# Human-to-Robot Skill Transfer Via Teleoperation

G. Z. Grudic

Dept. of Electrical Engineering
University of British Columbia
2356 Main Mall
Vancouver, British Columbia
Canada, V6T 1Z4
gregg@ee.ubc.ca

P. D. Lawrence

Dept. of Electrical Engineering
University of British Columbia
2356 Main Mall
Vancouver, British Columbia
Canada, V6T 1Z4
peterl@ee.ubc.ca

## Abstract

*Programming robots to work in unstructured, changing environments, has proven to be difficult. Humans, however, effectively function in such environments. We propose a framework for programming robotic tasks using teleoperation based human-to-robot skill transfer. We assume that there exists a human expert who can accomplish a task in an unstructured environment solely through teleoperation based feedback. The human expert performs the desired task a number of times while her/his input/output pairs are being recorded. This recorded data is used to construct a mapping between these sensor inputs and actuator outputs. The mapping is then used to autonomously control the robot. In this paper we propose a set of characteristics which a human-to-robot skill transfer system should have. We then summarize the system we have implemented and present the results of some experiments we have done in skill transfer in unstructured environments.*

## 1 Introduction

The teleoperation of heavy machines such as excavators, log loaders and feller-bunchers is becoming technically feasible with the use of a computer-assisted human guidance system [9]. Through teleoperation, humans are capable of efficiently performing many tasks in highly unstructured environments. For example, the feller-buncher operator is required to locate a tree which is to be harvested, grasp the tree, cut it at the base, and finally move the tree to a specified location. This task is accomplished in an environment which is intrinsically unstructured; the logging location continually changes, the trees are all different, the site is dynamically changing as trees are removed, etc. This type of unstructured environment makes it difficult to program autonomous robots to perform similar tasks.

Robots are currently being widely used in manufacturing applications where the tasks being performed are repetitious and well defined. Two typical examples are welding robots in the automobile manufacturing industry and circuit board assembly robots in the electronics manufacturing industry. In both of these examples the robots have simple internal models of their environment and neither the model nor the environment is expected to change during the manufacturing process. Thus the robots simply follow a preprogrammed set of motions (which usually constitute the robot's internal model of the world) and require relatively few external sensors of their environment in order to accomplish the required tasks. However, if the environment or task change even to a small extent, this approach to programming robots becomes difficult and time consuming to implement.

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 expert's examples, even when the unstructured environments change after learning has occurred. This type of generalization of an unstructured task typically requires some kind of learning system which is capable of appropriately interpolating between learning examples.

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 [12] [11] where he successfully uses a 3 layer perceptron network to control the CMU ALVINN autonomous driving system as it drives along a road. Other related work includes that of Kosuge *et al* [8] 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 humans skills: for the deburring task Shimokura and Liu [14], and for compliance control Asada [2]. Lui and Asada [10] use process dynamics for transferring manipulative skills. Yang *et al* [17] use a multidimensional Hidden Markov Model to represent the transferred skill. Rouse *et al* [13] 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 describe our chosen approach to human-to-robot skill transfer. We then describe a small scale testbed we have implemented for running experiments in teleoperation based human-to-robot skill transfer. The paper ends with a summary of some of the experiments we have done on this testbed. The main goal of these experiments is to demonstrate how effectively a high dimensional, unstructured task can be transfered from human-to-robot given only a few examples of the tasks execution.

## 2 Formulation of the Problem

During teleoperation, the human operator has access to a set of sensors, symbolized by $\mathbf{s}$, and a set of actuators, symbolized by $\mathbf{a}$. Based on a time sampling of sensor inputs, and a time sampling of his or her previous actuator outputs, the human operator generates an actuator response, at time $t$, symbolized by $\mathbf{a}_t$. One can symbolize the mapping which the human operator is using as follows:

$$\mathbf{a}_t = \mathbf{M_H}(\mathbf{s}_t, \mathbf{I_H}(\mathbf{s}_{t-1}, ..., \mathbf{s}_{t-n_H}, \mathbf{a}_{t-1}, ..., \mathbf{a}_{t-m_H})), \tag{1}$$

where,

1. $(\mathbf{s}_t, \mathbf{s}_{t-1}, ..., \mathbf{s}_{t-n_H})$ is a time sampling history of sensor inputs back to time interval $t - n_H$;
2. $(\mathbf{a}_t, \mathbf{a}_{t-1}, ..., \mathbf{a}_{t-m_H})$ is a time sampling history of actuator outputs back to time interval $t - m_H$;
3. $\mathbf{I_H}(\bullet)$ is a function which defines the current state of the task sequence (i.e. what has been done by time interval $t-1$).
4. $\mathbf{M_H}(\bullet)$ is a mapping from the current state vector defined by $\mathbf{I_H}(\bullet)$ and the current sensor input $\mathbf{s}_t$, to the operator's actuator output $\mathbf{a}_t$.

