Data-Driven Haptic Perception for Robot-Assisted Dressing

Ariel Kapusta*[1], Wenhao Yu[2], Tapomayukh Bhattacharjee[1], C. Karen Liu[2], Greg Turk[2] and Charles C. Kemp[1]

Abstract—Dressing is an important activity of daily living (ADL) with which many people require assistance due to impairments. Robots have the potential to provide dressing assistance, but physical interactions between clothing and the human body can be complex and difficult to visually observe. We provide evidence that data-driven haptic perception can be used to infer relationships between clothing and the human body during robot-assisted dressing. We conducted a carefully controlled experiment with 12 human participants during which a robot pulled a hospital gown along the length of each person's forearm 30 times. This representative task resulted in one of the following three outcomes: the hand missed the opening to the sleeve; the hand or forearm became caught on the sleeve; or the full forearm successfully entered the sleeve. We found that hidden Markov models (HMMs) using only forces measured at the robot's end effector classified these outcomes with high accuracy. The HMMs' performance generalized well to participants (98.61% accuracy) and velocities (98.61% accuracy) outside of the training data. They also performed well when we limited the force applied by the robot (95.8% accuracy with a 2N threshold), and could predict the outcome early in the process. Despite the lightweight hospital gown, HMMs that used forces in the direction of gravity substantially outperformed those that did not. The best performing HMMs used forces in the direction of motion and the direction of gravity.

I. INTRODUCTION

Robotic assistance with activities of daily living (ADLs) [1] could potentially enable people to be more independent. This may improve quality of life and help address societal challenges, such as aging populations, high healthcare costs, and shortages of healthcare workers found in the United States and other countries [2], [3].

In this paper we focus on robot-assisted dressing. Studies indicate that more older adults receive human assistance with dressing than other common ADLs, except for bathing/showering, and that over 80% of people in skilled nursing facilities require assistance with dressing [4]. Due to the complexity of dressing, healthcare professionals have developed specialized techniques for dressing assistance that can vary based on a person's disabilities. Likewise, there exist a number of specially designed assistive devices to help people maintain their independence. However, current assistive devices for dressing, such as reachers, dressing sticks, long-handled shoehorns, and sock aids, provide limited support and rely on the user having substantial cognitive, perceptual, and motor capabilities [5], [6]. Robots cold
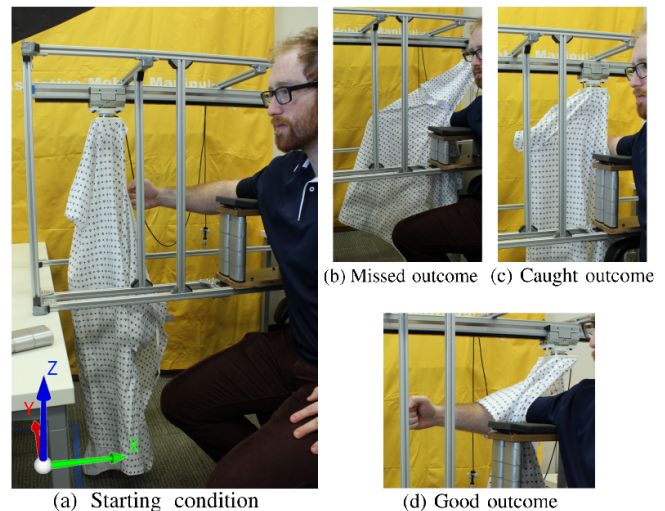

Fig. 1: An experimenter demonstrating: (left) the starting configuration of our robot-assisted dressing subtask with labeled axes; (right) the three outcomes of the task.

potentially serve as more versatile assistive devices for dressing and thereby grant people greater independence and a higher quality of life.

During complex manipulation tasks like dressing, line-of-sight sensors often cannot directly view task-relevant aspects of the environment due to occlusion from objects and the robot's own body. Similarly, critical information, such as applied forces, may not be visually evident. Providing assistance with dressing accentuates these challenges, since the task requires that the robot: physically interacts with complex, deformable, and articulated objects; avoids human discomfort and injury; and works with clothing that visually occludes the human body.

Haptic perception via the robot's end effector may be well-matched to robot-assisted dressing, providing a useful channel of information that complements other sensory modalities. Haptic perception does not rely on line of sight, and measured forces should be directly related to task-relevant phenomena, such as the forces applied by the clothing to the human body. When considered as a time series, forces might also have useful information about the state of the task, such as the current relationship between clothing and the human body. However, there has been a lack of controlled studies focused on the role of haptic perception during robot-assisted dressing.

