Human-to-Robot Skill Transfer Using the SPORE Approximation

Gregory Z. Grudic

Dept. of Electrical Engineering
2356 Main Mall
University of British Columbia
Vancouver, British Columbia
Canada, V6T 1Z4
FAX: (604) 822-9209
gregg@ee.ubc.ca

Peter D. Lawrence

Dept. of Electrical Engineering
2356 Main Mall
University of British Columbia
Vancouver, British Columbia
Canada, V6T 1Z4
FAX: (604) 822-9209
peterl@ee.ubc.ca

Abstract

Humans are capable of accomplishing a variety of complex tasks in environments which have continuously changing characteristics. Robots, however, generally require highly structured environments which allow the programmer to make a large number of assumptions, making the robot's task easier to encode. We propose a framework for programming robotic tasks using human-to-robot skill transfer. We assume that there exists a human expert who can accomplish a task in an unstructured environment by using various sensor displays and controls. The human expert performs the desired task a number of times while her/his input/output pairs are being recorded by the robot. The robot then uses this recorded data to construct a mapping between these sensor inputs and actuator outputs. This mapping must be general enough to allow the robot to accomplish the same task, in similar but not identical, dynamic, unstructured, environments. This paper presents a testbed for human-to-robot skill transfer which is based on the teleoperated control of a small mobile robot working in an unstructured environment. The skill which is transferred from human-to-robot is loosely based on the tree tending task, a task which was chosen for its inherently unstructured nature. The SPORE approximation is proposed as a means for creating the robot's mapping from sensor inputs to actuator outputs.

1 Introduction

We approach the problem of programming robots that can work in unstructured environments from the point of view of human-to-robot skill transfer. In this paradigm, the human transfers his or her skills to a robot by giving the robot a finite number of examples of how the desired manipulation task is done. The key to the success of this type of robot programming is that the robot be able to generalize the desired task, given only the human experts examples, even when the unstructured environments change after learning has occurred.

Other researchers have had success with similar approaches to the robot programming problem. We mention only a few examples here. Perhaps the most related work is that of Pomerleau [6] [5] where he successfully uses a 3 layer perceptron network to control the CMU ALVINN autonomous driving system as it drives along a road. Pomerleau's approach differs from that presented here in that 1) we use a nonparametric approximation approach and 2) the human expert supplying the training data uses only the sensor inputs being fed to the approximation (Pomerleau does not use this restriction). However, as with the approach presented here, Pomerleau uses a large number of sensor inputs (960) to effectively generalize the desired task.

Other related work includes that of Kosuge *et al* [3] where *a priori* knowledge of the task is incorporated with human generated examples to transfer a skill. Other researchers use neural networks to encode the human's skills: for the deburring task Shimokura and Liu [8], and for compliance control Asada [1]. Lui and Asada [4] use process dynamics for transferring manipulative skills. Yang *et al* [9] use a multidimensional Hidden Markov Model to represent the transferred skill. Rouse *et al* [7] also consider symbol processing models as useful tools for capturing human skills and knowledge.

The approach to human-to-robot skill transfer proposed in this paper assumes the following:

> The human operator can accomplish the desired manipulation task by using only the robot's sensors to obtain information about the world, and using only the robot's actuators to manipulate world objects. Thus, the human has only the robot's perspective on gathering world information and manipulating world objects, and no other. This implies that if the human can accomplish the task, then the robot should also be able to accomplish it.

This assumption has at least one important implication for the choice of a learning system, or approximation framework, for generating the sensor to actuator mappings. Humans are usually only experts at manipulation tasks which they can anthropomorphize to their natural frame of reference. In other words, we are accustomed to particular human sensory and action modalities, and as the information for the task becomes more sparse, the task becomes more difficult. As a result, when a human expert is demonstrating a task which is sufficiently complex, it is likely that he or she is using a large number of sensor inputs of various types, including both visual and tactile. Thus, the type of human-to-robot skill transfer proposed here requires an approximation framework which is highly flexible with respect to the number and variety of sensor inputs, and yet is practical to implement.

In this paper, we briefly describe a nonparametric approximation framework (SPORE) which we propose to use in human-to-robot skill transfer. We then describe a human-to-robot skill transfer testbed which we have implemented. Finally, we report the results of our first skill transfer experiment done in an unstructured environment using the proposed approximation. The skill which we have chosen to demonstrate is loosely based on the tree tending task, and was chosen due to its inherently unstructured characteristics. However, the framework we are in the process of formulating should be generalizable to a large variety of human skills.