Within this context, the goal of teleoperation based human-to-robot skill transfer is the following:

> Given $q$ time samples, $\mathbf{s}_0, ..., \mathbf{s}_q, \mathbf{a}_0, ..., \mathbf{a}_q$, of sensor input/actuator output pairs produced by a human operator executing a given task, determine time intervals $n_R$ and $m_R$, as well as mappings $\mathbf{I_R}(\bullet)$ and $\mathbf{M_R}(\bullet)$, such that the following control relationship
>
> $$\mathbf{a}_t = \mathbf{M_R}(\mathbf{s}_t, \mathbf{I_R}(\mathbf{s}_{t-1}, ..., \mathbf{s}_{t-n_R}, \mathbf{a}_{t-1}, ..., \mathbf{a}_{t-m_r})), \tag{2}$$
>
> allows the robot to autonomously accomplish the same task, in a similar but not identical environment.

One should note that the most basic assumption made here is that a time sampling of an operator's sensor input/actuator output pairs, contains enough information to generalize the task being performed. In Section 4, we experimentally show that this is a valid assumption for the human-to-robot skill transfer experiments we have conducted.

The goal of human-to-robot skill transfer as stated above not a trivial one to attain. This type of human-to-robot skill transfer scenario has a number of characteristics which must be addressed in order to become feasible. Listed below are 5 characteristics which we feel are significant.

**High dimensional input space**: As stated in the introduction, it is likely that for the human operator to teleoperate a task, he/she will most likely require a large number of sensor inputs. For example, in the task demonstrated in Section 4. the human uses a video image consisting of 1024 pixels, this means that, at a minimum, one will likely need to form an approximation which has 1024 input variables. Such high dimensional approximations are not, in general, easy to construct. Thus a technically feasible approach to reliably building approximations of $\mathbf{I_R}(\bullet)$ and $\mathbf{M_R}(\bullet)$ needs to be formulated.

**Large sample set**: Depending on the complexity of the task being performed, a large number of sensor input/actuator output sample pairs may be required in order to effectively generalize the task. Hence, not only are $\mathbf{I_R}(\bullet)$ and $\mathbf{M_R}(\bullet)$ high dimensional, they may also require large numbers of sample pairs to construct.

**Sensor noise**: Typically, the human teleoperator is observing noisy sensor signals while accomplishing a task. For example, in the task demonstrated in Section 4 the NTSC video image the teleoperator used while accomplishing a task was extremely noisy. Because the robot uses the same same sensor inputs, the mappings $\mathbf{I_R}(\bullet)$ and $\mathbf{M_R}(\bullet)$ must be constructed in such a way as to effectively account for this noise.

**Human inconsistency (non unique solutions)**: Human actions while accomplishing a task are not always consistent. For example, when a teleoperator controls a robot in order to avoid colliding with an obstacle, it may not matter whether the robot goes to the left or right of the obstacle, as long as it is avoided. In fact the operator may maneuver to the left of the vehicle sometimes, and to the right at other times. In order to effectively transfer such an "avoid collision" skill, the mappings $\mathbf{I_R}(\bullet)$ and $\mathbf{M_R}(\bullet)$ must be constructed in such a way as to account for this.

**"Internal" operator state**: Finally, consider the state function $\mathbf{I_R}(\bullet)$. This state function is required when a current action depends on a previous action, the results of which are not directly visible in the sensor inputs $\mathbf{s}_t$. Given simply a finite number of teleoperator input/output pairs, it is not possible to directly observe the "state" of the teleoperator. Therefore, either the teleoperator must somehow indicate current state, or the system must be able to observe it from a the given examples. Note that, in theory, given enough input/output pairs, internal state is observable, and thus $\mathbf{I_R}(\bullet)$ can always be constructed. However, given the high dimensional input space, this may not be easy to implement in practice.

In order to address the above characteristic difficulties of human-to-robot skill transfer, we propose to divide the problem into 2 closely related subproblems. The first we refer to as the *knowledge representation* or function approximation problem, which deals with the problem of how the functions $\mathbf{I_R}(\bullet)$ and $\mathbf{M_R}(\bullet)$ are represented and stored. The second subproblem we refer to as the *knowledge transfer* problem, which defines the interface between the human operator demonstrating the desired task, and the knowledge representation system used to learn the task.

