Training Robots for Tactical Navigation Using Spatial Analysis

Marjorie Skubic, James Keller and Pascal Matsakis
Department of Computer Engineering and Computer Science
University of Missouri-Columbia
skubic@cecs.missouri.edu

Motivation and Overview of the Interactive Robot Training Approach

Programmability distinguishes robots from other forms of automation, but the actual programming process is difficult and has limited robots from achieving wide use. To address this problem, we have been investigating new methodologies for transferring skills to robots – namely, training them through interactions with human trainers in much the same way that human workers are trained for purposeful tasks (e.g., showing, guiding, or using positive and negative reinforcement). In particular, we have been addressing robot training for tactical navigation, which provides a means of achieving purposeful, task-oriented motion. For this investigation, we assume that the environment is known qualitatively but that we do not have a precise model. That is, we know whether it is indoors or outdoors, in a house or an office or a hospital, on city streets or in a dirt field. Qualitative information allows us to make assumptions about the type of objects found in the environment although we do not assume knowledge of their location. Our goal is to facilitate the use of robots as semi-autonomous tools in remote or dangerous environments. In particular, the goal is to develop robot training methods that allow domain experts to define their own task use of robots as intelligent tools. We consider the users to be experts in the task complexities, but they are not expert robotics engineers.

A recent approach to this problem has been to use a human operator to explicitly demonstrate a behavior or task (e.g., through teleoperation) and teach the robot how to imitate the demonstration (also called programming by demonstration - PbD). From our previous work in robot PbD [Skubic97,98, Chronis00], we have concluded that the following inherent problems exist and must be overcome for successful implementations: (1) variable delays in response time of the human demonstrator, (2) inconsistent and sometimes unintentional responses of the demonstrator, (3) uncertainties in measurement, and (4) inadequate interfaces and demonstration platforms. These problems aggravate the skill learning process because training data often has inconsistent mappings of sensory input to actuator commands. In addition, past PbD approaches have put an exceptional pressure on the demonstrator to produce numerous and/or perfect demonstrations, resulting in an impractical skill acquisition method. In order for this method to be feasible, we must consider the human trainer and provide an interface that facilitates quick and easy skill transfer and allows the user to know immediately whether the robot is learning the intended behavior or task.

Our approach, as outlined in Figure 1, utilizes two-way communications between the trainer and the robot. Derived from an interactive teaching method for human students [Collins82], the IRT framework relies on a sensor-based qualitative state (QS) which provides context information and acts as a symbolic link between the human trainer and the robot. For each QS, the trainer directs the robot to achieve the desired behavior; sensor signals and commands are monitored and the robot extracts a skill model. During the training process, feedback is provided to the trainer in the form of the recognized QS. In addition, the trainer can also observe the skill model as it is being generated and see exactly what the robot is learning as the training session progresses.

As a pragmatic approach to achieving easy and intuitive, two-way communications, we have been developing an interface between a robot and a PDA such as a PalmPilot device. Commands are issued to the robot using the stylus on the small, handheld PDA via a wireless, radio connection. Returned information can be displayed to the human trainer in multiple forms including sensory data, perceived qualitative conditions or recognized objects. Our PDA interface was motivated both by the ineffectiveness of the mouse and keyboard interface and also the desire for a compact and unrestrictive interface. The stylus of the PDA provides an easier control method than a mouse or joystick, allowing for smoother trajectories; this not only makes control easier for the trainer but also allows us to capture more consistent training data, which makes the skill learning process easier.