With respect to the proposed approximation framework, the main goal of this paper is to answer the following questions:

1. Our intention is to feed the sensor inputs used by the human operator directly to the SPORE approximation, with no intermediate signal/image processes. The question then becomes: given only a small number of examples of the skill being transferred, how well does the SPORE approximation generalize this skill beyond the conditions of the training set?
2. The SPORE approximation is constructed using training data which takes a specific amount of memory to store. Due to the nonparametric nature of the SPORE approximation, the size (in terms of bytes) of the learned approximation is directly dependent on the training data (i.e. it's nonlinearity or complexity). Thus the question then becomes: what is the *compression ratio* of the approximation? For our purposes the compression is defined as follows:

$$compression\ ratio = \frac{size\ of\ approximation\ (bytes)}{size\ of\ training\ data\ (bytes)} \tag{1}$$

In answering the above questions we hope to gain some insight into the efficacy of the SPORE approximation for human-to-robot skill transfer.

2 The Spore Approximation Framework

The sensor to actuator mappings are represented using a new function approximation framework which is here referred to as the **S**pace **P**artitioning, self-**O**rganizing and dimensionality **RE**ducing, or *SPORE* approximation [2]. The SPORE approximation represents functions by summing cascades of 2 dimensional functions. The approximation takes the following form: let $f(x_0, ..., x_{N-1})$ be an N dimensional function, then the SPORE approximation of f is given by (see Figure 1 for a 4 dimensional example of this structure):

$$\begin{aligned}\hat{f}(x_0, ..., x_{N-1}) = g_{\rho_1\sigma_1}(h_{\rho_1\sigma_1}(x_0, x_1), x_2) + \\ g_{\rho_1\sigma_1\rho_2\sigma_2}(h_{\rho_1\sigma_1\rho_2\sigma_2}(h_{\rho_1\sigma_1}, x_2), x_{((3)\mathrm{mod}(N))}) + \\ \textstyle\sum_{l=3}^{l_{max}} g_{\rho_1\sigma_1\ldots\rho_l\sigma_l}(h_{\rho_1\sigma_1\ldots\rho_l\sigma_l}(\\ h_{\rho_1\sigma_1\ldots\rho_{l-1}\sigma_{l-1}}, x_{((l)\mathrm{mod}(N))}), x_{((l+1)\mathrm{mod}(N))}),\end{aligned} \tag{2}$$

where,

1. $A \bmod N$ is the integer remainder (modulus) of $(A \div N)$, and is used as an index to cycle through the independent variables of $f(x_0, ..., x_{N-1})$;
2. the subscript notation $\rho_1\sigma_1 ... \rho_l\sigma_l$ indexes constraints combined to form disjoint regions (referred to as *branches*) of the input space. It

serves to identify the two dimensional functions $g_{\rho_1\sigma_1\ldots\rho_l\sigma_l}$ and $h_{\rho_1\sigma_1\ldots\rho_l\sigma_l}$ which approximate the target function, f, on the specific sub-domain of the input space referred to by $\rho_1\sigma_1 \ldots \rho_l\sigma_l$;

3. l_{max} is an integer defining the maximum number of terms required in order to achieve some bounded approximation error $\max_{\mathbf{x}\in\mathcal{D}} |\hat{f}(\mathbf{x}) - f(\mathbf{x})| < \epsilon$ where ϵ is an arbitrary, small, positive, real number, and $\mathcal{D}$ is the domain of the function f.