In order to assess the potential usefulness of haptic perception for robot-assisted dressing, we conducted a controlled study with 12 human participants (see Figure 1) and used

[1]A. Kapusta, T. Bhattacharjee, and C. C. Kemp are with the Healthcare Robotics Lab, Georgia Institute of Technology, Atlanta, GA.

[2]W. Yu, K. Liu, and G. Turk are with the School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA.

*A. Kapusta is the corresponding author (akapusta@gatech.edu)

well-established algorithms for pattern recognition (HMMs). We trained the HMMs exclusively on time series of measured forces. Overall, our results suggest that haptic perception can play a useful role in robot-assisted dressing. In particular, forces measured at the robot's end effector were highly informative about the relationship between clothing and the human body. Classifiers performed well when evaluated with people and end-effector velocities outside of the training data. Classifiers also performed well when we limited the maximum force that the robot could apply. They accurately classified the relationship between the clothing and the human body early in the dressing task. Notably, our results suggest that forces in the direction of gravity can be especially informative.

II. APPROACH

In order to assess the potential usefulness of haptic perception for robot-assisted dressing, we conducted a carefully controlled study with 12 human participants (see Figure 1) and used well-established algorithms for pattern recognition (HMMs). In order to design an informative and replicable experiment, we deliberately focused on a representative subtask of dressing with a commonly used article of clothing. In our experiment, a robot pulled a hospital gown along the length of each person's forearm. This representative task resulted in one of the following three outcomes: the hand missed the opening to the sleeve; the hand or forearm became caught on the sleeve; or the full forearm successfully entered the sleeve (see Figure 1).

A. The Dressing Subtask

We selected a hospital gown because it is a common article of clothing with which health professionals provide assistance. The dressing subtask of pulling a sleeve onto a person's forearm is representative of many dressing tasks that involve pulling a fabric tube over a part of the body, including pulling a sleeve over an upper arm, a pant leg over the lower leg, a sock over a foot, or a stocking cap over a head. In more general terms, the three outcomes for our task relate to initial insertion of a body part into a fabric tube and sliding a fabric tube along a body part until reaching a joint. The body part can miss the opening to the fabric tube. The body part can become caught on the fabric tube prior to subtask completion. Or, the full body part can successfully enter the fabric tube.

For a robot, these three outcomes have distinct implications. If the body part misses the opening, the robot would likely try again. If the body part becomes caught, the robot should avoid causing discomfort to the human. If the full body part successfully enters, then the robot would likely prepare for the next subtask, or finish the dressing task.

B. Data-driven Haptic Perception

Our goal was to classify the outcome of the task based on the time series of forces. The HMMs used the raw forces measured at the robot's end-effector in the direction of movement and upward direction (X and Z directions, respectively). We investigated the performance of the classifiers using univariate and bivariate models using the force in the X direction, the

force in the Z direction, or both, and present these results in Section V-A.4.

III. RELATED WORK

This work lies at the intersection of robot-assisted dressing, anomaly detection, and data-driven haptic perception for robotic assistance.

A. Robot-Assisted Dressing

Researchers who have previously investigated robot-assisted dressing have focused on using kinematics and vision. Gao et al. presented work on user modeling during robot-assisted dressing. Their focus was on visually identifying the pose of the user, modeling the movement space of the upper body joints, and selecting where to place the openings of a sleeveless jacket [7]. Tamei et al. proposed a method using reinforcement learning for a robot to learn trajectories to dress a mannequin in a t-shirt, focusing on topological relationships between the mannequin and shirt. Their experiments were initiated with the mannequin's arms in the shirt sleeves [8]. Klee et al. introduced an approach for a robot manipulator to perform robot-assisted dressing and demonstrated the approach, having a Baxter robot place a hat on two human participants. They represented dressing tasks as a sequence of goal poses with respect to the user. The robot requested that the user reposition themselves if the robot decided that the goal was infeasible. The robot also modeled the user's constraints to determine where to reposition the user [9].

Other researchers have focused more on the problem of visual perception of clothing for dressing tasks. Work by Koganti et al. used a depth sensor to estimate the mannequin-cloth topology when pulling a t-shirt onto a mannequin using reinforcement learning [10]. More recent work by Koganti et al. 2015 further investigated the problem of visual perception of cloth, particularly in the robot-assisted dressing task of pulling a t-shirt onto a mannequin. They used a dynamic model of the cloth to aid in determining the human-shirt topology [11]. Yamazaki et al. presented a method with which a humanoid robot assisted a single participant in putting on pants. Their method primarily used optical flow to detect failure cases. They also used force sensing to detect task failure, but did not describe that part of their method in detail in the paper [12].