### 2.1 Knowledge Representation

Consider first the *knowledge representation* or function approximation problem. As stated above, the input space of both $\mathbf{I_R}(\bullet)$ and $\mathbf{M_R}(\bullet)$ is typically large, noisy, and large numbers of examples of these mappings may be required to effectively generalize the desired task. We view the problem of knowledge representation as having two aspects. The first is concerned with the representation structure or architecture, and the second is concerned with the associated learning algorithm which generates the structure based on a finite number of examples of the target function. There are many types of multidimensional approximation schemes currently in the literature. A good summary of the more classic approaches can be found in [1]. There are 2 distinct types of approaches to the approximation problem. The first is the parametric approach where the number of parameters in the knowledge representation structure is fixed *a priori*, and the function of the learning algorithm is to set the values of these parameters based on the training examples. Some examples of parametric approximations are polynomials, radial basis functions, spline approximations, and perceptron networks. The second approach is the nonparametric approximation approach where the number of parameters used is not set *a priori*, and the learning system must determine the number of parameters, as well as their values. Nonparametric approximation is usually harder to efficiently implement, and suffers from slow rates of convergence [15]. However, nonparametric approximation is potentially a more powerful approach because it is often not known *a priori* how many parameters are required to appropriately fit the learning data. Examples of nonparametric approximation techniques are GMDH [5], MARS [6], CART [3], and additive regression [16]. Most existing approximation techniques don't have learning algorithms which theoretically guarantee convergence to the target function, while others are not practical to implement when the dimension of the problem is high, or the number of training data examples becomes too large. This makes most existing techniques not suitable for the type of teleoperation based human-to-robot skill transfer proposed here.

In order make human-to-robot skill transfer feasible in practice, we propose that the knowledge representation system used should have the following characteristics:

1. The knowledge representation system should be capable of representing any bounded multidimensional function, to within any desired accuracy. If this condition is not met, then failure to transfer a skill may simply be the result of not being able to represent it within the knowledge representation framework.
2. The complexity of evaluating the mappings $\mathbf{I_R}(\bullet)$ and $\mathbf{M_R}(\bullet)$, once represented within the knowledge representation framework, should not make it impractical to implement in practice when large numbers of sensor inputs are used, or when large numbers of examples are required to generate a sufficient approximation.
3. There should exist a learning algorithm associated with the knowledge representation system, which is able to form uniformly convergent approximations of any bounded continuous multidimensional function. This is an important characteristic because, even if the knowledge representation system is capable or representing the desired function, this is no guarantee that there will in practice exist a learning algorithm which can take examples of the desired function and generate a representation which approximates this function. Furthermore, this learning algorithm should have the following characteristics:
    (a) The algorithm should be relatively insensitive to, in the minimum, Gaussian noise.
    (b) The algorithm should produce nonparametric representations. This is essential because it is often not known how many parameters are necessary to form a sufficient approximation. A trial and error approach to choosing

*a priori* how many parameters to use for a given skill would simply be too time consuming to do.

(c) The rate of convergence of the algorithm, defined as the number of sample points required to achieve a desired level of approximation, should be relatively independent of the number of dimensions (or sensor inputs) used in the knowledge representation. With most nonparametric approximation techniques, the rate of convergence of the algorithm makes it infeasible to implement in practice [15][16], when a large number of input variables are used.

(d) The complexity of forming an approximation should not make the learning algorithm impractical to implement as the dimension of the problem increases. This is another essential characteristic, because often algorithms which are able to effectively form 2 dimensional approximations cannot in practice be made to work on higher dimensional problems [7].

(e) Finally, the complexity of forming an approximation should not make the learning algorithm impractical to implement as the number of learning samples becomes large. This is also an essential characteristic, because often learning algorithms will fail when a large number of learning examples are used [7].

The knowledge representation system we have chosen to use has been specifically designed to meet all of the above goals. A description of the system can be found in [7].

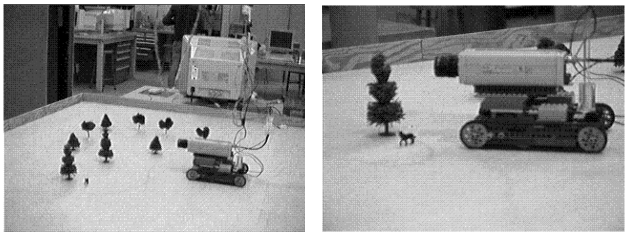

Figure 1: The Human to Robot Skill Transfer Testbed

## 2.2 Knowledge Transfer

Next consider the *knowledge transfer* problem. As described above, a humans actions while teleoperating a task may not always be consistent, the state of the task my not be directly observable, and large numbers of teleoperation examples may be required to effectively transfer a task. As a result, we propose that the knowledge transfer system used by the human operator should have the following characteristics.

