unfortunately, the processing speeds for sophisticated algorithms on existing hardware currently preclude its use for real-time navigation. Nonetheless, experiments were conducted (Chapter 8), typically in lurch mode with the vehicle pausing during image processing, showing the ability of the robot to follow paths and to avoid obstacles using vision.

A fast-line finding algorithm was developed for both path edge extraction and vehicle localization purposes. The line-finder can be tuned to anticipated line orientation based on the availability of *a priori* knowledge or the results of previous frames. The use of a feedforward strategy to track features over time is typical of the vision algorithms.

As in all the vision work, this algorithm had to be accompanied by large amounts of software to accomplish the navigational tasks. This included:

- Grouping strategies to interpret the results of the image processing in the context of the navigational need (e.g. agglomeration of the line fragments to produce the path edges).
- Utilization of both *a priori* knowledge and data obtained from previous frames to reduce computational demand. (e.g. windowing the image).
- Interpretation of the output of the algorithm to produce intermediate results that reflect the navigational goals (e.g. determination of center road line).

- Conversion of the interpreted image into suitable motor commands (i.e. visual servoing).
- Communication of those commands to the vehicle.
- Production of new image expectations based upon the motor commands given the vehicle (e.g. where will the path appear in the next image now that the robot has moved).

A fast region segmentation algorithm (Chapter 6) is also available and is used to provide information for path following. Through knowledge of the spectral characteristics of the path or landmark in question, these features can be successfully extracted. The use of color will greatly enhance the ability of this algorithm. Chapter 8 describes an experimental run using this technique.

The depth-from-motion algorithm described in Chapter 6 provides the location of obstacles in the robot's environment. These detected obstacles can then be instantiated as **avoid-obstacle** motor schemas and subsequently conduct the robot's navigation. Chapter 8 presents experiments in both indoor and outdoor environments. This particular approach, unlike fast line finding and fast region segmentation, is particularly susceptible to problems in image registration. When this problem and others (such as the ability to extract the focus of expansion reliably and accurately) are solved, considerably better results are expected. Nonetheless, the ability of the algorithm to track feature points over multiple frames is indicative of the significant potential for this approach to obstacle avoidance.

Ultrasonic data is the most efficient means for demonstrating schema-based navigational abilities. Unfortunately, due to the limited resolution and discriminatory power of ultrasonic sensors, semantic interpretation is limited. Nonetheless, recognition of obstacles and tracking behavior are both readily implemented using this sensor form. Chapter 6 discusses the use of ultrasound in AuRA while Chapter 8 describes the experimental results obtained in the motor schema testbed used with HARV. Obstacle avoidance, wandering, impatient waiting, follow-the-leader, and other interesting behaviors have been demonstrated using this sensor modality.

Shaft encoders provide general directional ability as well as approximations of distance traveled. Although it would be a mistake to rely too heavily on the encoders due to relatively large errors that arise during path execution, they can still provide, in most cases, a reasonable coarse estimate of the robot's whereabouts. This is especially useful in the context of the uncertainty management subsystem described in chapter 7 and as applied in the localization experiment in Chapter 8. Chapter 6 discusses some of the general issues in using HARV's shaft encoders.

§*1.4 Spatial uncertainty management*

In order to provide expectations for the robot's perceptions and to make the computational requirements for perceptual processing more tractable, a means of managing both positional and orientational uncertainty was developed. The spatial uncertainty map consists of two components – the first a convex polygon representing the extent of position uncertainty of the robot relative to the global map, and the second a compass wedge explicitly representing the orientation uncertainty. Growth in uncertainty occurs when the robot moves, based on statistical data reflecting the terrain type being traversed. Reduction in uncertainty occurs through the recognition of landmarks and the subsequent pruning of the spatial uncertainty map in a manner appropriate to the particular landmark. The spatial uncertainty map itself provides the basis for the Expecter process to provide image windows for landmark recognition by active perceptual schemas.

The uncertainty map manager, responsible for both the growth and reduction of the spatial uncertainty map, has been implemented in LISP. The ties connecting the LTM structure (coded in C) have been completed. Chapter 7 presents the spatial uncertainty management approach used within AuRA and Chapter 8 includes an experiment illustrating how vehicle localization is accomplished with this software.

Additional areas of uncertainty management, outside of the area of spatial uncertainty, undoubtedly will need to be addressed as the AuRA architecture grows. The treatment of symbolic uncertainty at the mission planner level, in particular, must be dealt with.