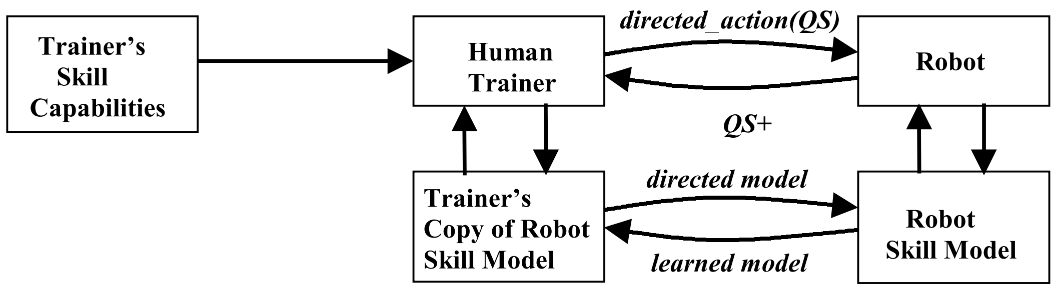

Figure 1. The Interactive Robot Training Framework, utilizing the sensor-based qualitative state (QS)

The skill acquisition and control architecture is shown in Figure 2. Learning can occur in each of the three main components shown in the figure. An overview of the components is listed below:

(1) Supervisory Controller: We consider procedural tasks (i.e., a sequence of steps) and represent task structure as a Finite State Automata (FSA) in the supervisory controller, following the formalism of the Discrete Event System (DES) [Ramadge89]. The FSA models behavior sequences that comprise a task; the sensor-based qualitative state (QS) is used for task segmentation. The change in QS is an event that corresponds to a change in the behavior. Thus, the trainer demonstrates a desired task as a sequence of behaviors using the existing behavior primitives and identifiable QS's, and the task structure is extracted in the form of the FSA. During the demonstration, the QS and the FSA is provided to the trainer to ensure that the robot is learning the desired task structure. With an appropriate set of QS's and primitive behaviors, the FSA and supervisory controller is straightforward. Also, this task structure is consistent with structure inherently used by humans for procedural tasks, making the connection easier for the human trainer. We have used this approach in learning force-based assembly skills from demonstration, where a qualitative contact state provided context [Skubic97,98]. For navigation tasks, spatial relations provide the QS context.

(2) State Classifier: The robot is provided with the ability to recognize a set of qualitative states, which can be extracted from sensory information, thus reflecting the current environmental condition or context. The initial set is preprogrammed or learned off-line. For navigation skills, robot-centered spatial relations provide context (e.g., there is an object to the *front-left*). Adding the ability to recognize classes of objects provides additional perception (e.g., there is a person to the *front-left*). In addition to the preprogrammed and prelearned QS's, the trainer can add to the set by showing examples of specific objects that can be used for context information (e.g., there is a person named Joe to the *front-left*). Qualitative spatial relations can also be used to analyze maps and desired robot trajectories and to extract appropriate robot behavior sequences. In the next section, we describe previous work on image spatial analysis and proposed extensions for tactical robot navigation.

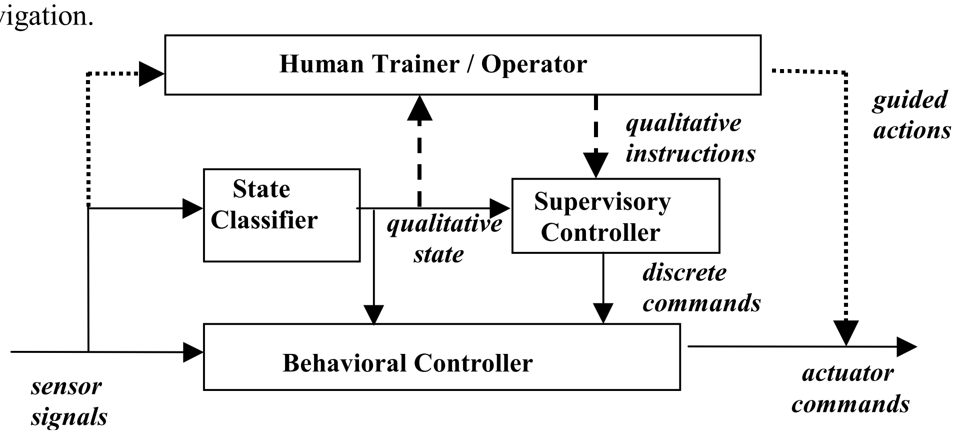

Figure 2. The Skill Acquisition and Control Architecture. Dotted lines show the links for training. The qualitative state and qualitative instructions can be used as low bandwidth control for on-line operations or for training.