B. Anomaly Detection

Our work has similarities with anomaly detection from data such as trajectories and forces [13], [14], [15], [16], [17]. Pastor et al. presented an approach for a robot to acquire new motor skills from demonstration using reinforcement learning. Within their framework they included a method for predicting failure by monitoring for statistical outliers [18]. In contrast, we are interested in classification between a few outcomes in the specific task of robot-assisted dressing, not the detection of anomalies in general.

Fig. 2: The robot end effector holding the hospital gown with rounded rectangular plastic pieces sandwiching the fabric. (right) Annotated version of the image on the left.

C. Data-driven Haptic Perception for Robotic Assistance

Redmond et al. collected haptic data from various activities of daily living and found statistical differences between similar tasks (e.g. writing with a pen vs pencil) [19].

The data-driven algorithms we use are similar to the algorithms we used in prior work to categorize object properties based on time series of tactile sensing data [20]. We have also demonstrated in prior research the value of using data-driven models of forces in other assistive tasks, such as shaving a person's face [21] and opening doors [22]. Unlike these tasks, our robot-assisted dressing task involves the robot indirectly interacting with the human body via a complex deformable object (cloth).

IV. METHODS

In this section, we describe the robot we designed for this study, the experiment we conducted with human participants, and the classifiers we implemented using HMMs.

A. The Robot

We designed and built a specialized robot to collect data (see Figure 1). The robot consists of a 1.2 meter long Festo linear actuator (part number DGE-25-1200-SP-KF-GV) driven by an Animatics SmartMotor (part number SM2315DT) within an aluminum frame. An ATI Mini45-ERA force-torque sensor sits between the linear actuator's slider and a special 3D printed end effector that holds the hospital gown. The hospital gown is a Medline Patient Gown (part number MDTPG3RABPAS) 1.32 meters long and weighing ~250 grams. People often use two hands to hold a sleeve open when assisting another person with dressing. We designed the robot's single end effector to be wide enough (12 cm) to hold the sleeve open. The end effector holds the top of the gown at the shoulder above the right sleeve (see Figure 2). This allows the robot to pull the right sleeve onto the participant's right arm.

We oriented the X axis of the force-torque sensor to be in the direction of movement as the robot pulls the hospital gown along the length of the participant's forearm. The force-torque sensor's Z axis is parallel to gravity and points upward. We recorded force measurements from the force-torque sensor at 100Hz. Figure 1(a) shows the robot with a participant in the starting condition before the robot pulls the hospital gown along the length of the participant's forearm.

B. The Experiment

We conducted a study with 12 people, since we expected anatomical variations and other differences between people to affect the forces measured by the robot. We conducted this research with approval from the Georgia Institute of Technology Institutional Review Board (IRB), and obtained informed consent from all participants.

1) Recruitment: We recruited 14 able-bodied participants via word of mouth and email, but only included data from 12 participants (N = 12) due to one participant's large arm motions during trials, and another participant's limited arm mobility, which resulted in the participant not attaining the required arm posture. We obtained written informed consent from all participants according to our experimental protocol. We required participants to meet the following inclusion/exclusion criteria: ≥ 18 years of age; have not been diagnosed with ALS or other forms of motor impairments; and fluent in written and spoken English. Of the 12 participants, 5 were female and 7 were male. Their ages ranged from 22 to 30 years, and the lengths from their inner elbows to the ends of their fists ranged from 0.288 m to 0.35 m.

2) Experiment Design: The participant held his/her arm in a specified posture at four heights. The specified posture was: upper arm and forearm parallel with the floor; central axes of the upper arm and forearm at 90 degrees with respect to one another; forearm and hand aligned with the linear actuator's axis of movement; fingers curled into a loose fist; and palm parallel with X-Z plane. In addition to verbal instructions, the experimenters placed a piece of metal with a right-angled cross-section in a specific location within the robot's frame, set using a spacing bar, and had the participant position the right angle snugly within the bend of his/her elbow. Once the arm was positioned, the piece of metal was removed. We uniformly set arm heights using an adjustable armrest. The participant adjusted his/her height by varying the height of the stool on which he/she was sitting in order to have the upper arm parallel to the ground when resting on the armrest.

At each height, the robot pulled the gown along the length of the forearm 10 times in total at two speeds: 5 times at 0.1 m/s, and 5 times at 0.15 m/s. The initial state for each trial had the robot stationary with the gown only touching the robot end effector. The robot stopped moving if the forces in the X or Z direction reached 10 N or if it traveled 85 cm.