§*1.5 The AuRA framework*

AuRA, the Autonomous Robot Architecture, is the structure in which these individual components are linked. Chapter 3 presents an overview of the AuRA architecture and includes a discussion of the system issues involved. The overall architecture is quite ambitious, ultimately incorporating the UMASS VISIONS scene interpretation system running on the UMASS Image Understanding Architecture. The first pass implementation of AuRA is much more modest, tying the available components together through blackboard-like global data structures (clipboards) or through the use of other expedient interim techniques. The AuRA's system architecture is not complete at this time, containing only a rudimentary mission planner, limited environmental knowledge embedded in LTM, a partial set of pilot rules for schema selection, some limitations on the current implementation of the uncertainty management subsystem, and limitations in the interfaces between components. The scope of this project, however, extends several years beyond this dissertation as the AuRA architecture becomes more fully integrated.

§2. Future research

Mobile robotics is a science in its infancy. Much remains to be done in almost all areas. There are several aspects of future research that would be quite appropriate for embedding in AuRA. These are discussed below.

§2.1 *Additional sensors*

AuRA's modularity makes it straightforward to embed new sensors and sensor algorithms. The use of inclinometers for vehicle orientation in a non-planar world would assist in the difficult problem of image registration. A literal "hill-climbing" schema could be embedded if the robot had the ability to measure its pitch and roll.

A laser range finder is highly desirable, but currently prohibitively expensive. Future funding may enable its acquisition thus enhancing the vehicle's ability to detect obstacles.

Color vision would greatly expand the possibilities for the fast region segmentation algorithm. Much remains to be done to apply the use of color in the domain of action-oriented perception.

§2.2 *Extension to three dimensions*

Much of this architecture is generalizable to robots capable of maneuvering in three dimensions. This would include space and undersea vehicles. Reformulation of the potential field equations used by the motor schemas to three dimensions is straightforward. The same would hold for the uncertainty management subsystem except for possible problems in accounting for vehicle drift in the presence of ocean currents or the like. The meadow map representation, although capable of being extended to three dimensions may pose some mathematical difficulties, but certainly bears investigation. Some preliminary work on these extensions to provide for navigation in three dimensions is presented in [15].

Another area for exploration is a 3D ground-based world using the meadows as topographical facets with representations of their slope for path planning purposes. This would result in a piecewise planar world as opposed to a single ground plane.

§*2.3 Learning*

Several different levels of learning can be addressed within AuRA. The movement of information stored in short-term memory to long-term memory is a form of learning, as the representation persists after the robot has moved out of the area. In order to do this effectively however, semantic interpretation of the objects in question must be undertaken to distinguish whether the object to be learned is dynamic in nature (hence will move and is thus inappropriate for embedding in LTM as it may not be present the next time the robot is in the area) or is static and worth retaining in LTM. This is not a simple task by any means. Continued investigation into scene analysis (e.g. the VISIONS natural scene interpretation system) to provide both semantic interpretation of previously unrecognized environmental objects as well as construction of world models to guide sensor processing will potentially "open the eyes" of the robot.

Another form of learning involves the assimilation of motor skills. Just as robots are currently taught using a teaching pendant in industrial environments, perhaps through the combination of perceptual learning (e.g. by a reward/punishment regimen) with trained motor behavior, the need for mathematical formulae to represent potential fields at the programming level would disappear.

Adaptation of the gains on the motor schema outputs is also an important area for the application of learning theory. In this manner, the robot can learn from its experience as to which schema gains are best for a particular task. For a task requiring a delicate balancing of schema outputs, such as door entry, the tuning of these outputs is important.

In one sense, learning already exists in several of the perception algorithms used within AuRA. By utilizing the knowledge obtained from previous frames to guide subsequent processing (e.g. path edge position), the robot has acquired a model of its environment.

This information, however, is not currently used to update long-term memory representations and is discarded at the end of the navigational task.

§2.4 *Homeostatic control*

The niche is already present for homeostatic control within the AuRA framework (Chapter 3). Certainly the need exists for robots capable of managing their own internal resources, if they are to operate in environments that are too hazardous for man. By equipping the vehicle with temperature and energy sensors, the motor schema approach can be extended (as described in chapter 3) to include signal schemas. These signal schemas would modify the dynamic motor behavior of the robot based on continuous internal sensing. Instead of using worst-case analysis in the formulation of global plans, dynamic replanning would occur based on available information. This would result not only in safer robots, but also more efficient robotic systems.

§3. General contributions

The general contributions for this work include:

- **Extensions of techniques for representing world knowledge to include multi-terrain representations.**

  Previous approaches, typically using a regular grid representation, suffered from high memory requirements, large search times, and digitization bias, all of which are avoided with the meadow map approach (Chapter 4).