(3) Behavioral Control: The robot is equipped with a set of primitive behaviors and behavior combinations. Some behaviors may be preprogrammed and some may be learned off-line using a form of unsupervised learning. The trainer can add to the set of behaviors by demonstrating new behaviors which the robot learns through supervised learning, thus allowing desired biases of the domain expert to be added to the skill set. We have investigated both supervised and unsupervised learning methods. For example, we have used evolutionary computing to learn a set of fuzzy rules for robot navigation, e.g., a composite set of potentially competing behaviors, as in heading for a target while avoiding obstacles [Chronis99]. We have also investigated supervised learning with neural networks, e.g., learning a nonlinear, sensor-to-actuator mapping for a corridor-following behavior using demonstrated training data [Chronis00]. Collecting consistent training data has been problematic for the reasons mentioned above and has in part motivated the PDA interface.

Methods and Current Status

The IRT framework is a modular architecture which separates the main perceptual and control functions. The key to making the interactive robot training work is the QS, especially in the following ways: (1) the ability to perceive an often ambiguous context based on sensory conditions, especially in terms that are understandable for the human trainer, (2) choosing the right set of QS's so as to communicate effectively with the trainer, and (3) the ability to perform self-assessment, as in knowing how well the QS is identified which helps in knowing when to get further instruction. For our work, we consider the QS to be a cluster in the sensor feature space, although identifying the features used for the clustering is not always obvious and clustering for identification is not necessarily warranted.

Previously, we have investigated mobile robot navigation behaviors (both pre-programmed and learned); to further the robot training objectives, the next key step is to focus on the state classifier. For identifying the QS, we will leverage previous work on extraction of spatial relationships from image data and extend this to spatial relations for robot navigation. Spatial relationships provide powerful cues for humans to make decisions; thus, it is plausible to investigate their use as a qualitative state for robot skills, as well as a linguistic link between the human and the robot.

Spatial Relations for Image Analysis

In our current work, we have been investigating basic spatial relationships among regions of a two-dimensional image. We have demonstrated in [Wang97] that spatial relations using fuzzy set definitions form powerful features for object recognition. In recent work funded by the Office of Naval Research [Wang99ab], one goal was to build fuzzy rule-based systems to let the computer provide linguistic scene descriptions that are understandable by humans. Since fuzzy rule-based systems simulate the human's way of reasoning, it is natural to use them in scene description. Perceptual responses from human subjects were collected and used to train a neural network, generating confidence values for five spatial relations (*left*, *above*, *right*, *below*, and *surround*). These five spatial relation values were then used as inputs to a fuzzy rule-based system for scene description.

Figure 3(a) shows the intensity component of a standard color outdoor scene. The image was automatically segmented and labeled using the methods presented in [Yang94, Keller94]. This supervised approach was based on fuzzy aggregation networks and actually produced a fuzzy segmentation and labeling of the image. The closest crisp partition is displayed in Figure 3(b). Spatial relation membership values were generated from the neural network approach [Wang99a] and this information was fed to the fuzzy rule base. After application of the rule-based system, outputs were automatically converted to linguistic labels, shown in Figure 3(c).

The description in Figure 3 is typical of those obtained in [Wang99b]. How do we judge the performance of high level scene description? We only argued from an intuitive basis. That is, do we believe that the description captured the *essence* of the scene? Appealing to intuition has been a standard means of justifying output of spatial relationship definitions, but a better method is warranted. One approach involves using this system to produce a description to match against one generated by another means, e.g., a person or a different description package. Preliminary data is encouraging [Keller99a]. However, the *language* used in

The roof is right of the tree.
The wall is right of the tree.
The sky is right of the tree.
The wall is below the roof.
The sky is above the roof.
The lawn is below the wall.
The road is below the wall.
The road is below the lawn.

