Experiences with an Autonomous Robot Attending the AAAI Conference

François Michaud, Jonathan Audet, Dominic Létourneau,
Luc Lussier, Catherine Théberge-Turmel and Serge Caron
LABORIUS - Research Laboratory on Mobile Robotics and Intelligent Systems
Department of Electrical and Computer Engineering
Université de Sherbrooke, Sherbrooke (Québec Canada) J1K 2R1
laborius@gel.usherb.ca, http://www.gel.usherb.ca/laborius

Initiated in 1999, the AAAI Mobile Robot Challenge is to make a robot attend the National Conference on AI: *"In this event, a particularly challenging task is defined which is well beyond current capabilities, will require multiple years to solve, and should encourage larger teams and collaborative efforts. ...the task is for a robot to be dropped off at the front door of the conference venue, register itself as a student volunteer, perform various tasks as assigned, and talk at a session. The Challenge will require integration of many areas of artificial intelligence as well as robotics".*

Different research projects are underway at our laboratory with the objective of designing software and hardware capabilities that will allow autonomous systems to assist humans in real life conditions. For instance, a mobile robot that can read symbols would be able to derive useful indications from signs located in public places. The use of artificial emotions would allow the robot to adapt more efficiently to the environment by monitoring the progression in the accomplishment of its goals. Recharging itself whenever necessary is also something required for mobile robots to be used over prolonged periods of activity. A touch screen is a simple interface that allows to display visual information and get information from users. Our research is also oriented toward the validation of an architectural methodology that integrates these abilities in a coherent and efficient manner.

Motivated by the idea of making a robot attend a conference, we decided to take on the Challenge by addressing a simplified version of the task from start to end, and see how far we could go with the robotic platform and material we had. Our AAAI 2000 robot entry is a Pioneer 2 DX robot with 16 sonars, a Pan-Tilt-Zoom (PTZ) camera with a frame grabber, a RF Ethernet-modem connection and a Pentium 233 MHz PC-104 onboard computer. We also installed a touch screen monitor on top of the robot. The robot is able to dock into a charging station to ensure its energetic autonomy.

Research Projects on Integrated Intelligence

Intelligence manifests itself in many ways: reacting to what is perceived, planning actions to take to accomplish a task, modelling the world, reasoning about goals and events, etc. Different mechanisms are required to reproduce on machines these abilities while still preserving their underlying principles, and they need to be integrated appropriately into a common framework. It is not yet understood how exactly this should be accomplished [1] and it is by experimentation that we are able to see what works and what does not. In our case, our work follows the guidelines of an hybrid reactive-deliberative architectural methodology [2] represented in Figure 1. Behavior-producing modules (or *Behaviors*) are the basic components that allow to control the robot's actuators. The behaviors selected reflect what are the intentions of the robot, which come from *Implicit* conditions determined only by sensory inputs, from a hierarchical organization of the goals of the robot (managed by the *Egoistical* module), and from reasoning done by the *Rational* module using innate or acquired knowledge about the task. Note that information can be exchanged between the *Rational* module and the behaviors. Processes for these three modules all run independently and concurrently to derive behavior recommendations, i.e., by indicating if a behavior is desirable or undesirable. This ensures that emergence is also preserved for the selection of behaviors. The *Behavior Selection* module combines all recommendations and simply activates behaviors that are more desirable than undesirable. Finally, *Motives* are used to manage the robot's goals, while *Emotions* are there to monitor the accomplishment of the goals over time. One important source of information for motives is the observation of the effective use of the behaviors, represented by the link *Behavior.Exploitation*, which can serve as an abstraction of the robot's interactions within the environment. An active behavior may or may not be used to control the robot, depending on to the sensory conditions it monitors and the arbitration mechanism used to coordinate the robot's behaviors. So, an active behavior is exploited only when it provides commands that actually control the robot.

Different approaches can be used in the architecture's modules to implement the necessary intelligent decision mechanisms based on the robot's capabilities and purposes. With our Pioneer 2 robots, we are interested in working on new capabilities or capabilities not frequently addressed, as follows:

Reading symbols. In real life settings, we can find all kinds of symbols and signs (like room number, exit signs, etc.) that help us find our way or provide useful information about the environment. With all the research activities involving handwritten or printed symbol recognition going on for quite some time now [3], we think that it would be interesting to give this capability to a mobile robot. Making a robot recognize printed symbols is an interesting idea because it can be a method shared by different types of robots, as long as they have a vision system. The information is also accessible by humans, which is not possible when electronic media are used for deriving information about the environment.