We positioned the participant's arm at several heights during the trials in order to have examples of each outcome of our robot-assisted dressing task. At the lowest height the arm missed the opening to the sleeve, passing below it, in every trial. At the middle height, either the arm would miss

Fig. 3: The mean (solid line) and one standard deviation (shaded area) for the three outcomes. (top graph) Forces in the direction of movement. (bottom graph) Forces in the upward direction.

the opening of the sleeve, the fist would become caught on the opening, or the the opening of the sleeve would become caught on the participant's forearm. At the highest height, the arm entered the sleeve successfully and was caught at the elbow in every trial. We also conducted trials with a height between the middle and highest heights. However, we could not confidently label the outcomes for all of these trials, and thus excluded this height from our study.

3) Safety: Safety for the participant was an important consideration during the experiments. The robot uses a ball-and-screw linear actuator to move its end effector with high precision. The robot can apply force to the person's arm via the gown. The forces can increase quickly when the gown becomes caught.

We designed the experiment to mitigate risks associated with interacting with the robot. We positioned the right arm of the participant without obstructions behind it, so in the case of high forces, he/she could move his/her arm to reduce the forces. In informal tests with lab members, we found that there is variation in how people hold their hands while donning a sleeve. To reduce the risk of jammed or caught fingers, we asked the participant to hold his/her right hand with fingers curled into a loose fist.

We programmed the robot's controller to stop movement if the change in force from the start of the trial exceeded 10 newtons in any single direction. Throughout the experiments, an experimenter held an emergency stop switch that would kill power to the robot and immediately stop its movement.

4) Labeling: We assigned one of the following three labels to each trial: *missed*, *good*, or *caught*. We assigned the label *missed* when the participant's arm did not enter the opening to the sleeve and the gown continued to move along the participant's arm unimpeded, which happened in all 120 trials at the lowest height and 29 trials at the middle height. We assigned the label *good* when the sleeve successfully covered the participant's forearm and stopped in contact with the person's upper arm, which happened in all trials at the highest height (120 trials). We assigned the label *caught* when the sleeve became caught on the participant's fist or forearm (91 trials), which only happened at the middle height. In 76 of the *caught* trials, the opening to the sleeve became caught on the participant's fist, but sometimes (15 trials) the sleeve became caught on the forearm.

5) Preprocessing Data: For all measured forces, we subtracted the force reported by the force-torque sensor at the start of the trial, and thus our analysis and system only used the change in force from the start of the trial. We expect this change in force to primarily be associated with physical interactions between the garment and the participant's arm, although forces due the inertia of the cloth and other phenomena would also influence the signals.

We measured the length of each participant's arm from the inner side of the elbow to the end of the fist. We used this distance to make the time series for each trial start when the person's fist was below the front edge of the robot's end effector, which is prior to any contact occurring

between the garment and the person's body. We did not use any data recorded prior to this configuration. In terms of experimental methodology, we did not control for variability due to oscillations of the clothing. This preprocessing leaves natural variation in the state of the pendulum-like oscillations of the garment, since the location of the front of the fist varies across participants, but the dynamics of the cloth should remain relatively consistent with respect to distance along the robot's linear actuator.

We also truncated data after the 10 newton threshold was reached or after the robot's movement had stopped after moving 85 cm in the *missed* outcome, because participants frequently relaxed their arms after the robot stopped.

C. The Classifiers

We used hidden Markov models (HMMs) [23] to model the time series force data associated with the dressing task, building on our prior work on tactile sensing [24] and the work of others [25].

We used the General HMM library (http://ghmm.org/) to train and run left-right HMMs. We trained one HMM for each of the three task outcomes, with the training set labeled as *missed*, *good*, or *caught* as described in Section IV-B.4. Once trained, the three HMMs served as a classifier. Given a new time series of forces, the classifier would find the likelihood of each of the three HMMs generating the times series of forces. The classifier then assigned the category label associated with the HMM with the highest likelihood.

1) Selecting the number of states: We created univariate and bivariate HMMs using the force in the X direction, the force in the Z direction, or both. We evaluated their accuracy using various numbers of hidden states (3, 5, 10, 15, 20, 30, and 40) using the cross-validation method described in Section V-A.1. We found that 10 states for the bivariate HMMs and 3 states for the univariate HMMs gave the highest accuracy. For this paper, we only report results with 10 state bivariate HMMs and 3 state univariate HMMs.

V. RESULTS

In this section, we present classification performance using data-driven haptic perception. Unless explicitly stated, all results are from classifiers using bivariate HMMs. In our results, we frequently refer to the accuracy of the classifiers, which we use to mean the percentage of trials that were correctly labeled, where each trial can only be given one label. Figure 3 shows the force data in the X and Z directions from all 360 trials (12 participants x 3 heights x 2 velocities x 5 repetitions).