**Task subdivision**: Tasks can often be divided up unto a number of steps. For instance, when loading a truck with containers, the operator must first locate an appropriate container, pick it up, locate and move towards the truck, and finally place the container in an appropriate spot on the truck. In our experiments we have found that it is much easier to transfer the skill embodied in any one of these individual steps, than it is to transfer the entire skill at one time. One reason for this is the large number of sample points required to differentiate and convey the entire skill sequence, as opposed to only one part of it. Thus, we believe that it is important that the knowledge transfer system allow the teleoperator to transfer skills in individual stages. The knowledge transfer system should have some automated way of determining what part of the skill has already been acquired so the this part is not re-learned or forgotten when a larger task is learned. The overwriting of a learned task sequence is particularly evident when the human is inconsistent in teleoperating a task [4].

**Automated and operator assisted state identification**: Whenever the current sensor inputs don't allow the teleoperator to directly determine which part of a task has been executed and which hasn't, then identification of the state function, $\mathbf{I_R}(\bullet)$, becomes particularly important. As an example, consider the problem of frying an egg without using a temperature sensor. In order to fry an egg, one must first heat the frying pan to the appropriate temperature. This usually requires one to turn on the burner, and then wait a few minutes for the frying pan to heat up before putting in the egg. In order to transfer this skill to a robot, the teleoperator needs to convey the concept of time either by allowing the robot access to a timer and using some automated procedure for identifying $\mathbf{I_R}(\bullet)$, or, have the operator in some way implicitly convey this state information. We conjecture that an effective knowledge transfer system needs to allow for both types of state identification approaches.

**Incremental learning**: Often, while a skill is being transfered, the robot may partially learn the task sequence. Hence, under certain conditions the robot may not act correctly. When this occurs, it is important that the operator be able to interact with the robot to *incrementally* improve its operation during the times the robot is functioning incorrectly, while not destroying the knowledge that has already been acquired. We believe the ability for the teleoperator to interact with the robot to incrementally improve the robot's ability to execute any one part of the task sequence, is an important part of the knowledge transfer system. This allows the operator to focus examples on parts of the task which the approximator has not generalized well, rather than repeat whole task sequences more times.

It should be noted that to date we have not implemented a knowledge transfer system which has the above characteristics. The characteristics given above are derived from human-to-robot skill transfer experiments we have performed on an experimental testbed we have implemented (described in the next section). However, we are currently working on implementing a knowledge transfer system which has these characteristics.

## 3 The Experimental Testbed

A small scale proof-of-concept testbed for human-to-robot skill transfer has been implemented (see Figure 1), 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 1) to allow it to grasp objects; and 3) an NTSC color camera. The leftmost image in Figure 1 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 demonstrates the desired task a number of times via teleoperation. The human operator observes images received from the camera located on the robot (at a 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.

## 4 Experimental Results

The human-to-robot skill transfer experiments involve the following scenario. An object is picked which a teleoperator can distinguish from the background (the background being the unstructured environment of the robotics lab). The task of the teleoperator is to maneuver the robot in a counterclockwise direction until the chosen object is centered in the camera's field of view, and then approach the object until it is within gripper range. The teleoperator then once more rotates the robot counterclockwise in search of another example of the same object. The human-to-robot skill transfer experiment involves the teleoperator executing this object approach sequence 4 times, for 4 different random starting positions of the object and vehicle. The input/output examples generated by the teleoperator during these 4 approaches are recorded and then an approximation of the mapping $\mathbf{M}_{\mathbf{R}}(\bullet)$ is constructed off-line, based on these examples. Note that no state information is required for this task, and so $\mathbf{I}_{\mathbf{R}}(\bullet)$ is not approximated. Once $\mathbf{M}_{\mathbf{R}}(\bullet)$ has been constructed, the ability of the robot to autonomously execute the same task is evaluated.

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 2 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 approximation; what the human sees is exactly what is passed to the approximator. In order to generalize a 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.