Artificial emotion. The concept of artificial emotion can contribute greatly in designing autonomous robotic agents, mostly by making robots respond emotionally to situations experienced in the world or to interactions with humans [4, 5]. However, they can also play an important role in solving what are called *universal problems of*

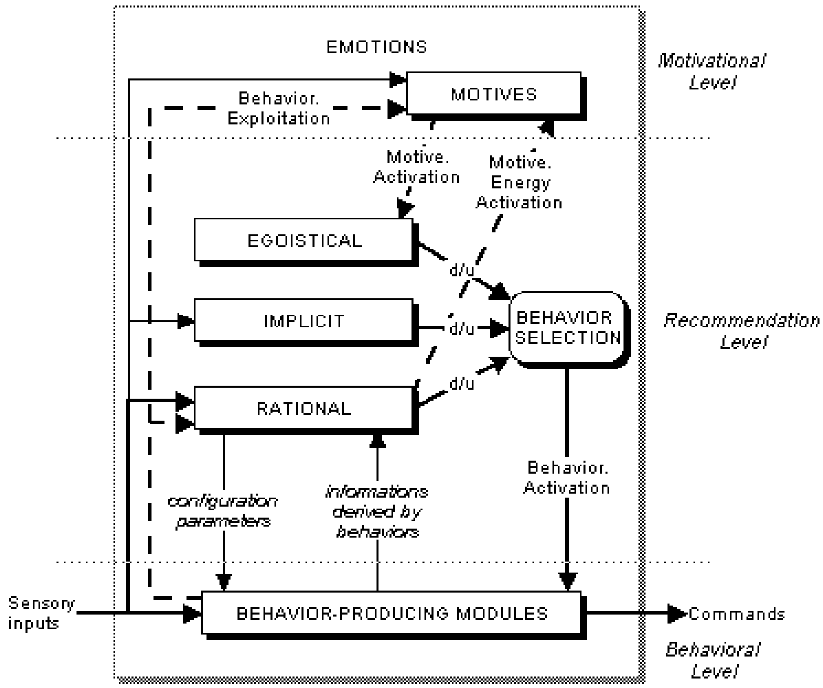

Figure 1: Software architecture used with our robot. Behavior recommendations are referenced as d/u, for behavior desirability/undesirability

adaptation [6]: hierarchy, territoriality, identity and temporality. For the Challenge, temporality, i.e., taking into consideration the limited duration of an individual's life, is the main adaptation problem. Artificial emotion can help a robot manage its own limitations in interacting with the world, interact socially with others and establish a shared meaning for interpersonal communication [7]. From an engineering point of view, our objective is to provide mobile robot with such capabilities.

Autonomous recharge. Autonomous robots have limited energy and sooner or later their batteries need to be recharged. We designed a charging station that can be perceived by the robot using an infrared beacon. A rign of infrared detectors located at the back of the robot allows it to dock into the charging station.

Touch screen interaction. Touch screens are frequently used as simple devices to interact with people. Installed on top of our robot, the touch screen monitor has two purposes: generate simple facial expressions to monitor the emotional states of the robot, and interact with people using menu screens.

Strategy of our AAAI 2000 Robot Entry

Since the AAAI Mobile Robot Challenge is an ambitious task, it does not have a hard-and-fast specification; a robot can attempt just a few of the possible sub-tasks, or a simplified version (based on assumptions made by the teams and approved by the judges) of the whole task. Here is the strategy that we planned on taking:

- **Registration**. We assumed that the robot starts from the front door of the conference center. Since no maps are allowed for reaching the registration desk, we decided to use symbols to indicate the direction to take. Left and right arrows indicate that the robot must turn in that direction. An up arrow is for continuing in that direction, while a down arrow indicates that the robot should go the other way. The registration line is identified with the symbol **L**. The robot must then wait in line and follow the person in front (using its sonars), until it arrives in front of the line. Then, the robot starts to search for the appropriate registration desk, identified with its last name initial, **H**. Once reached, instructions for its presentation (like where and when) can be entered using the touch screen. During the registration phase, symbols other than arrows, **L** or **H** may be encountered but would not be considered by the robot. Also, when it takes too long to find the registration desk, the robot asks somebody to guide it to the registration line by following the person, using the same mechanism for waiting in line.