A. Cross-Validation Performance of HMMs

1) Generalization to new participants: We recorded trials with 12 participants. We used leave-one-subject-out cross validation to evaluate how well classifiers could generalize to new participants. This evaluation trained a total of 12 classifiers. We trained each classifier on 330 trials from 11 participants and then tested it on the remaining 30 trials from the one participant who was left out. For training, we used

Fig. 4: Confusion matrix showing the bivariate HMMs' performance with leave-one-subject-out cross validation.

Fig. 5: Confusion matrix showing the bivariate HMMs' performance with leave-one-velocity-out cross validation.

the entire time series of force data for both speeds with time-warping to sync the two speeds. We interpolated points for the slower speed to obtain synced data. For testing, we used the entire time series of force data for both speeds without time warping.

Since each of the 360 trials was classified exactly once, we report the overall accuracy and provide a confusion matrix in a manner analogous to the performance of a single classifier applied to all 360 trials. Our HMM classifiers had an overall accuracy of 98.61%. Note for comparison that we would expect a random classifier to achieve 33% accuracy and a majority classifier to achieve 40% accuracy. Figure 4 shows the resulting confusion matrix. These results show that data-driven haptic perception can accurately categorize outcomes of our robot-assisted dressing subtask, and that it can generalize to new people.

2) Generalization to new velocities: We recorded trials with 2 velocities, 0.1 m/s and 0.15 m/s. We used leave-one-velocity-out cross validation to evaluate how well classifiers could generalize to new velocities. This evaluation trained a total of 2 classifiers. We trained each classifier on 180 trials

TABLE I: The accuracy of the bivariate HMMs as the test data is truncated using various force threshold. These values are included on the graph shown in Figure 6.

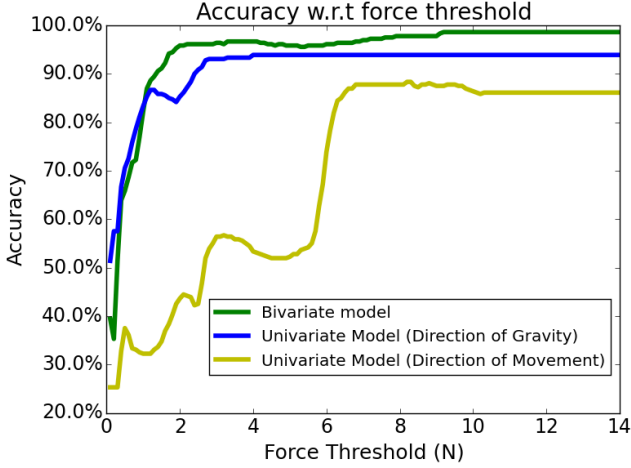| Force Threshold	Accuracy
1.0	82.2%
2.0	95.8%
3.0	96.4%
5.0	95.8%
10.0	98.6%

Fig. 6: The accuracy of the bivariate and univariate HMMs as the test data is truncated using various force thresholds.

from 1 velocity and tested it on the remaining 180 trials from the velocity that was left out. For both training and testing, we used each trial's entire time series of force data for both training and testing without time warping.

Since each of the 360 trials was classified exactly once, we report the overall accuracy and provide a confusion matrix in a manner analogous to the performance of a single classifier applied to all 360 trials. Our HMM classifiers had an overall accuracy of 98.61%. Figure 5 shows the resulting confusion matrix. These results suggest that our approach can generalize to new velocities.

3) Performance with limited forces: In our experimental protocol we limited forces to 10 N in either the X or Z directions, stopping the actuator if this limit was exceeded. Note that the threshold used in the experiments of 10 N in either X or Z directions can result in a force magnitude greater than 10 N. During robot-assisted dressing, further limiting the magnitude of the forces could be beneficial for comfort and safety.

To evaluate how well our classifiers could perform with lower forces, we truncated each of the original 360 trials when the magnitude of the in-plane (in the X-Z plane) forces reached various thresholds. This results in trials comparable to what we would expect to obtain if we were to run our experiment with lower force thresholds, which would often result in the robot stopping earlier. We then used these truncated trials to test the performance of the 12 classifiers

we trained with leave-one-subject-out cross validation in Section V-A.1. Note that we made sure to only apply a classifier to trials that had not been used to train the classifier.

Table I and Figure 6 report the overall accuracy of our HMM classifiers for various thresholds on the force magnitude. These results show that data-driven haptic perception can still perform well when the magnitude of force applied by the robot is kept lower.

4) Comparing Univariate and Bivariate HMMs: We also performed the evaluation with limited forces from Section V-A.3 with classifiers composed of univariate HMMs. These univariate HMMs only used forces in the X direction or Z direction. This required training a total of 24 classifiers using leave-one-subject-out cross validation, 12 with forces in the X direction and 12 with forces in the Y direction. We present their performance in Figure 6.