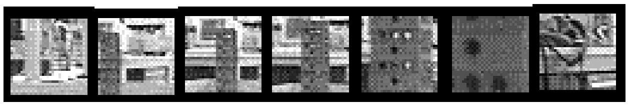

Figure 2: A red and Blue Striped Tower

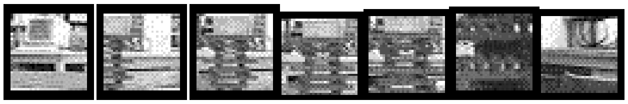

Figure 3: A red and Blue Striped Tower with Holes

Figure 2 and Figure 3 show grey-scale camera frame examples of approach sequences to 2 objects we have experimented with. The object in Figure 2 is a red and blue striped tower, and the object in Figure 3 is also a red and blue striped tower but has a number of holes through it which allow background scenes to be observed. For the object in Figure 2, the teleoperator generated 4632 camera images during 4 approaches. Similarly, for the object in Figure 3, the teleoperator generated 4805 images during 4 example approaches. The variation in camera frame counts is simply a result of the random placement of the vehicle and object. A nonparametric approximation of $\mathbf{M}_{\mathbf{R}}(\bullet)$ was constructed for each object individually as described in [7]. The generation of the 4 approach examples took the teleoperator approximately 5 minutes to generate for each object, while the off-line learning of $\mathbf{M}_{\mathbf{R}}(\bullet)$ required approximately 4 additional minutes on the SPARCstation 20, for each object. The size of the approximation generated for the object in Figure 2 is 426 K-bytes, while the size of the approximation generated for the object in Figure 3 is 516 K-bytes. The time required to evaluate either approximation (once it has been constructed) for an observed image is less than 120 milliseconds for either object (run on a SPARCstation 20).

In order to evaluate the ability of the robot to autonomously execute an approach sequence toward an object, we run 30 experiments for each of the two objects described above. Each experiment involved placing the vehicle and object in random positions, and then seeing whether the robot would autonomously be able to located and approach the object. For the object shown in Figure 2, the robot successfully executed the task 28 out of 30 trials, while for the object shown in Figure 3, the robot was successful in approaching the object 25 out of 30 trials. Note that the initial distance between the vehicle and target object was always less than 1.5 meters, because at distances greater than this, the object is not easily distinguishable, even by a human operator.

Given 1) the high dimensionality of the skill being transfered (1024 sensor inputs), 2) the high level of video noise (see Figure 2 and Figure 3), and, 3) the fact that only 4 example sequences of the task were generated, the more than 80% task success rate of the autonomous robot shows significant promise. We therefore conjecture that teleoperation based human-to-robot skill transfer is a feasible approach to programming robot's that work in unstructured environments.

## 5 Conclusion

We have outlined a set of characteristics which we believe are useful to have in a human-to-robot skill transfer system. In addition, we have demonstrated the efficacy of human-to-robot skill transfer for the task of locating and approaching an object in high dimensional, visually cluttered environments. At present we are in the process of running additional experiments in human-to-robot skill transfer. Specifically, we are interested in studying the efficacy of skill transfer 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 in this paper. We would also like to thank Dan Chan for building the experimental testbed.

## References

[1] P. Alfeld. Scattered data interpolation in three of more variables. In T. Lyche and L. L. Schumaker, editors, *Mathematical Methods in Computer Aided Geometric Design*, pages 1–33. Academic Press, Boston, MA, 1989.

[2] 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.

[3] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and Regression Trees.* Wadsworth International Group, Belmont, California, 1984.

[4] E. Dupius, Sheng Liu, and H. Asada. On the consistancy of complience control laws aquired by humans. In *ASME 1992 Winter Annual Meeting.* ASME, 1992.

[5] Stanley J. Farlow. The gmdh algorithm. In Stanley J. Farlow, editor, *Self-Organizing Methods in Modeling: GMDH Type Algorithms*, pages 1–24. Marcel Dekker, Inc., New York, NY, 1984.

[6] J. H. Friedman. Multivariate adaptive regression splines. *Ann. Stat.*, 19:1–141, 1991.

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

[8] 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.

[9] P. D. Lawrence, F. Sassani, B. Sauder, N. Sepehri, U. Wallersteiner, and J. Wilson. Computer-assisted control of excavator-based machines. In *1993 International Off-Highway and Powerplant Congress and Exposition*, 1993.

[10] 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.

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

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

[13] 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.

[14] 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.

[15] C. J. Stone. Optimal global rates of convergence for nonparametric regression. *Ann. Stat.*, 10:1040–1053, 1982.

[16] C. J. Stone. Additive regression and other nonparametric models. *Ann. Stat.*, 13:689–705, 1985.

[17] 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.