Learning Force Sensory Patterns and Skills from Human Demonstration *

Marjorie Skubic and Richard A. Volz

Department of Computer Science
Texas A&M University
College Station, TX 77843-3112

Abstract

The motivation behind this work is to transfer force-based assembly skills to robots by using human demonstration. For this purpose, we model the skills as a sequence of contact formations (which describe how a workpiece touches its environment) and desired transitions between contact formations. In this paper, we present a method of identifying single-ended contact formations from force sensor patterns. Instead of using geometric models of the workpieces, fuzzy logic is used to learn and model the patterns in the force signals. Membership functions are generated automatically from training data and then used by the fuzzy classifier. This classification scheme is used to learn desired sequences of contact formations which comprise a force-based skill. Experimental results are presented which use the technique to extract skill information from human demonstration data.

1 Introduction

Robot programming by using human demonstration provides a natural approach to human-robot interaction and at the same time, avoids much of the computational complexity found in model-based program generation. Several such programming systems have been proposed and developed. Vision-based systems have been used to observe and extract an assembly plan from human demonstrators (e.g., [10]). While useful for observing high-level task information, vision has so far has not proven adequate for observing intricate assembly operations. Capabilities of such vision systems could be expanded with the use of force-based skills.

While not yet coupled with vision, force signals have been incorporated into programming by demonstration systems, in order to learn compliant task strategies for assembly operations. Early work is presented by Hirzinger and Landzettel in [7]. More recent work includes that of Delson and West [2]. In contrast, we are not attempting to learn an entire task trajectory. Rather we will focus on learning low-level force-based skills. We assume the existence of some other sensory mechanism to position the robot close to the area of interest. This may take the form of a vision system or a teleoperated system with a human operator in the loop. Our goal is to learn how to execute force-based skills once the robot is close to the area of contact.

An approach to assembly operations which we believe can be combined with robot programming by demonstration is to model the assembly task as a sequence of discrete states, representing the contact formations (i.e., contact state) [3] between workpieces. In this way, the learning process can be simplified and the learned program should be more robust to discrepancies in position and orientation. The utility of this approach has been demonstrated by Asada et al in [1]. However, the method used to identify the discrete contact states required a geometric model of the workpieces [6], a burden we would like to avoid in robot programming by demonstration.

McCarragher et al have used the discrete event approach to assembly, incorporating system dynamics and qualitative reasoning [11]. In [8], FFT's are used to identify contact states; however, this approach cannot easily be run in real-time. In [9], Hidden Markov Models are used to model skills, which include the contact states and discretized velocity commands representing the transitions between contact states. This provides a 2-step approach to mapping forces to velocity commands. The first step is to recognize the discrete state (i.e., the contact state), and the second step is to determine the velocity commands which move the robot from one state to the next.

We will borrow McCarragher's approach to skill modeling; the focus of this paper will be on monitoring and identifying the contact states from force signals, without using geometric models. We introduce the concept of a single-ended contact formation (SECF) and present a fast sensor-based algorithm for identifying SECF's. We will utilize previous work [12], which is based on a quasi-static contact condition. We extend the approach to now consider the dynamic system and illustrate how a sequence of SECF's can be extracted.

*This work has been supported by NSF under grants CDA-9115123, CDA-9422123, and EID-9017249 and by the State of Texas under TATP grants 999903-267 and 999903-095.

Similar to the strategy used by Hara and Yokogawa [5], our method uses fuzzy logic to model and recognize the sensory patterns and also resolve ambiguities. In addition to its pattern recognition strengths, fuzzy logic provides a linguistic interface to the human demonstrator, which allows us to assign a meaningful and recognizeable name to each membership class. We believe this will provide a natural method for transferring force-based assembly skills from human experts to robots.

In Section 2, the sensory patterns of single-ended contact formations will be discussed. Section 3 will describe the algorithm for identifying SECF's. Section 4 will present experimental results and discussion. Conclusions will be given in Section 5.

2 Sensory Patterns of Single-Ended Contact Formations

The contact formation [3] provides a qualitative description of how 2 or more objects make contact with each other (e.g., edge 1 of the grasped object touches side B of the environment). As such, the possible contact formations depend on the relative size and shapes of the objects, reflecting the geometric relationship. In contrast, a single-ended contact formation (SECF) provides a one-sided description of how the grasped object touches its environment (e.g., edge 1 of the grasped object touches any side in the environment). It does not attempt to distinguish between similar contact points in the environment.