The bivariate HMMs performed better in almost all cases than either of the univariate HMMs. This result suggests that the two directions of force provide complementary information. However, the force in the direction of gravity appears to be more informative.

5) Performance with limited distance: In our experimental protocol, we started the robot's end effector 0.44 m from the front of the participant's upper arm and moved it 0.85 m or until the forces in either the X or Z directions exceeded 10 N. During robot-assisted dressing, classifying the outcome earlier in the process could be advantageous. For example, the robot could more quickly take corrective actions or transition to the next subtask.

To evaluate how well our classifiers could perform earlier in the process, we truncated each of the original 360 trials at various distances of end effector travel. We then used these truncated trials to test the performance of the 12 classifiers we trained with leave-one-subject-out cross validation in Section V-A.1. Note that we made sure to only apply a classifier to trials that had not been used to train the classifier.

Figure 7 shows the performance of our bivariate HMM classifiers with respect to the travel distance of the robot's end effector. The figure also shows the positions of various relevant landmarks. For example, after traveling ~0.44m, the robot end effector would be above the average position of the inner elbow. The average stop position for the *caught* outcome is ~6.5cm past this point. The accuracy of the HMM classifiers at this position was 94.7%. These results show that data-driven haptic perception can accurately classify outcomes earlier in the dressing process.

VI. LIMITATIONS

We have provided evidence for the value of data-driven haptic perception for robot-assisted dressing through a carefully controlled experiment. In order to design an informative and replicable experiment, we deliberately focused on a representative subtask of dressing with a commonly used article of clothing. Given the strength of our results, we expect that haptic perception will be useful for other dressing subtasks with different articles of clothing. However, a number of factors might diminish the performance of haptic perception
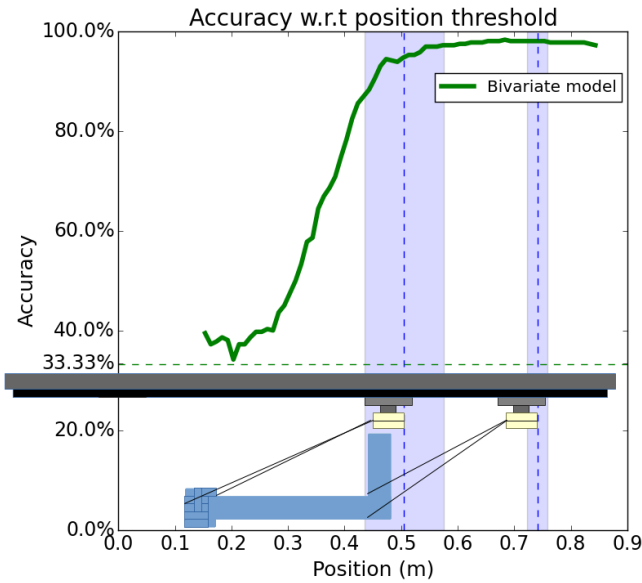
Fig. 7: This figure shows the performance of classifiers when the test data is truncated at various distances traveled by the robot's end effector. A diagram depicting an average human arm and the robot's end effector holding the gown is overlaid. The dotted lines from left to right represent the average stop position for the *caught* outcome and the average stop position for the *good* outcome. The end effector positions correspond to average stop points. The shaded regions represent one standard deviation of these stop locations. We marked 33% accuracy as the expected result of a random classifier.

or require different methods. Garments vary widely for functional and stylistic reasons. The geometry and materials of other garments could result in distinct dressing subtasks and measured forces. For example, an elastic cuff at the wrist of a sleeve would likely alter the forces over time. Likewise, jewelry could changes forces or catch on garments. During our experiments, an armrest supported the participant's upper arm, and the participant's forearm was initially aligned to the robot's motion. In other dressing tasks, body parts could have greater freedom to move with respect to the robot, resulting in more variability of the forces measured by the robot. Likewise, a robot might hold a garment in place while the person's body moves, which might increase the variability of the forces measured by the robot.

VII. DISCUSSION

Our method's high accuracy using HMMs, a well-established algorithm, supports the feasibility of data-driven haptic perception for robot-assisted dressing. Higher accuracy might be attainable with other algorithms. Likewise, using different features, instead of raw forces, might improve the performance. Also, moments and out-of-plane forces measured at the robot's end effector might be informative. We expect our results with conventional HMMs and raw force measurements in two directions to serve as a useful baseline for research that investigates other pattern recognition

algorithms and features.