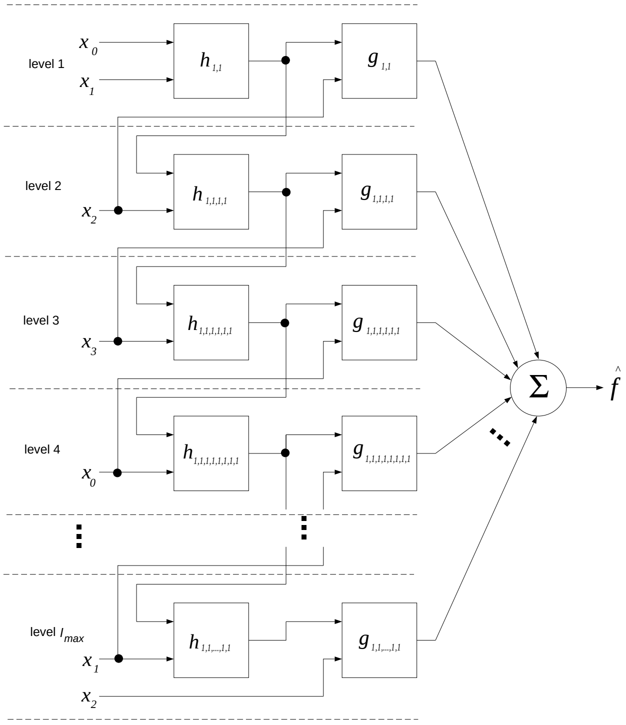

Figure 1: 4 Dimensional Example of the SPORE Structure Along a Single Branch

The SPORE approximation is constructed one term at a time until the approximation error has been reduced to the desired level. The number of terms constructed by the learning algorithm is not chosen *a priori*, but is determined as learning progresses, making the SPORE approximation nonparametric. Note that each term requires the approximation of only a pair of 2 dimensional functions ($g_{\rho_1\sigma_1\ldots\rho_l\sigma_l}$ and $h_{\rho_1\sigma_1\ldots\rho_l\sigma_l}$), which is one reason why SPORE is tractable, even for large N.

Beginning with the entire input domain, the sequences of ρ's and σ's represent constraints which branch in a tree-like structure defining disjoint subdomains. Hence each level of the tree represents a partitioning of the input space $\mathcal{D}$. These constraints are constructed such that in each sub-domain or branch (referred to by the subscript notation $\rho_1\sigma_1 \ldots \rho_l\sigma_l$), equation (2) gives an adequate approximation of the target function. This domain subdivision is not chosen *a priori*, but is controlled by the learning algorithm as learning progresses.

The SPORE approximation framework has a number of characteristics which make it a potentially good candidate for representing the high dimensional functions required in human-to-robot skill transfer. These characteristics include the following [2]:

1. The complexity of evaluating the SPORE approximation for any one input vector $(x_0, \ldots, x_{N-1})$ is given by $2l_{max} \cdot C_{2D}$, where C_{2D} is the complexity of evaluating a 2 dimensional function. For the example given in Section 4, we chose to use third order polynomials to represent the 2 dimensional functions, and hence C_{2D} is the complexity of evaluating a third order, 2 dimensional polynomial. This choice of a 2 dimensional function representation is arbitrary and does not affect the theoretical convergence properties of the SPORE learning algorithm. However, it does affect the number of terms required to obtain a sufficient approximation, as well as the extent of domain subdivision generated by the learning algorithm.

2. The SPORE approximation is able to represent any bounded continuous N dimensional function ($N > 3$), to any arbitrary accuracy. The learning algorithm associated with the SPORE approximation has the following characteristics:
   (a) An algorithm exists which is theoretically guaranteed to construct a uniformly convergent minimax SPORE approximation for any bounded continuous N dimensional function ($N > 3$). The approximation formed by the algorithm is nonparametric.
   (b) If $f(x)$ is bounded, continuous, and at least once differentiable, then the rate of convergence of the learning algorithm (defined as the number of sample points required to achieve a given level of approximation) for the construction of each branch of equation (2), is independent of dimensionality, and is given by $\|f(\mathbf{x}) - \hat{f}_n(\mathbf{x})\|_\infty < O((cn)^{-1} \log cn)^r$, where $\| \cdot \|_\infty$ is the L^∞ norm; f is the function being approximated

and $\hat{f}_n$ is the SPORE approximation based on n random samples of f; $c > 0$ is some constant independent of dimension N; and, $r = p/(2p+3)$, where p is the number of times that f can be differentiated.

(c) The complexity of learning along one branch of the SPORE approximation to level l_{max} is given by $l_{max} \cdot (C_g + C_h)$, where C_g is the complexity of learning a $g_{\rho_1\sigma_1\ldots\rho_l\sigma_l}$ function and C_h is the complexity of learning an $h_{\rho_1\sigma_1\ldots\rho_l\sigma_l}$ function. For the example given in Section 4, we chose to use third order polynomials to represent the 2 dimensional functions, and hence C_g and C_h are defined by the complexity of learning a third order, 2 dimensional polynomial.

