

Towards Robotic Feeding: Role of Haptics in Fork-based Food Manipulation

Tapomayukh Bhattacharjee, Gilwoo Lee*, Hanjun Song*, and Siddhartha S. Srinivasa

Abstract—Autonomous feeding is challenging because it requires manipulation of food items with various compliance, sizes, and shapes. To understand how humans manipulate food items during feeding and to explore ways to adapt their strategies to robots, we collected a rich dataset of human trajectories by asking them to pick up food and feed it to a mannequin. From the analysis of the collected haptic and motion signals, we demonstrate that humans adapt their control policies to accommodate to the compliance and shape of the food item being acquired. We propose a taxonomy of manipulation strategies for feeding to highlight such policies. As a first step to generate compliance-dependent policies, we propose a set of classifiers for compliance-based food categorization from haptic and motion signals. We compare these human manipulation strategies with fixed position-control policies via a robot. Our analysis of success and failure cases of human and robot policies further highlights the importance of adapting the policy to the compliance of a food item.

Index Terms—Haptics and Haptic Interfaces, Force and Tactile Sensing, Perception for Grasping and Manipulation

I. INTRODUCTION

NEARLY 56.7 million (18.7%) among the non-institutionalized US population had a disability in 2010 [1]. Among them, about 12.3 million needed assistance with one or more activities of daily living (ADLs) or instrumental activities of daily living (IADLs). Key among these activities is feeding, which is both time-consuming for the caregiver and challenging for the care recipient to accept socially [2]. Although there are several automated feeding systems in the market [3]–[6], they have lacked widespread acceptance as they use minimal autonomy, demanding a time-consuming food preparation process [7] or pre-cut packaged food.

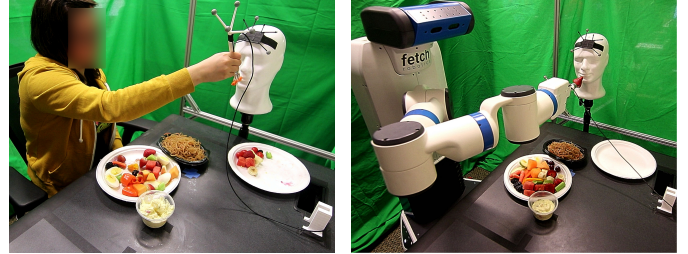
Eating free-form food is one of the most intricate manipulation tasks we perform in our daily lives, demanding robust nonprehensile manipulation of a deformable hard-to-model target. Automating food manipulation is daunting as the universe of foods, cutlery, and human strategies is massive. In this paper, we take a small first step towards organizing the science of autonomous food manipulation.

Manuscript received: September, 9, 2018; Revised December, 5, 2018; Accepted January, 3, 2019.

This paper was recommended for publication by Editor Allison M. Okamura upon evaluation of the Associate Editor and Reviewers' comments. This work was funded by the National Institute of Health R01 (#R01EB019335), National Science Foundation CPS (#1544797), National Science Foundation NRI (#1637748), the Office of Naval Research, the RCTA, Amazon, and Honda.

*These authors contributed equally to the work. All the authors are with Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, Washington 98195 {tapo, gilwoo, hanjuns, siddh}@cs.washington.edu

Digital Object Identifier (DOI): see top of this page.



(a) Human feeding experiment (b) Robot feeding experiment

Fig. 1: Examples of a feeding task with a dinner fork.

First, we collect a large and rich dataset of human strategies of food manipulation by conducting a study with humans acquiring different food items and bringing them near the mouth of a mannequin (Figure 1). We recorded interaction forces, torques, poses, and RGBD imagery from 3304 trials leading to more than 18 hours of data collection which provided us unprecedented and in-depth insights on the mechanics of food manipulation.

Second, we analyze our experiments to build a *taxonomy* of food manipulation, organizing the complex interplay between fork and food towards a feeding task. A key observation was that the choice of a particular control policy for bite acquisition depended on the *compliance* of the item. For example, subjects tilted the fork to prevent a slice of banana from slipping, or wiggled the fork to increase pressure for a carrot. Other feeding concerns, such as how the target would bite, were reflected in the manipulation strategies during both bite acquisition and transport. This key idea that people use compliance-based strategies motivated us to explore compliance-based food categorization. Food classification based on haptic and motion signals instead of only vision-based classification [8]–[10] is beneficial during food manipulation, as visually similar items may have different compliance and therefore may need different control policies. Temporal Convolutional Network [11] most successfully categorized food items in our experiments.