Our goal is to recognize the contact state from force sensor data only, without using geometric knowledge or position information. It is not possible to determine the complete contact formation from force data alone. However, it is possible to determine a one-sided description of the contact state, using force signals from the robot force sensor. From the human demonstrator's point of view, this is the haptic information needed. A human performing an assembly operation would use vision to position an object close to the desired point of contact and then use haptic senses to complete the task. In this situation, there is no need to determine the complete contact state from haptic senses alone. We wish to provide the same type of feedback to the robot.

To meet our goal of recognizing the SECF, we examine force sensor data, looking for distinguishable features. The set of force and moment vectors that comprise all possible vectors for a given SECF form a cone, which reflects the motion constraints existing between the grasped object and its environment. A theoretical discussion is provided by Hirai and Asada in [6]. Figures 1 and 2 show example force and moment patterns for several measurements on 2 SECF's. This data was collected using the small, plastic block, shown in Figure 3, in contact with a flat surface. Unlike a machined aluminum block, this plastic block does not have sharp edges, and the bottom surface has a texture. Even with these imperfections, the SECF force and moment vectors form distinct patterns.

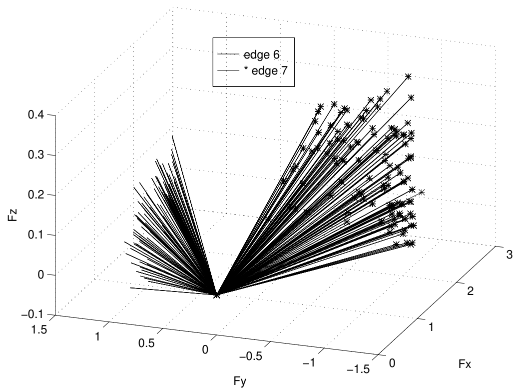

Figure 1: **Force Vector Patterns for Two Edge Contact Formations**

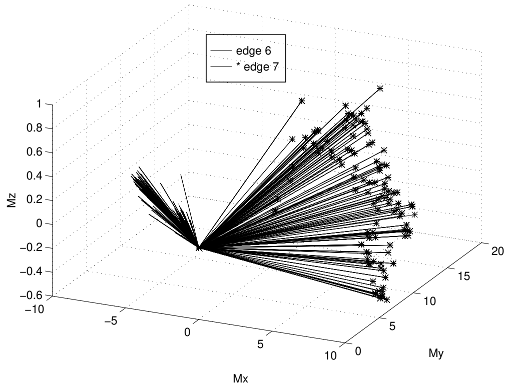

Figure 2: **Moment Vector Patterns for Two Edge Contact Formations**

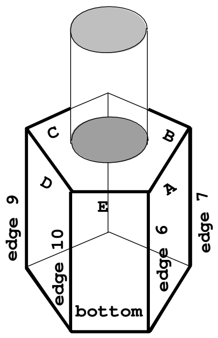

Figure 3: **Pentagon Block with 11 Contact Formations:** 5 sides, 5 edges plus the bottom surface

3 Identification of Single-Ended Contact Formations

Because of the sensory patterns, a classifier can be constructed and used to identify the SECF, based on the force and moment sensory data. Each class represents a unique SECF. We propose, as a sensor-based pattern classifier, the use of fuzzy logic. Empirical training data, instead of geometric models, are used to learn the sensory patterns. Fuzzy membership functions are used to represent the patterns.

3.1 The Fuzzy Classifier

In order to build a classifier, one or more features must be identified which can be used to distinguish between different SECF's. As illustrated in Figures 1 and 2, one significant feature is the direction of the vector. [1] To capture the direction, the data vectors are normalized, yielding a projection onto a unit sphere. Let f_x, f_y, and f_z be the force components of the sensor signal. Then, the normalized force components, f_{xn}, f_{yn}, and f_{zn}, are

$$f_{in} = \frac{f_i}{|f|} \qquad \text{for } i{=}x,y,z$$

$$\text{where} \quad |f| = \sqrt{f_x^2 + f_y^2 + f_z^2}$$

Similarly, let m_x, m_y, and m_z be the moment components of the sensor signal. Then, the normalized moment components, m_{xn}, m_{yn}, and m_{zn}, are

$$m_{in} = \frac{m_i}{|m|} \qquad \text{for } i{=}x,y,z$$

$$\text{where} \quad |m| = \sqrt{m_x^2 + m_y^2 + m_z^2}$$

An additional class, called C_0, is added for the condition where the grapsed object is not contacting the environment. The force magnitude, $|f|$, will be used to identify this class.