We have used human participants, providing results that we expect to be generalizable to other robot-assisted dressing tasks. Some benefits of using human participants rather than an artificial limb are that human arms have fairly specific and interesting haptic responses. For example, for many participants, when the sleeve caught on the fist, the arm rotated from horizontal until the force experienced by the robot reached 10 newtons. Some participants allowed their arms to be rotated up to 45 degrees. Our participants spanned a range of arm sizes, shapes, and body hair. Additional HMMs for various *caught* subclasses might increase the accuracy of the system and allow a robot to better select recovery methods.

Although we designed our robot specifically for the experiment we conducted, specialized robots that operate in similar ways might be useful for dressing assistance, such as a simple robot for helping with socks and shoes.

A. Complexity of the Task

The measured forces did not clearly match patterns that we expected prior to conducting the experiment. Despite the controlled conditions of the experiment, we found the phenomena involved to be complicated. Figure 3 shows the means and standard deviations of the forces in the X direction and the Z direction during the three dressing outcomes. Although we are unsure of the underlying causes of the measured forces, we will now attempt to provide some physically plausible interpretations.

We would expect three main types of phenomena to influence the measured forces. First, we would expect positive changes to forces in the Z direction due the weight of the garment being transferred from the robot's end effector to the human body. Second, we would expect negative changes to forces in the X and Z directions due to the garment sliding along the human arm. Third, we would expect negative changes to forces in the X and Z directions due to the sleeve catching on the fist, elbow, or forearm.

We would expect negative changes to the forces in both the X and Z directions from sliding and catching due to the angle of the taut cloth between the arm and the end effector. As seen in Figure 1, the cloth tends to be pulled diagonally down (negative Z direction) and to the left (negative X direction) by the human arm. Given the flexibility of the cloth, one might think of the taut cloth as an ideal cable that can only apply a tensile load to the robot's end effector, resulting in a force pointing down and to the left. Since the gown is not made from a stretchy fabric, if it becomes caught and taut and the end effector continues to move, we would expect the forces to rapidly increase in magnitude.

Aspects of the specific outcomes in Figure 3 can be interpreted in this manner. For the *good* outcome, the change in force in the Z direction becomes positive, which is likely due to the weight of the gown being transferred from the robot's end effector to the human body. Interestingly, a similar effect in the *missed* outcome is not apparent. For the *caught* outcome, the change in force in the X and Y directions rapidly

becomes negative. At the end of the *good* outcome, the sleeve catches on the elbow, which appears to result in a less rapid negative change to the force in the X direction. Interestingly, an associated change is not apparent in the Z direction. For the *missed* outcome, the cloth does not catch and there is not a rapid negative change in the forces. For the *caught* outcome, the decreasing, increasing, and decreasing pattern may be related to some participants' forearms rotating upwards once the cloth became caught on their fists.

Despite the controls in place, there was also variation in the stop positions for the *good* and *caught* outcomes. When the fist caught on the opening of the sleeve, the average stop position was 0.475m with a standard deviation of 0.034m. The average stop position for the *good* outcome was 0.741m with a standard deviation of 0.018m. In less controlled environments, stop positions and force profiles might vary more widely.

Overall, the measured forces are complex and not easily interpretable. Attempting to use simple mechanical models or physical intuition to create classifiers by hand could be challenging, especially prior to seeing real-world data. Approaches that train haptic perception systems from data can circumvent these challenges. The good performance in our evaluation suggests that data-driven haptic perception is a viable way to handle the complexity of haptic signals measured during robot-assisted dressing.

VIII. CONCLUSION

We assessed the potential usefulness of data-driven haptic perception for robot-assisted dressing. Overall, our results suggest that haptic perception can help robots infer the relationship between clothing and the human body. For our dressing subtask, bivariate HMM classifiers performed well when tested with people and end-effector velocities outside of the training data. They also performed well when we limited the maximum force that the robot could apply. In addition, they accurately classified the relationship between the clothing and the human body early in the dressing task. Notably, our results suggest that forces in the direction of gravity can be especially informative.

ACKNOWLEDGMENT

REFERENCES

[1] J. M. Wiener, R. J. Hanley, R. Clark, and J. F. V. Nostrand, "Measuring the activities of daily living: Comparisons across national surveys," *Journal of Gerontology: SOCIAL SCIENCES*, vol. 45, no. 6, pp. S229–237, 1990.

[2] Institute of Medicine, *Retooling for an Aging America: Building the Health Care Workforce*. The National Academies Press, 2008.

[3] H. J. Goodin, "The nursing shortage in the United States of America: an integrative review of the literature," *J Adv Nurs*, 2003.

[4] T. L. Mitzner, T. L. Chen, C. C. Kemp, and W. A. Rogers, "Identifying the potential for robotics to assist older adults in different living environments," *International Journal of Social Robotics*, pp. 1–15, 2013.