3 The Experimental Testbed

A small scale proof-of-concept testbed for human-to-robot skill transfer has been implemented (see Figure 2), consisting of a tracked mobile robot which is approximately 35 centimeters long and 23 centimeters wide. The robot is equipped with 1) left and right tracks, each with independent proportional forward/backward motion similar to that found in forestry excavators; 2) a gripper (not shown in Figure 2) to allow it to grasp objects; and 3) an NTSC color camera. The leftmost image in Figure 2 shows the robot's workspace and the rightmost image shows a closeup of the mobile robot. Initially this robot is intended to operate in a flat 2.5 meter by 1.8 meter workspace which has randomly placed obstacles and grippable objects. The intended task of the robot is to manipulate the target objects (while avoiding the obstacles) in a manner defined by a human operator who first accomplishes the desired task a number of times via teleoperation. The human operator observes images received from the camera located on the robot (at a frame rate of approximately 20 to 30 frames per second), and operates two joysticks which control the robot's left/right forward/backward motion, as well as the open/close gripper motion. A SPARCstation 20, equipped with a frame grabber, is used to display the on-board camera image to the human operator as he or she demonstrates the desired task. The learning and subsequent autonomous control of the robot are also performed by the SPARCstation 20.

As shown in Figure 2, the testbed currently uses model "trees" as objects which can be both manipulated and avoided. The "trees" are currently of three different types (tall spruce, short spruce, and short birch), and our goal is to teach the robot to differentiate between them. This testbed is considered unstructured because 1) although there are only 3 different types of trees, no two trees of the same type are identical either in height, shading of green, or in texture, making manual encoding of a model difficult; and, 2) the image background of the trees is a robotics lab which is continually changing as equipment is moved around and people move throughout the lab.

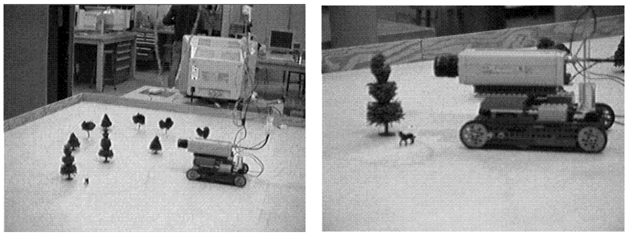

Figure 2: The Human to Robot Skill Transfer Testbed

4 Experimental Results and Discussion

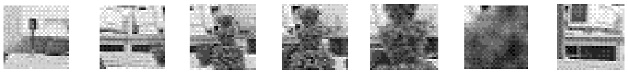

Figure 3: "Tree Tending Skill" — select a tree of a certain size and species for inspection.

The first skill which has been successfully transferred from human to robot is described in Figure 3. This skill is loosely related to the "tree tending skill," which requires the selection of a certain size and species of tree for inspection. The skill involves turning the robot counterclockwise in place until at least one tall spruce is within the camera's field of view, and then maneuvering the robot into a position where the tree can be inspected. Figure 3 shows several typical camera views which the operator sees, as the task of locating and moving towards a tall spruce is performed. Starting at the leftmost image of Figure 3, the robot is turning counterclockwise searching for a tree; in the third image a tall spruce has been found and the robot begins to move towards it; in the sixth image the spruce is within inspecting distance and the robot turns off to the left in search of another tall spruce.

Note that this first task does not involve avoiding obstacles or actually stopping to inspect the tree, it simply requires the robot to find and move towards a tall spruce tree in the visually cluttered environment of the robotics lab.

The image which the human operator, as well as the robot, observes while executing the desired task, consists of 1024 pixels derived by sampling a 32 pixel by 32 pixel region in the center of the original NTSC image (see Figure 3 for gray-scale examples). The sampled pixel resolution is 8 bits, 3 bits for red, 3 bits for green, and 2 bits for blue. There is no image processing done prior to feeding these pixel values to the SPORE approximation; what the human sees is exactly what is passed to the approximator. In order to generalize the desired task the robot must approximate 2 mappings, one to control the robot's forward and backward motion, and one to control the robot's left and right motion. Both of these mappings have 1024 inputs and 1 output.

The training data generated by the human operator consisted of demonstrating the above skill a total of 4 times, from 4 random starting positions of the robot, and 4 random positions of the spruce. Due to image resolution limits, the distance between the robot and the "tree" was always 1.5 meters or less. It took approximately 2 minutes (equivalent to 2180 camera frames or input/output examples) for the human operator to do these 4 demonstrations. In the SPORE approximation constructed based on this training data, the 2 dimensional functions $g_{\rho_1\sigma_1 \dots \rho_l\sigma_l}$ and $h_{\rho_1\sigma_1 \dots \rho_l\sigma_l}$ were modeled using third order polynomials. The SPORE approximation consisted of 1139 terms (see equation (2)) for each of the 2 mappings generated, and no domain subdivision occurred. The length of time to evaluate the resulting approximation for one instance of sensor inputs is 191 milliseconds.