- **Taking the elevator**. The next step is to take the elevator, with its location identified by the letter **E**. No other symbols are considered useful for this phase. The robot tries to find the elevator by itself, but depending on the time spent and the time remaining, the robot can also ask somebody to guide it to the elevator by following the person. Once the **E** symbol recognized, since we did not used any mechanisms to locate the elevator doors (like edge detection for instance), the robot assumes that the elevator door is on the right. The robot turns and waits in front of the elevator, asking for somebody to open the doors. When the doors open, the robot enters and goes to the back of the elevator, makes a 180° turn, stops, asks to go to the appropriate floor and to be told when it is time to leave the elevator. The robot then leaves by continuously going in the direction that is free of obstacles.

- **Schmooze**. If time permits, the robot can schmooze before going to the conference room. It does that by tracking important people using badge color, going toward such person and interacting using the touch screen. A snapshot of the person is taken based on the position of the arms and face of the person, derived from skin tone.

- **Guard**. Again, if time permits, the robot can guard an area if somebody makes such request using the touch screen. The robot simply stops and tracks people based on skin color. Snapshots are taken when appropriate.

- **Navigate to the conference room**. When the robot believes it is time to go to the assigned conference room, it starts to follow directions given by arrow signs and room number. When it finds the appropriate room number, it turns in the direction assumed for the position of the door (left in our implementation), and wait to be told that it is its turn to present.

- **Presentation**. When its turn comes, the robot starts to look for the charging station, because that is where it is going to make its presentation. It does so by looking for the charging symbol or by sensing the infrared beacon of the charging station. If it finds the charging symbol, the robot makes a 180° turn and backs up into the charging station. The robot then takes a snapshot of the audience and generates a HTML report describing what it experienced during its journey.

Note that when the robot executes a command according to a symbol detected, it positions itself relative to the symbol based on the position of the Pan-Tilt-Zoom camera and sonar readings. Also, the robot may decide to find the charging station when appropriate (for instance when the amount of energy is getting too low), anytime during the phases explained previously. All of the images taken for symbol recognition are also memorized in the HTML report along with messages describing the decision steps made by the robot.

Behavioral Level

To implement such strategy, Figure 2 shows the organization of the behaviors used. Behavior arbitration is done using Subsumption [8]. These behaviors do the following: *Safe-Velocity* makes the robot move forward without colliding with an object; *Random-Turn* is used for wandering; *Follow-Color* tracks objects of a specific color, like for schmoozing; *Skin-Tracking* controls the camera to capture the face and arms of a person using color detection; *Follow-Wall* makes the robot follow walls, which is useful to help the robot find the charging station; *Sign-Tracking* tracks, using a PID controller, a symbol of a specific color (black in our case) surrounded by another color (orange), and zooms in to get an image of maximum resolution for the symbol to identify; *Wait-in-Line* makes the robot follow the object in front, using its sonars; *Direct* changes the position of the robot according to specific commands generated by the *Rational* module; *Recharge* makes the robot dock into the charging station; *Avoid* makes the robot move away from obstacles; *Unstall* tries to move the robot away from a position where it has stalled; *Free-Turn* turns the robot in the direction free of obstacles; *Backup* makes the robot move backward; *Rest* stops the robot; and *Control-Display* controls the touch screen interface. The behaviors control the velocity and the rotation of the robot, and also the Pan-Tilt-Zoom camera. Note that these behaviors can be purely reactive, can have internal states, can be influenced by emotions or the *Rational* module, or can provide information to the *Rational* module.

Recommendation Level

For behavior recommendation, the *Implicit* module continuously recommends the activation of *Safe-Velocity*, *Random-Turn*, *Direct*, *Avoid*, *Unstall* and *Control-Display*. This makes the robot wander safely around in the environment, allowing to respond to commands associated with a symbol detected whenever appropriate (according to the plan states) and to interact with people. The *Egoistical* module derives recommendations based on