- **The ability to conduct path planning in diverse terrain types in an efficient manner.**

  By only resorting to path relaxation at terrain boundaries and by using the path improvement strategies mentioned below, quality paths can be obtained at low computational costs (Chapter 4).

- **The invention of path improvement techniques that refine initially coarse and unacceptable paths into "reasonable" paths.**

By adding tightening and straightening techniques to remove unnecessary twists and turns in the raw global path, significant improvement over previous path planning work using meadow maps is evidenced. These improvement strategies enhance the navigator's ability to produce safe or short paths depending on the mission planner's goals (Chapter 4).

- **The embedding of knowledge to guide uncertainty management and perceptual processing into the long-term memory representation.**

  Landmarks, terrain statistical data, spectral characteristics, etc., can be tied to the LTM representation through the use of the feature editor. The hybrid free-space vertex-graph representation (meadow map) saves information on both the obstacles and free space, unlike several other representation strategies (e.g. pure vertex graph and Voronoi diagrams), making the addition of further knowledge much easier (Chapter 4).

- **Decomposition of reactive/reflexive behaviors into a set of simple motor schemas which when combined react intelligently to sensory events.**

  By having individual motor schemas produce outputs that are relevant to the overall system goal, intelligent navigation can be achieved. The pilot selects these schemas based upon the current navigational subgoal, and parameterizes them to fit the task at hand (Chapter 5).

- **Making path execution amenable to distributed processing by allowing multiple concurrent schema instantiations operating asynchronously to produce the motor control for the robot.**

  The individual schema instantiations can run on separate processors each posting their output on a blackboard which is monitored by the move-robot SI. The combined results are then forwarded to the motor subsystem for execution (Chapter 5).

- **Applying potential field methodology to sensory-based navigation and incorporating the uncertainty in perception into the resultant velocity fields.**

  Representing a motor goal by a velocity field is an important means for accomplishing a variety of tasks. Fields have been specified for obstacle avoidance, goal attraction, path following, etc. The computation to produce the resultant vector is quite simple and can be performed rapidly (Chapter 5). Experimental results in Chapter 8 bear out these conclusions.

- **The embedding of action-oriented perceptual strategies specific to the task within the motor schema itself.**

  Action-oriented perception can only be accomplished by associating perceptual techniques with specific motor skills. A further advantage is the added capability to use multiprocessing for perception due to the natural partitioning of perceptual tasks based upon specific motor needs (Chapters 5 and 6).

- **The application of new visual perception techniques to the navigation problem (i.e. path finding, obstacle avoidance and localization) including fast line finding, fast region segmentation and a depth-from-motion algorithm.**

  Experimental results using these techniques are presented in Chapter 8. Some of the algorithms are quite robust (e.g. fast line finding) while others are more sensitive to environmental conditions such as camera roll (e.g. depth-from-motion). Much has been learned from both the successes and difficulties encountered by the application of these algorithms to navigation (Chapter 6).

- **A new technique for the management of spatial uncertainty.**

  By explicitly representing the degree of location and orientation uncertainty relative to the *a priori* world map information, the ability to establish expectations for perceptual processing is gained. The use of specialized uncertainty growth techniques based on empirical terrain data and uncertainty reduction techniques based on landmark recognition provide the feedback necessary for uncertainty management (Chapter 7).

- **The AuRA framework itself which ties together, in a cybernetic context, the individual components of the architecture.**

  Schema-based navigation and action-oriented perception are the principal tenets on which AuRA is founded (Chapters 3, 5, and 6).

§4. Directions

The field of mobile robotics foretells a bright and promising future. As inroads are made into this artificial intelligence problem, more and more applications will arise. Already there exist sentry robots, industrial cleaning robots, and mobile robots that operate in the manufacturing domain. The military is strongly supporting research into autonomous vehicles for battlefield, resupply, and reconnaissance operations.

There seems to be some innate fascination in people regarding autonomous machines. Even though Hollywood has come to treat autonomous robots as a solved problem (e.g. R2D2), the mystique continues. What perhaps is most important to the future of this field is that researchers do not resort to a series of contrived demonstrations to prove that their robots work. We, as artificial intelligence researchers, must learn from the mistakes that our predecessors in the 1950s and 60s made, especially in promising too much too soon.

The key, in my estimation, is to make steady progress through fundamental research, both in perception and motor control. Too much emphasis on applications will detract from the real issues involved. Fortunately, there exist excellent mobile systems (animals), provided by God, that can be studied to yield insights into the mobility problem. The biological successes that already are in place should not be ignored by the roboticist.

In conclusion, if our expectations are realistic, the support is consistent, and work in other domains is drawn upon, a new era in intelligent machines awaits.