Our classifier is structured as one logical rule per class, which describes the mapping between sensory data and class. The force magnitude and the 6 normalized sensor signals are used as input features for the classifier. The set of possible signals for each feature is represented as a fuzzy membership function. The membership function for class C_0 is represented by *Zero* (force magnitude). The membership functions for each contact class C_i are represented by Fx_i, Fy_i, Fz_i, Mx_i, My_i, and Mz_i. Fuzzy conditions about sensor signals are combined using a fuzzy conjunction operator. The rule for class C_0 is shown below:

$$\begin{array}{lll} \text{if} & |f| & \text{is } Zero \\ \text{then} & class & \text{is } C_0 \end{array}$$

All of the remaining rules have the following form:

$$\begin{array}{llll} \text{if} & |f| & \text{is } NotZero & \text{and} \\ & f_{xn} & \text{is } Fx_i & \text{and} \\ & f_{yn} & \text{is } Fy_i & \text{and} \\ & f_{zn} & \text{is } Fz_i & \text{and} \\ & m_{xn} & \text{is } Mx_i & \text{and} \\ & m_{yn} & \text{is } My_i & \text{and} \\ & m_{zn} & \text{is } Mz_i & \\ \text{then} & class & \text{is } C_i & \end{array}$$

The membership functions are generated automatically using supervised learning. Training data is acquired by demonstrating each SECF class and collecting the force and moment signals generated. For each class, the mean, μ, and standard deviation, σ, are calculated. These are used as parameters to generate the membership functions, $\pi(r)$, as shown below:

$$\pi(r) = 1 - e^{-(3\sigma/|\mu - r|)^3}$$

Similarly, the membership functions for force magnitude are as follows:

$$Zero(|f|) = 1 - e^{-(3\sigma/|f|)^3}$$

$$NotZero(|f|) = 1 - Zero(|f|)$$

A fuzzy conjunction operator is used to combine the 6 membership functions. In this case, the Hamacher product is used in its reduced form [12]. The combination of the 7 feature membership functions is interpreted as the possibility (i.e., a fuzzy number in [0,1]) of this sensor data set being in the specified class. In the classification process, possibility values are calculated for each class; the highest possibility number represents the identified CF class. This has the effect of resolving classification ambiguities by selecting the class with the nearest prototype vector, which would be at the center of the combined fuzzy membership functions. In practice, a threshold is also used to determine whether any classification can be made. In the current tests, the threshold is empirically set to 0.5. If no possibility value is above the threshold, then the data vector is said to be unclassified.

3.2 Classifying Time Series Data

The next step is to identify the SECF when the robot is moving and the grasped object makes contact with its environment. We would like to use an efficient method that can be used both to extract a sequence of SECF's off-line from demonstration data and also to monitor the contact state in real-time during the execution of the skill.

[1] Imagine holding the edge of a block against a table. Pressing the block harder against the table will yield a higher magnitude in force and/or moment, but the directions of the vectors will still be within the cones for that SECF.

For efficiency, the classifier should be run only when the contact condition changes. This is analogous to the force event defined by McCarragher [11], which is a gain or loss of contact resulting in a new contact formation. The force event is observed as a significant change in force. By monitoring the change in force (performed as a threshold check), we can decide when the classifier should be run to detect a change in SECF.

When the change in force is used to identify the contact formation (as McCarragher has done), the setting of the threshold value is important. If the threshold is set too high, it is possible to miss an event. If the threshold is set too low, it is possible to incorrectly interpret dynamic effects as a change in contact. With our approach, the value of the threshold is not so critical. It is still possible to set the threshold too high as to miss significant force events. However, setting the threshold too low will merely result in a loss of efficiency but will not affect the correctness of classification.