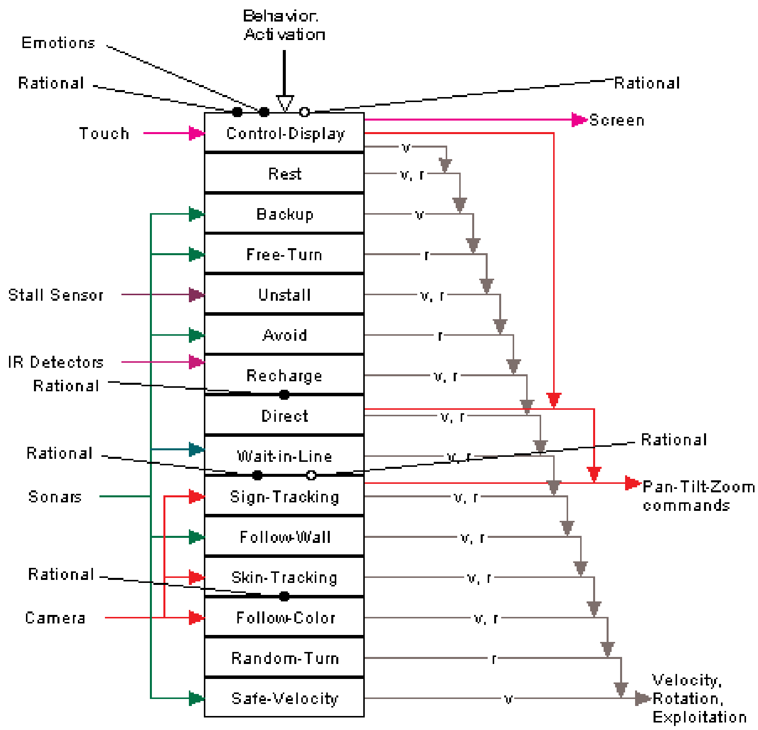

Figure 2: Behaviors used to control the robot.

the following motives: *Energize* for the behaviors *Recharge* and *Wall-Following*, *Distress* for *Free-Turn*, *Backup* or *Rest* depending on the activation level of the motive, and *Schmooze* for *Skin-Tracking*.

The most complex recommendation module for these experimentations is the *Rational* module because it determines behavior recommendations based on the different phases the robot has to go through for the Challenge. As shown in Figure 3, it makes recommendations based on symbol recognition, human-robot interactions, and a manager.

Symbol recognition. Our symbol recognition technique works in four steps: 1) image segmentation using colors, 2) positioning, 3) features extraction of color segments and 4) symbol identification using artificial neural networks [9]. Image segmentation is achieved using 32 colors and commodity hardware [10] in RGB space. Symbol perception is done by looking for a black blob completely surrounded by an orange background. Each recognizable symbol is contained in one segment, i.e., all pixels of the same color representing the symbol must be connected together to avoid recombination of boundary boxes. The positioning phase consists of stopping the robot and positioning the Pan-Tilt-Zoom camera to get the image with maximum resolution of the symbol perceived. Part of the 320×240 image delimitated by the orange region is scaled down to a 13×9 image. Symbol identification can then proceed, using three standard backpropagation neural networks. These networks differ by their number of hidden units (5, 6 and 7 respectively). A majority vote (2 out of 3) determines that the symbol is correctly

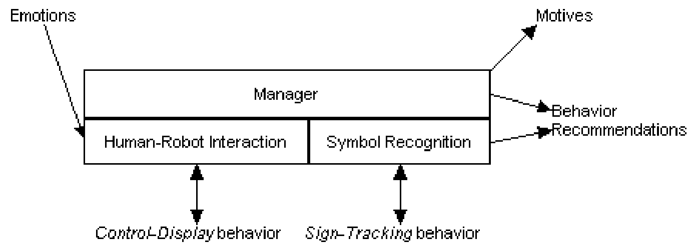

Figure 3: Decision modules used for *Rational* behavior recommendations.

recognized or not. A recognized symbol can then be processed by the manager, otherwise it is dismissed. In relation to our work, Adorni et al. [11] uses symbols (surrounded by a shape) with a map to confirm localization. But their approach uses shapes to detect a symbol, black and white images and no zoom.

Human-robot interaction. Human-robot interaction is done using the touch screen monitor to display messages and information about the internal states of the robot, and also get user's request using menu screens. By default, the face shown in Figure 4 is displayed on the screen. The face illustrates the emotional state of the robot by changing the orientation and the size of the mouth and of the eyebrows. Colors represent the difference between emotional pairs *Joy/Sadness* and *Anger/Fear*. The position and the size of the eyes changes according to the Pan-Tilt-Zoom controls. When a user touches the screen when the face is displayed, the *Control-Display* behavior makes the robot stop. The user then has three options: to make the robot guard or execute a specific rotation (to help position the robot), or to go to the main control panel. The main control panel is only used for debugging the different parts of the implemented strategy, by allowing us to explicitly change the state of the robot. Otherwise, other menu screens allow the robot to receive specific input from people according to the plan states. Group of sub-menus can be associated with a particular plan state, like generating the sequence of messages for receiving the instructions for the presentation at the registration desk. Finally, when the robot needs assistance, the *Control-Display* behavior displays a smaller version of the face with a message to request for assistance. The robot continues to move while displaying this message, until somebody touches the screen to provide the assistance requested.