(a) (b) (c)

Figure 3. (a) Intensity image of a color outdoor scene, (b) Automatic segmentation and labeling into 6 classes by fuzzy aggregation networks, (c) Linguistic output from the fuzzy rule-based system

[Wang99b] is very coarse and the range of individual opinions cannot be covered. In fact, most definitions of spatial relations are based on the sole notion of the histogram of angles, which provides a representation of the relative position of an object with regard to another. In [Matsakis99a], we introduced the notion of the histogram of forces. It generalizes and supersedes that of the histogram of angles. Using histograms of forces, many representations of the relative position between image objects can be obtained, and new definitions of spatial relations result [Matsakis99b]. They provide complementary opinions that can be combined in different ways to model differing human perceptions. They also allow the language presented in [Wang99b] to be refined by means of linguistic hedges (descriptive phrases) and a self-assessment, and also to be tailored to match individual human users. The linguistic descriptions generated in [Keller99b] are composed of the primary direction (e.g*., A is to the right of B*), the secondary direction that supplements the description (e.g., *but a little above*), and a self-assessment (e.g., *the description is satisfactory*). This indicates the extent it is necessary to turn to other spatial relations (e.g., *surrounds*). A simple set of rules has been implemented to test out and validate the descriptive nature of this approach.

Spatial Relations for Robot Navigation

The image spatial analysis work has provided the appropriate granularity for communicating with human observers. We plan to leverage this work, extending the spatial analysis for robot navigation. In particular, spatial relations can provide qualitative information for navigation. The QS classifier can be trained off-line; once trained, the actual execution is fast enough for real-time navigation.

There are three ways in which we propose to use spatial relations for robot navigation. First, egocentric spatial relations provide a robot-centered view of the environment. Consider a top view of the robot. A simple ring of sonar sensors provides a 360 degree proximity view around the robot, which can be converted into a description using spatial relations (e.g., there is an object on the *left*, or there is something in *front*, or combinations such as *front-left*, and including the relation *surrounded*.) Camera images provide a third dimensional view above or below the sonar sensors, capturing height, so that descriptions can include relations such as *above-front*. These spatial relations can then be used in fuzzy rules to generate more complex QS's, e.g., if a long object is on the left and a long object is on the right but nothing is in front or behind, then I am in a corridor. Using data accumulated over time can be used to handle sensor uncertainties and also to provide increased knowledge for spatial relations. Second, the image spatial analysis can be used to provide a direction relative to something in the image. For example, the trainer can issue instructions such as go to the *right* of the tree, or go *behind* the tree. Self-assessment is important here so the robot knows how well it is doing during navigation and also knows when it has reached the desired QS. Finally, spatial relations can be used to analyze maps and facilitate their use in communicating desired programs and navigation tasks. A map can either be generated with sensory information and displayed on the PDA, or the trainer can draw an approximate map directly on the PDA. In either case, the trainer then draws the desired robot trajectory also on the PDA, relative to the map. Spatial relations similar to those developed previously for image analysis can then be used to convert the relative trajectory to an appropriate navigation behavior sequence.

References

[Chronis99] G. Chronis, J. Keller and M. Skubic, "Learning Fuzzy Rules by Evolution for Mobile Agent Control", *Proceedings of the 1999 IEEE International Symposium on Computational Intelligence in Robotics and Automation*, Monterey, CA, Nov., 1999.

[Chronis00] G. Chronis and M. Skubic, "Neural Network Control of a Mobile Agent", technical report, University of Missouri-Columbia, in preparation.

[Collins82] A. Collins and A. Stevens, "Goals and strategies of inquiry teachers", In *Advances in Instructional Psychology*, R. Glaser, editor, chapter 2, pp. 65-119, Lawrence Erlbaum Associates, Hillsdale, NJ, 1982.