Third, we highlight the importance of choosing a compliance-based control policy by analyzing the performance of a fixed position-control strategy on a robot. The robot had more failures in picking up soft and hard-skinned items compared to human subjects who adapted their control policies to the item's compliance.

Food manipulation promises to be a fascinating new challenge for robotics. Our main contributions in this paper are a rich dataset, an analysis of food manipulation strategies towards a feeding task, an intuitive taxonomy, and a haptic analysis. We envision that a future autonomous robotic feeding system will use the data and taxonomy to develop a set of discrete manipulation strategies that depend on the class of



Fig. 2: Experimental setup with an instrumented fork (Forque) to acquire bites of different food items and feed a mannequin.

food items, methods from haptic classification to categorize a food item to one of these classes, and insights from the robot experiment to implement the control policies. This paper does not address the subtleties of interactions with an eater. We are excited about further work that builds upon these contributions towards a science of food manipulation.

II. RELATED WORK

Our work connects three areas of research: food manipulation, manipulation taxonomies, and haptic classification.

1) *Food manipulation*: Studies on food manipulation in the packaging industry [12]–[15] have focused on the design of application-specific grippers for robust sorting and pick-and-place. Crucially, not only did they identify the need for haptic sensing as critical for manipulating non-rigid food items, but they also pointed out that few manipulators are able to deal with non-rigid foods with a wide variety of compliance [12]–[15].

Research labs have explored meal preparation as an exemplar multi-step manipulation problem, baking cookies [16], making pancakes [17], separating Oreos [18], and preparing meals [7] with robots. Most of these studies either interacted with a specific food item with a fixed manipulation strategy [16], [17] or used a set of food items for meal preparation which required a different set of manipulation strategies [7]. Importantly, all of these studies emphasized the use of haptic signals (through joint torques and/or fingertip sensors) to perform key sub-tasks.

2) *Manipulation Taxonomies*: Our work is inspired by the extensive studies in human grasp and manipulation taxonomies [19]–[22] which have not only organized how humans interact with everyday objects but also inspired the design of robot hands and grasping algorithms [23].

However, unlike most of these studies, our focus is to develop an application-specific taxonomy focused on manipulating deformable objects for feeding. We believe this focus is critical as feeding is both a crucial component of our everyday lives and uniquely different in how we interact with the world. In that regard, our work echoes the application-specific work in human-robot interaction on *handovers*, also a crucial and unique act [24], [25], where the analysis and taxonomy of human-human handovers laid the foundation for algorithms for seamless human-robot handovers [24]–[26].

3) *Haptic Classification*: Most of the studies on haptic classification use specialized or distributed sensors on robot hands or fingertips for direct robot-hand and object interactions. Our work focuses on using a tool (Forque) to record forces and

motions of Forque-food interactions and addresses the problem of classifying food items. Researchers have previously used haptic signals to classify haptic adjectives [27], categorize rigid and deformable objects [28], recognize objects [29], [30] and for inferring object properties such as elasticity of deformable objects [31], hardness [32], and compliance [33], [34].

In a related work on meal preparation application, Gemici and Saxena [7] learn physical properties of 12 food items using end-effector forces, torques, poses, joint torques, and fingertip forces. However, they carefully designed the robotic actions (e.g. cut, split, flip-turn) using multiple tools (knife, fork, spatula) to extract meaningful sensor information to infer physical properties such as hardness, plasticity, elasticity, tensile strength, brittleness, and adhesiveness. Our objective is to classify food items into compliance-based categories using a variety of forces and motions that people use naturally when manipulating different food items for feeding.

III. HUMAN STUDY SETUP

We built a specialized test rig (Figure 2) to capture both motions and wrenches during a feeding task.