With this in mind, we use the concept of a force event to classify time series data. A force event is defined as a significant change in any of the 3 force components. That is, a threshold check is performed for each of the 3 individual force components. If a force event is detected, then the force magnitude is calculated and the force and moment vectors are normalized. The fuzzy classifier is then executed using these seven features as described in the previous section. This provides a fast, efficient means of monitoring SECF's. (See Section 4.2 for timing information.)

4 Experimental Results

To test the identification strategy, experiments were conducted in the Texas A&M robotics lab. We were particularly interested in observing the effects of an imperfect environment, including sensor noise.

4.1 The Demonstration Environment

The demonstration equipment included an American Robot Merlin robot arm controlled by a PC, running lynxOS and RCCL. A JR3 force sensor was mounted on the robot arm. For these tests, the robot grasped a plastic pentagon-shaped object, as shown in figure 3, in a consistent orientation, although not necessarily aligned with the robot axes. The pentagon-shaped block was chosen over a square block because the geometric features of the block cannot all align with the axes of the force sensor. This provides a more rigorous test of the classification method. To accomodate the robot's 3-pronged gripper, the workpiece was fitted with a cylindrical handle. The environment was an open wooden, rectangular box.

For the human-controlled demonstration, the PHANToM haptic interface was used as a hand controller to generate 6 degree of freedom position commands to the robot. The PHANToM provides haptic feedback to the operator in the form of 3-dimensional force signals from the JR3 sensor. That is, haptic feedback occurs as a point source of force with no torque feedback.

The PHANToM interface to the robot was designed to be intuitive to the operator, so that movements of the PHANToM tip corresponded directly to movements of the robot's end effector. To control the robot, the operator held a pen-like stylus located at the end of the PHANToM arm, in a similar configuration as the robot holding the handle of the object. The servo loop providing haptic feedback to the operator ran at a frequency of 1200 Hz. Position and force information were logged at a frequency of 120 Hz.

4.2 Training the Classifier

Training data was collected using the PHANToM interface to drive the Merlin robot. With the robot grasping the workpiece, the operator moved the robot so that the workpiece made contact with the environment along a single SECF. Holding the same SECF, the operator rotated the robot in an attempt to cycle through the range of possible positions and forces. This was repeated for each of the 11 SECF classes.

The fuzzy classifier was trained using half of the training data and tested using the other half. Of the 1845 data points in the test set, 1822 were correctly classified, 15 were incorrectly classified, and 8 could not be classified, yielding a success rate of 98.8%. Running on a 133 MHz pentium, the test program took 0.539 sec to classify all 1845 data points, which is an average of 0.29 msec per point. Thus, the algorithm is more than fast enough for real-time force-based systems.

4.3 Learning a Sequence of Contact Formations

After the classifier was trained, test runs were made using sequences of SECF's. The first sequence consisted of 7 steps, as follows: no contact, edge 9, no contact, side C, no contact, edge 8, and no contact. Figures 4 and 5 show the force and moment profiles for this sequence.

Using the algorithm discussed, the SECF's were classified. Figure 6 shows the unfiltered classification results. For purposes of graphing the results, an index is assigned to each SECF. Cases of no contact are shown with an index of 0. Cases where no classification could be made are shown with an index of -1.

A filtering algorithm is applied to the classification results, yielding new results as shown in Figure 7. The filtering algoithm eliminates the cases of unclassification and also filters out transitory classes. This has the effect of eliminating undesired human-generated actions. Human factors studies have shown that the bandwidth for human limb motion is at most 5 Hz for internally generated trajectories and around 10 Hz for reflex actions [4]. In these experiments, we have filtered out transitory states occurring faster than 15 Hz.