[Keller94] J. Keller, R. Krishnapuram, Z. Chen, Z., and O. Nasraoui, "Fuzzy Additive Hybrid Operators for Network-Based Decision Making", *International Journal of Intelligent Systems*, Vol. 9, No. 11, 1994, pp. 1001-1024.

[Keller99a] J. Keller, P. Gader and X. Wang, "LADAR Scene Description using Fuzzy Morphology and Rules", *Proceedings, IEEE Workshop on Computer Vision Beyond the Visible Spectrum: Methods and Applications,* in conjunction with CVPR '99, Fort Collins, CO, June, 1999, pp. 120-129.

[Keller99b] Keller J. M., and Matsakis, P., "Aspects of High Level Computer Vision Using Fuzzy Sets", *Proceedings of the 8th IEEE Int. Conf. on Fuzzy Systems*, Seoul, Korea, August 1999, pp. 847-852.

[Khab99a] M. Khabou, P. D. Gader, and J. M. Keller, "LADAR Target Detection Using Morphological Shared-Weight Neural Networks", *Machine Vision and Applications*, 1999.

[Matsakis99a] Matsakis, P., and Wendling, L., "A New Way to Represent the Relative Position between Areal Objects*", IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 21, no. 7, pp. 634-643, July 1999.

[Matsakis99b] Matsakis, P., Keller, J. M., and Wendling, L., "F-Histograms and Fuzzy Directional Spatial Relations", LFA'99 (French-Speaking Conference on Fuzzy Logic and Its Applications), Valenciennes, France, October 1999.

[Ramadge89] P.J. Ramadge and W.M. Wonham, "The control of discrete event systems", *Proceedings of the IEEE*, vol. 77, no. 1, pp. 81-97, Jan., 1989.

[Skubic97a] M. Skubic, "Transferring assembly skills to robots: Learning force sensory patterns and skills from human demonstration", Ph.D. Dissertation, Texas A&M University, College Station, TX, 1997.

[Skubic97b] M. Skubic, S. Castrianni and R. Volz, "Identifying Contact Formations from Force Signals: A Comparison of Fuzzy and Neural Network Classifiers", In *Proceedings of the 1997 IEEE International Conference on Neural Networks*, vol. 3, pp. 1623-1628, Houston, TX, June, 1997.

[Skubic98] M. Skubic and R. Volz, "Learning Force-Based Assembly Skills from Human Demonstration for Execution in Unstructured Environments", In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, Leuven, Belgium, May, 1998.

[Skubic99] M. Skubic and R. Volz, "Identifying Single-Ended Contact Formations from Force Sensor Patterns", *IEEE Transactions on Robotics and Automation*, accepted for publication.

[Wang97] X. Wang, J. Keller, and P. Gader, "Using spatial relationships as features in object recognition", *Proceedings of NAFIPS'97*, Syracuse, NY, September, 1997, pp. 160-165.

[Wang99a] X. Wang and J. Keller, "Human-based Spatial Relationship Generalization Through Neural/Fuzzy Approaches", *Fuzzy Sets and Systems*, Vol. 101, No. 1, 1999, pp. 5-20.

[Wang99b] X. Wang and J. Keller, "A Fuzzy Rule-based Approach for Scene Description Involving Spatial Relationships" *Computer Vision and Image Understanding*, revised, 1999.

[Won97] Y. Won, P. D. Gader, and P. C. Coffield, "Shared-Weight Neural Networks based on Mathematical Morphology with Applications to Automatic Target Recognition," *IEEE Trans. Neural Networks*, Vol. 8, No. 5, pp. 1195-1204, Sept. 1997.

[Yang94] H. Yang, "Fuzzy Aggregation Networks of Hybrid Units with Generalized Yager Operators and Their Applications to Function Approximation and Decision Making", University of Missouri-Columbia, Masters Thesis, 1994.