Manager. The manager is a simple planner implemented as a finite-state machine, and determines the steps the robot as to go through for accomplishing the task. Transitions between plan states are derived from symbols recognized or inputs from the touch screen. For instance, when the robot is searching for the registration line and a **L** symbol is recognized, the *Rational* module evaluates the commands the robot has to do to position itself in the line, sends these commands to the *Direct* behavior, and changes the plan state to reflect that the robot is now waiting in line to register. Note that if a symbol is incorrectly recognized, the manager may change the plan state only when the symbol recognized is allowed in the current state. If this happens, unfortunately no self-recovery mechanism from such error is implemented. This is also true if the user does not give accurate information to

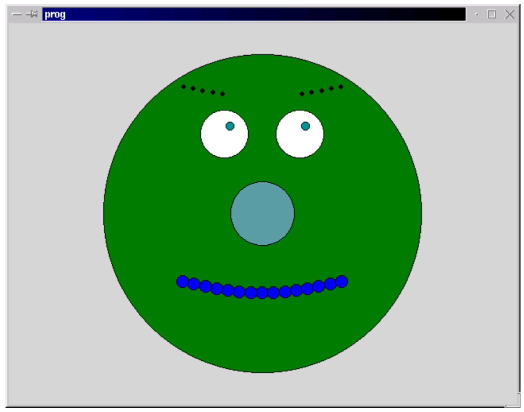

Figure 4: Face use to express the emotional state of the robot.

the robot. However, the user can explicitly set the appropriate plan state using the main control panel. The manager also includes a time manager to determine when the robot should change some of the plan states based on approximations of the time required to accomplish the different steps of the Challenge.

Motivational Level

Motives and emotions are the two elements of the architecture's motivational level.

Motives. Works on motivational systems [12, 5] demonstrates that a good balance between planning and reactivity for goal-management can be achieved using internal variables that get activated or inhibited by different factors. So in this implementation we use a similar approach for motives. Each motive is an activation variable associated with a particular goal of the robot. The activation level of a motive can influence other modules of the architecture and affect the overall behavior of the robot. A motive m has an energy level E and a mapping function M that determine its activation level A according to the formula: $A_m = M(E_m)$. The energy level and the activation level of a motive range between 0 and 100%. Factors that can influence the energy level of a motive are: sensory conditions, exploitation of behaviors associated with the motive, activation of other motives, *Rational* influences, emotions, and time. The energy level is computed by the equation $E_m = \sum_{j=1}^{n} w_j \cdot f_j$, where f represents n influencing factors, weighted by w (that must be predetermined). In these experiments, mapping from E to A can be directly transposed ($E = A$) or A is set to 100% when E reaches 100%, and to 0% when E reaches 0%.

For the challenge, the motives used are *Energize* to decide when and for how long recharging is required, *Distress* to monitor that the robot is able to move freely, *Schmooze* to interact with people when time permits, *Guard* to guard an area, and other motives representing the different plan states: *Registration*, *Taking the elevator*, *Navigate to the presentation room*, and *Presentation*. *Distress* is one motive that makes efficient use of the observation of behavior exploitation over time. For instance, long exploitation of *Avoid* is a sign that the robot is having difficulty moving freely in the environment. To prioritize the influences of motives in other modules of the architecture, we use Maslow's Hierarchy of Needs Theory [13] in which physiological needs are primary, security needs are secondary, followed by social and accomplishment needs.

Artificial emotion. Our objective is to design a generic emotional model that would work the same whatever the goals pursued by the robot, and not respond specifically to particular states and situations. To do so, the energy level of motives serves as an abstraction of the progression toward the accomplishment of the goals associated with activated motives. Monitoring the energy level of motives makes the approach generic, since the emotions can be used in different contexts (i.e., goals) according to the motives activated and their priority. This explains also why some motives also represents the different plan states for accomplishing the Challenge.