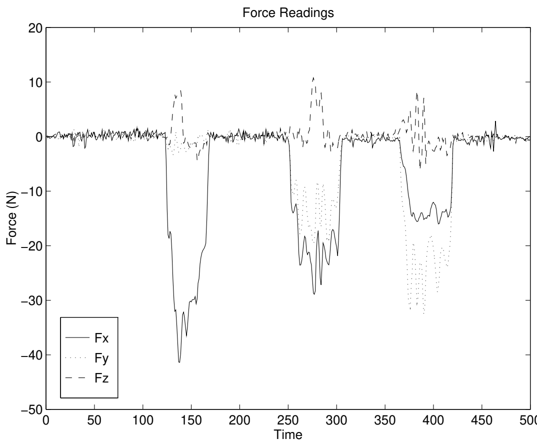

Figure 4: **Force Profile for Sequence 1**

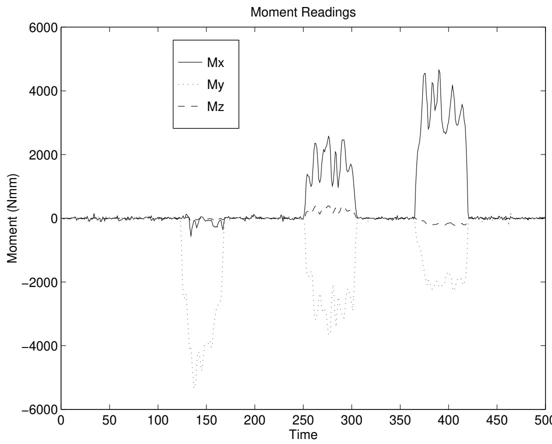

Figure 5: **Moment Profile for Sequence 1**

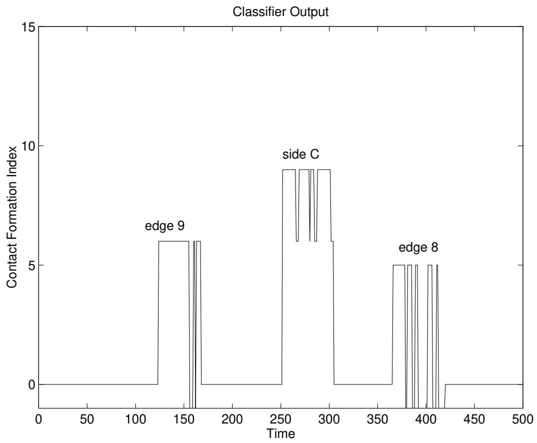

Figure 6: **Unfiltered Classification Results for Sequence 1**

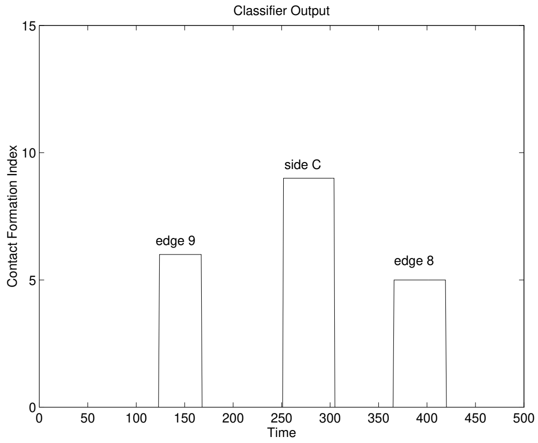

Figure 7: **Filtered Classification Results for Sequence 1**

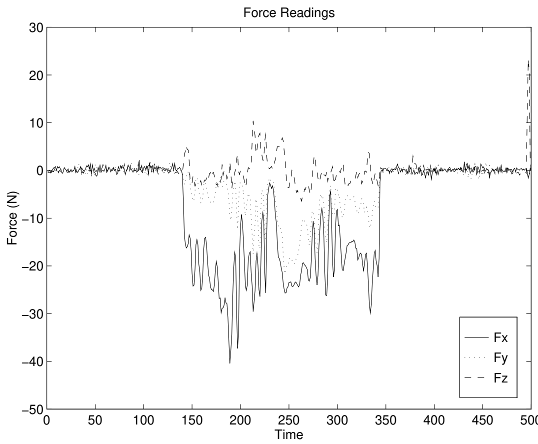

Figure 8: **Force Profile for Sequence 2**

The second sequence consisted of 6 steps as follows: no contact, edge 9, side C, edge 9, no contact, and bottom. Force and moment profiles are shown in Figures 8 and 9. The unfiltered and filtered classification results are shown in Figures 10 and 11.