[5] Canadian Partnership for Stroke Recovery. "Stroke Engine". [Online]. Available: http://www.strokengine.ca/

[6] The Wright Stuff, Inc. "The Wright Stuff healthcare products that make life easier". [Online]. Available: http://www.thewright-stuff.com/

[7] Y. Gao, H. J. Chang, and Y. Demiris, "User modelling for personalised dressing assistance by humanoid robots," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 1840–1845.

[8] T. Tamei, T. Matsubara, A. Rai, and T. Shibata, "Reinforcement learning of clothing assistance with a dual-arm robot," in *Humanoid Robots (Humanoids), 2011 11th IEEE-RAS International Conference on*. IEEE, 2011, pp. 733–738.

[9] S. D. Klee, B. Q. Ferreira, R. Silva, J. P. Costeira, F. S. Melo, and M. Veloso, "Personalized assistance for dressing users," in *Social Robotics*. Springer, 2015, pp. 359–369.

[10] N. Koganti, T. Tamei, T. Matsubara, and T. Shibata, "Estimation of human cloth topological relationship using depth sensor for robotic clothing assistance," in *Proceedings of Conference on Advances In Robotics*. ACM, 2013, pp. 1–6.

[11] N. Koganti, J. G. Ngeo, T. Tomoya, K. Ikeda, and T. Shibata, "Cloth dynamics modeling in latent spaces and its application to robotic clothing assistance," 2015.

[12] K. Yamazaki, R. Oya, K. Nagahama, K. Okada, and M. Inaba, "Bottom dressing by a life-sized humanoid robot provided failure detection and recovery functions," in *System Integration (SII), 2014 IEEE/SICE International Symposium on*. IEEE, 2014, pp. 564–570.

[13] D. Park, Z. Erickson, T. Bhattacharjee, and C. C. Kemp, "Multimodal execution monitoring for anomaly detection during robot manipulation," in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 2016.

[14] C. Piciarelli and G. L. Foresti, "On-line trajectory clustering for anomalous events detection," *Pattern Recognition Letters*, vol. 27, no. 15, pp. 1835–1842, 2006.

[15] A. Rodriguez, D. Bourne, M. Mason, G. F. Rossano, and J. Wang, "Failure detection in assembly: Force signature analysis," in *Automation Science and Engineering (CASE), 2010 IEEE Conference on*. IEEE, 2010, pp. 210–215.

[16] O. Rosen and A. Medvedev, "An on-line algorithm for anomaly detection in trajectory data," in *American Control Conference (ACC), 2012*. IEEE, 2012, pp. 1117–1122.

[17] A. Jain and C. C. Kemp, "Improving robot manipulation with data-driven object-centric models of everyday forces," *Autonomous Robots*, vol. 35, no. 2-3, pp. 143–159, 2013.

[18] P. Pastor, M. Kalakrishnan, S. Chitta, E. Theodorou, and S. Schaal, "Skill learning and task outcome prediction for manipulation," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 3828–3834.

[19] B. Redmond, R. Aina, T. Gorti, and B. Hannaford, "Haptic characteristics of some activities of daily living," in *Haptics Symposium, 2010 IEEE*. IEEE, 2010, pp. 71–76.

[20] T. Bhattacharjee, J. M. Rehg, and C. C. Kemp, "Haptic classification and recognition of objects using a tactile sensing forearm," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, pp. 4090–4097.

[21] K. Hawkins, C.-H. King, T. Chen, and C. C. Kemp, "Informing assistive robots with models of contact forces from able-bodied face wiping and shaving," in *21st IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2012.

[22] A. Jain and C. C. Kemp, "Improving robot manipulation with data-driven object-centric models of everyday forces," *Autonomous Robots*, vol. 35, pp. 143–159, 2013.

[23] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," in *Readings in Speech Recognition*, A. Waibel and K. F. Lee, Eds., Kaufmann, San Mateo, CA, 1990, pp. 267–296.

[24] T. Bhattacharjee, A. Kapusta, J. M. Rehg, and C. C. Kemp, "Rapid categorization of object properties from incidental contact with a tactile sensing robot arm," in *Humanoid Robots (Humanoids), 2013 13th IEEE-RAS International Conference on*. IEEE, 2013, pp. 219–226.

[25] V. Chu, I. McMahon, L. Riano, C. G. McDonald, Q. He, J. Perez-Tejada, M. Arrigo, N. Fitter, J. C. Nappo, T. Darrell, and K. J. Kuchenbecker, "Using robotic exploratory procedures to learn the meaning of haptic adjectives," in *Proceedings of International Conference on Robotics and Automation*, 2013.