Our model is a two-dimension bipolar model with four emotions, *Joy/Sadness* and *Anger/Fear*, each defined from 0% to 100%. *Joy/Sadness* monitors a decrease/increase in the energy level of motives, indicating the presence/absence of progress toward the accomplishment of the goal associated with activated motives. *Anger/Fear* examines oscillations/constancy in the energy level, indicating difficult/no progress toward pursued goals. The amount of energy level and the priority of the motives also influence emotions. Evaluation of the emotions occurs concurrently and does not require any arbitration: the idea was to observe the validity of the model from facial expressions on the touch screen and for determining when the robot required assistance. For instance, when the robot takes too much time in trying to accomplish one of the phases of the Challenge (like finding the registration desk or the elevator), the energy level of the associated motive increases for so long that the *Sadness* emotion gets fully activated. This suggests that the robot is unable to accomplish its goal, and that it should try to get some assistance (by displaying a message, as described in the last section). If this does not work, the motive pursued is deactivated, and the robot just stops trying to participate in the conference. This mechanism allows the robot to recover from failure in the task, but allows it to continue to navigate safely in the environment and ensure its energetic needs. The same thing happens if the robot tries to find the charging station, but instead the robot just stops and requests to be recharged when its level of *Sadness* reaches 100%.

What Happened at the Challenge

We started to work on our robot entry seven months before the contest, and the use of the architectural methodology facilitated greatly the integration of the decision-making mechanisms implemented. The hardest part revealed to be the adjustments of the weighting factors of the motives, which required a lot of fine-tuning. By trial and error, we set these weights to allow sufficient time to accomplish the associated goal. Even though it was a difficult task, it was obviously not an impossible one.

Necessarily, like many other teams, we were still making adjustments to our code at the conference. But we were able to get all of the different phases of the strategy working together correctly, exactly as planned: the robot recognized symbols to guide itself in the convention center; moved around in the crowd and asked for assistance when it seems lost or for taking the elevator; recognized dignitaries from their badge color; interacted with people using a touch screen interface; went to the specified conference room; gave a short presentation in HTML about the whole experience; and recharged itself whenever necessary. Figure 5 shows our robot close to the registration desk, reading a symbol. We made several successful tests in the exhibition hall, in a constrained area and with constant illumination condition. We also ran two complete trials in the convention center, with people in the actual setting for the registration, the elevator and the conference rooms. Here are some observations made from these two trials:

Figure 5: Lollita Hall near the registration desk.

- The robot was able to identify symbols correctly in real life settings. Identification performance was around 83 % (i.e., 17 % of the images used for symbol identification were identified as unrecognizable), with no symbol incorrectly identified. The most difficult part was to adjust orange and black color segmentation for different illumination conditions in the convention center: some areas were illuminated by the sun, while others (like the entrance of the elevator) were very dark. Even though the superposition of two color regions for the localization of a symbol gave more robustness to the approach, it was still difficult to find the appropriate segmentation that worked in so diverse illumination conditions. So in some cases like for the elevator, we slightly change the vertical angle of the symbol to get more illumination.
- The robot was able to take the elevator correctly, following the guidelines described in the strategy. Our strategy for entering and for leaving the elevator revealed to be adequate, but obviously relies on the information provided by the user and the adequate positioning of the symbol in front of the elevator doors.
- Consideration of the time required by the robot for going through the different phases of the Challenge has a direct impact on the energy level of the robot. With the limited amount of knowledge exploited by the robot about its environment, it took a relatively long time to make the robot go through the phases

of the Challenge from start to end, compare to a human for instance. It all depended on the number of symbols encountered to go to a particular place and on interactions with people. For instance, it took 8 minutes for the robot to go from the entrance of the convention center to the registration desk. Note that we helped the robot going toward symbols when it was going in a direction with no indications. Also, it took 5 minutes to take the elevator, considering that the robot followed someone to guide it toward the **E** symbol for the elevator. Our two trials lasted around one hour and a half each, and considering that we did some color adjustments before the trials, it was inevitable that the robot needed to recharge along the way. Our charging station revealed to be greatly valuable in that regard, as for the ability of the robot to autonomously decide when it needed to recharge. Even during fine tuning of our approach the robot surprised us by suddenly starting to back up when we expected it to read a symbol, to then understand that the robot sensed the charging station at its back and needed to recharge.

- Our approach to position the camera to get an image of the face and upper body of a person during schmoozing needs to be improved, since it confused arms with legs.