4.4 Discussion

Although not used in the identification of SECF's, we have included sample position profiles for the purpose of discussion. Figures 12 and 13 show the translation and rotation profiles for the first sequence. One noticeable feature of the human-generated data is the inconsistency. Inconsistency exists in the position data, and this results in inconsistency in the force data. By matching corresponding time steps in Figures 6 and 13, we can see that the transitory edge 9 classifications occurring between time steps 250 and 300 have probably been correctly classified. For the same time period, the rotation profile shows movement about the z axis. This movement would correctly correspond to an unintended shift from the side

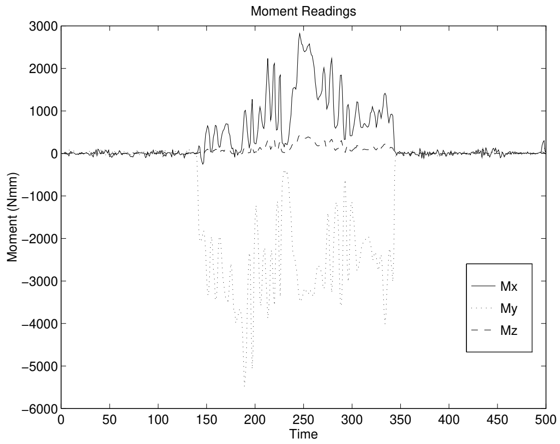

Figure 9: **Moment Profile for Sequence 2**

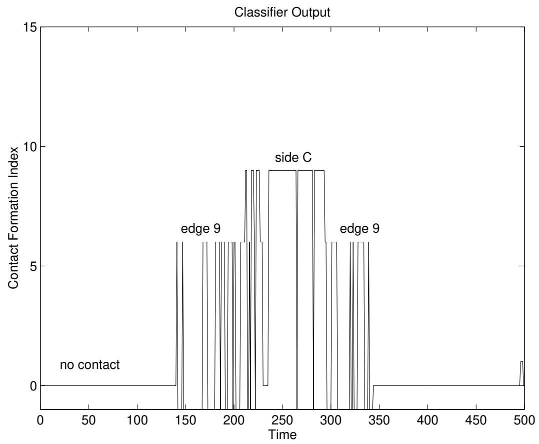

Figure 10: **Unfiltered Classification Results for Sequence 2**

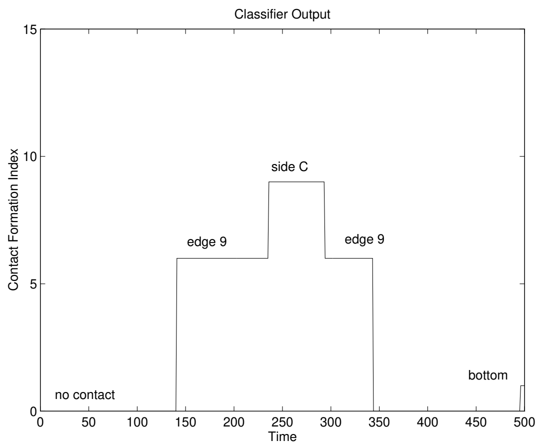

Figure 11: **Filtered Classification Results for Sequence 2**

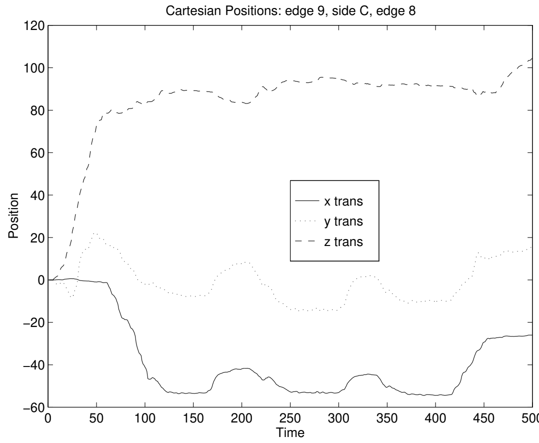

Figure 12: **Translation Profile for Sequence 1**

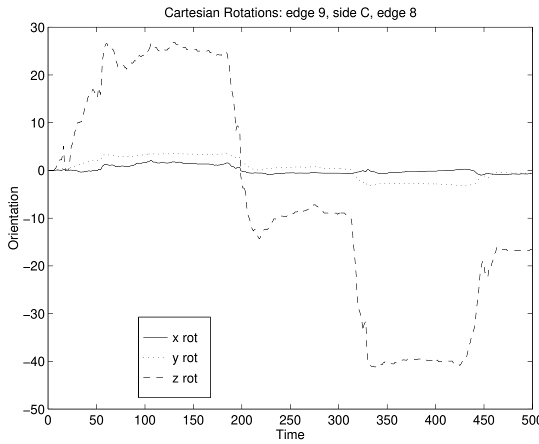

Figure 13: **Rotation Profile for Sequence 1**

C state to the edge 9 state. This provides additional justification for filtering out transitory states from the initial classification results. In the quest to transfer skills from human experts to robots, we would like to avoid including unintended and undesired actions.