Next consider the two questions posed at the end of Section 1. The first concerns the ability of the SPORE approximation to generalize given only a few examples of the desired skill, and raw sensor inputs used by the human. In the case of the above experiment, the human operator generated only 4 examples of the desired task, and the SPORE approximation was fed raw pixel data with no intensity compensation or any other type of signal preprocessing. In addition, one can see from Figure 3 that the image generated by the frame grabber inside the SPARCstation 20 is extremely noisy. Nonetheless, when we ran 30 experiments where the robot used the SPORE approximation to autonomously accomplish the transferred skill, the robot never failed to locate and move towards the tree. These 30 experiments were done with random placements of tree and robot (as with the training set, the spruce was always within 1.5 meters of the robot). It is interesting to note that the robot does not approach, for instance, upside-down spruce trees, and does not seem to be solely color sensitive; the shape of the tree and its texture are significant as well. We may therefore conclude that to the extent that the robot acquired the desired skill, it was able to execute it within an unstructured environment. From the experimental results we note that the success of the SPORE approximation in generalizing the task, given relatively few examples and raw pixel inputs, indicates that this approximation framework may indeed be a useful tool in human-to-robot skill transfer.

The second question posed at the end of Section 1 is concerned with the compression ratio of the SPORE approximation. The amount of memory needed to store the learned approximation is 182,240 bytes (22,780 double precision parameters). The original training set consisted of 2180 images each requiring 1024 bytes to store. In addition, the human experts actuator control outputs were recorded for each image, with each output consisting of 2 double precision voltage values. Thus, the training samples required 2,267,200 bytes, giving an approximation compression ratio (see equation (1)) of 0.08, while generating significant generalization. This means that the SPORE approximation compressed the training data by an order of magnitude. Compression in nonparametric approximation is an important consideration because if the approximation is as large or larger than the original set of training data, then the approximation is probably not generalizing noisy training data appropriately.

5 Conclusion

We have demonstrated human-to-robot skill transfer for the task of locating and approaching a model tree in a visually cluttered environment. We hypothesize that the SPORE approximation is a good tool for human-to-robot skill transfer because 1) it gives good generalization of a task done in an unstructured environment given only a small number of examples of the task (four examples in the case of the task described here), and, 2) the constructed approximation potentially takes up much less storage space than the original training set (in the case of the task presented here, the reduction was an order of a magnitude). We are currently in the process of running additional experiments in human-to-robot skill transfer. Specifically, we are interested in studying the efficacy of the SPORE approximation as the the number of sensor inputs is increased, thus giving the human operator the capacity to do more complicated tasks; and, as the number of training examples fed to the learning algorithm is increased, making the transfer of more complex tasks feasible.

Acknowledgments

We would like to thank Jane Mulligan for her extensive proof reading and clarification the concepts presented, and Dan Chan for building the testbed.

References

[1] H. Asada. Teaching and learning of compliance using neural nets: Representation and generation of nonlinear compliance. In *Proceedings of the 1990 IEEE International Conference on Robotics and Automation*, pages 1237–1244. IEEE, 1990.

[2] G. Z. Grudic and P. D. Lawrence. Uniformly convergent minimax approximation of multidimensional functions. *In Preparation*, 1995.

[3] K. Kosuge, T. Fukuda, and H. Asada. Aquisition of human skills for robotic systems. In *Proceedings of the 1991 IEEE International Symposium on Intelligent Control*, pages 469–474. IEEE, 1991.

[4] S. Lui and H. Asada. Transfering manipulation skills to robots: Representation and aquisition of manipulation skills using process dynamics. *J. Dynamic Syst. Measur. Contr.*, 114(2):220–228, 1992.

[5] D. A. Pomerleau. Input reconstruction reliability estimation. In *NIPS5*, pages 279–286. Morgan Kaufmann Pub., 1993.

[6] D. A. Pomerleau. *Neural Network Perception for Mobile Robot Guidance.* Kluwer Academic Publishers, Boston/Dordrecht/London, 1993.

[7] W. B. Rouse, J. M. Hammer, and C. M. Lewis. On capturing human skills and knowledge: Algorithmic approaches to model identification. *IEEE Trans. Syst. Man Cybernet.*, 19(3):558–573, May/June 1989.

[8] K. Shimokura and S. Liu. Programming deburrring robots based on human demonstration with direct burr size measurements. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pages 469–474. IEEE, 1991.

[9] Jie Yang, Yangsheng Xu, and Chiou S. Chen. Hidden markov model approach to skill learning and its application to telerobotics. *IEEE Trans. Robotics and Automat.*, 10(5):421–631, October 1994.