- In regard to the facial expression of the robot, the color representation is difficult to interpret. Playing with the mouth and the eyebrows was sufficient to determine the internal emotional state of the robot.

Benefits and Challenges of Integrated Intelligence

Overall, we are very satisfied by the robot performance since it achieved appropriately the planned strategy. The actual HTML report of the second trial in the convention center is available on the web (URL: http://www.gel.usherb.ca/laborius/AAAI2000). Even though our robot still requires much more capabilities to accomplish the task completely autonomously, we demonstrated that the Robot Challenge can be addressed as a whole. The Challenge is not a ranked event, but we are the first and only team that has presented a robot completing the task from start to end. In doing so, we found that the integration of autonomous navigation, symbol recognition, goal management, recharge, human-robot interaction and the HTML report gave a real sense that the robot is autonomous in its decision making for accomplishing the Challenge, running for hours and even helping us for debugging it by using the main console panel and the HTML report.

In fact, integration of different mechanisms required for intelligent decision making is fundamental to the design of autonomous systems that have to work in real life settings. The architectural methodology followed in our design contributed greatly in that regard, making a step toward combining behavior-based concepts with planning capabilities, motivation and emotion. In addition, the architecture facilitates the addition of useful mechanisms to improve the decision process of the robot. Other capabilities like the use of a speech interface, a generic planner, a map and a representation of the environment that would allow to avoid exploring the same area repetitively could be useful. These capabilities are mostly based on the use of knowledge about the environment and would be implemented in the *Rational* module of the architecture. The current implementation is surely a nice starting point for such projects, and the actual capabilities can make interesting additions to these new ones. For instance, coupling the cues derived from symbols detected in the environment with a map would surely improve the robot's localization capabilities. While we can anticipate that a robot like a Pioneer 2 will always have inherent limitations that makes it less efficient than humans in a task such as the Challenge, the objective is to make an autonomous robot do as much as it can with its capabilities. Such integration will certainly be a step toward some form of intelligent decision making required for robots operating in real life settings.

Acknowledgments

Research conducted by LABORIUS is supported financially by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canadian Foundation for Innovation (CFI) and the Fonds pour la Formation de Chercheurs et l'Aide à la Recherche (FCAR) of Québec. The authors would like to thank the reviewers for their helpful comments.

References

[1] R. C. Arkin, *Behavior-Based Robotics*, The MIT Press, 1998.

[2] F. Michaud, G. Lachiver, and C. T. Le Dinh, "Architectural methodology based on intentional configuration of behaviors," *Computational Intelligence*, vol. 17, no. 1, pp. 132–156, 2001.

[3] C.Y. Suen, C.C. Tappert, and T. Wakahara, "The state of the art in on-line handwriting recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 8, pp. 341–348, 1990.

[4] J. D. Velásquez, "A computational framework for emotion-based control," in *Workshop on Grounding Emotions in Adaptive Systems, Conference on Simulation of Adaptive Behavior*, 1998.

[5] C. Breazeal, "Infant-like social interactions between a robot and a human caretaker," *Adaptive Behavior*, 1998.

[6] R. Plutchik, "A general psychoevolutionary theory of emotion," in *Emotion: Theory, Research, and Experience*, vol. 1, chapter 1, pp. 3–33. Academic Press, 1980.

[7] F. Michaud, P. Prijanian, J. Audet, and D. Létourneau, "Artificial emotion and social robotics," in *Distributed Autonomous Robotic Systems*, Springer, 2000.

[8] R. A. Brooks, "A robust layered control system for a mobile robot," *IEEE Journal of Robotics and Automation*, vol. RA-2, no. 1, pp. 14–23, 1986.

[9] F. Michaud and D. Létourneau, "Mobile robot that can read symbols," in *Proceedings IEEE International Symposium on Computational Intelligence in Robotics and Automation*, 2001.

[10] J. Bruce, T. Balch, and M. Veloso, "Fast color image segmentation using commodity hardware," in *Workshop on Interactive Robotics and Entertainment*, 2000.

[11] G. Adorni, G. Destri, M. Mordonini, and F. Zanichelli, "Robot self-localization by means of vision," in *Proceedings EUROBOT*, 1996, pp. 160–165.

[12] P. Maes, "Situated agents can have goals," in *Designing Autonomous Agents: Theory and Practive form Biology to Engineering and Back*, pp. 49–70. The MIT Press, 1990.

[13] A. H. Maslow, *Motivation and Personality*, Harper and Row, 1954.