Simple test cases are shown here to illustrate the technique. Work is continuing with more complex cases; preliminary results look promising. Future research directions include work on an interactive user interface which is intended to make the demonstration process easier for the human teacher.

The identification of SECF's is also being used to extract more complete skill information from human demonstrations. The SECF's are used to segment the skill into subtasks. Transitions between SECF's are extracted from position data, in the form of reference velocity commands to the robot controller.

5 Conclusion

In this paper, we discussed the acquisition of force-based assembly skills from human experts. We assumed the existence of some sensory mechanism to position the robot close to the area of contact and concentrated on how the skill could be executed from that point. In this context, the concept of a single-ended contact formation (SECF) was presented and discussed.

To transfer and execute skills using SECF information, we needed a fast, efficient method of identifying them from force sensor data. For this purpose, we presented an algorithm which uses fuzzy logic to model and classify force sensory patterns, without using geometric information. Experimental results were presented which illustrated the use of our algorithm for identifying sequences of SECF's which were extracted from force profiles of human-demonstrated skills. It was clear that, not only is our classification highly reliable, but the efficiency of the approach is more than adequate for inclusion in real-time dynamic systems. It thus becomes the basis upon which we can build a system for transferring human force-based skills to robots.

References

[1] H. Asada and H. Izumi. Automatic program generation from teaching data for the hybrid control of robots. *IEEE Transactions on Robotics and Automation*, 5(2):166–173, April 1989.

[2] N. Delson and H. West. Robot programming by human demonstration: Adaptation and inconsistency in constrained motion. In *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, volume 1, pages 30–36, Minneapolis, MN, April 1996.

[3] R.S. Desai and R.A. Volz. Identification and verification of termination conditions in fine motion in presence of sensor errors and geometric uncertainties. In *Proceedings of the 1989 IEEE International Conference on Robotics and Automation*, pages 800–807, Scottsdale, AZ, May 1989.

[4] N.I. Durlach and A.S. Mavor, editors. *Virtual Reality: Scientific and Technological Challenges.* National Academy Press, 1995.

[5] K. Hara and R. Yokogawa. Recognition of state in peg-in-hole by fuzzy schema. *Journal of Advanced Automation Technology*, 4(3):134–139, 1992.

[6] S. Hirai and H. Asada. A model-based approach to the recognition of assembly process states using the theory of polyhedral convex cones. In *Proceedings of the 1990 Japan USA Symposium on Flexible Automation*, pages 809–816, Kyoto, Japan, 1990.

[7] G. Hirzinger and K. Landzettel. Sensory feedback structures for robots with supervised learning. In *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, pages 627–635, St. Louis, MO, March 1985.

[8] G.E. Hovland and B.J. McCarragher. Frequency-domain force measurements for discrete event contact recognition. In *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, volume 2, pages 1166–1171, Minneapolis, MN, April 1996.

[9] G.E. Hovland, P. Sikka, and B.J. McCarragher. Skill acquisition from human demonstration using a hidden markov model. In *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, volume 3, pages 2706–2711, Minneapolis, MN, April 1996.

[10] S.B. Kang and K. Ikeuchi. Toward automatic robot instruction from perception– temporal segmentation of tasks from human hand motion. *IEEE Transactions on Robotics and Automation*, 11(5):670–681, October 1995.

[11] B. J. McCarragher. Force sensing from human demonstration using a hybrid dynamical model and qualitative reasoning. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, volume 1, pages 557–563, San Diego, CA, May 1994.

[12] M. Skubic and R.A. Volz. Identifying contact formations from sensory patterns and its applicability to robot programming by demonstration. In *Proceedings of the 1996 IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 458–464, Osaka, Japan, November 1996.