A. Forque: A Force-Torque fork sensor

We instrumented a dinner fork, *Forque*, to measure wrenches and motions (Figure 2(b)) generated during food manipulation. We selected an ATI Nano25 F/T sensor for 6-axis force/torque (F/T) measurements due to its minimal size, weight, appropriate sensing range and resolution for food manipulation. We designed the end of the Forque handle to attach spherical markers for motion capture with the NaturalPoint Optitrack system [35]. We designed Forque’s shape and size to mimic the shape and size of a real dinner fork. We 3D printed the handle and the tip of the Forque in plastic and metal respectively. A wire connecting the F/T Sensor with its Net F/T box runs along the length of the Forque along a special conduit to minimize interference while feeding and was reported to have little impact on subjects’ motion. We embedded the F/T sensor on the Forque instead of under the plate to record wrenches independent of a food item’s position on the plate (edge of the plate, center of the plate, on top of another food item etc.) and to record wrenches during the transport phase.

B. Perceptual data

To collect rich motion data, we installed 6 Optitrack Flex13 [36] motion capture cameras on a specially-designed

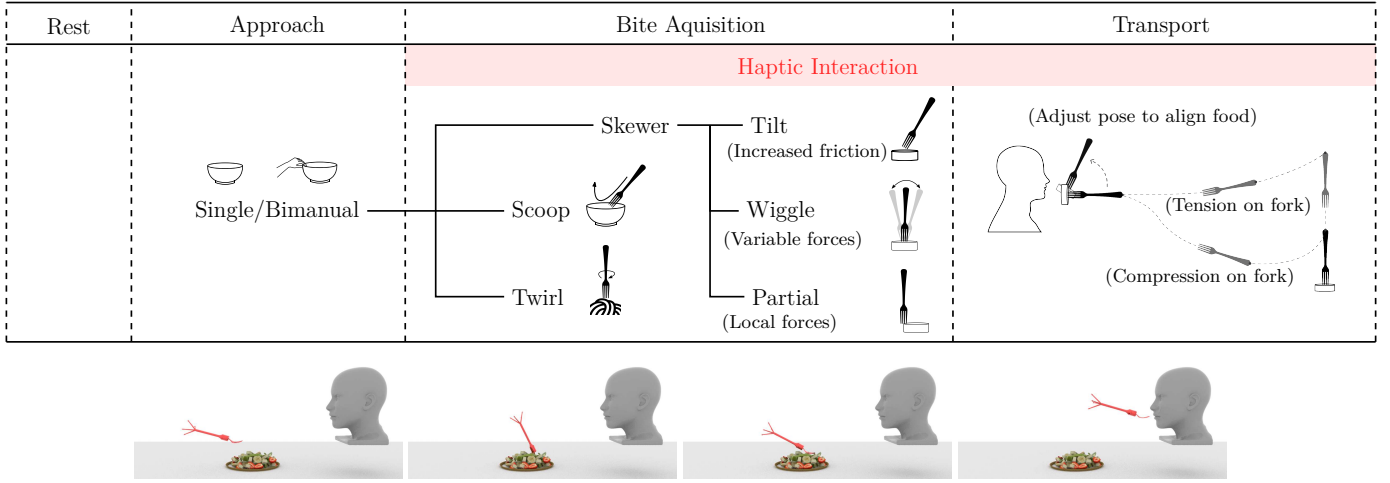


Fig. 3: A partial taxonomy of manipulation strategies relevant to a feeding task.

rig, with full coverage of the workspace. This provided full 6 DOF motion capture of the Forque at 120 frames per second (FPS). In addition, we installed a calibrated (both extrinsically and intrinsically) Astra RGBD [37] camera for recording the scene at 30 FPS, as well as a Canon DSLR RGB camera for recording videos for human labeling (Figure 2).

C. Data Collection

We selected 12 food items and classified them into four categories based on their compliance: *hard-skin*, *hard*, *medium* and *soft*. We had three food items for each of the four categories: *hard-skin* - bell pepper, cherry tomato, grape; *hard* - carrot, celery, apple; *medium* - cantaloupe, watermelon, strawberry; and *soft* - banana, blackberry, egg. We determined the classes of food items through mutual intercoder agreement. A primary and secondary coder (the main experimenter and the helper) independently skewered the food items and categorized the food items into arbitrary compliance-based classes. The primary and secondary coders completed two rounds of coding. After each round, the coders resolved discrepancies by adapting the number of classes and re-classifying the food items into these compliance-based categories. The second round of coding resulted in 100% intercoder agreement. Section VI-C further validates our categorization of the food items. In addition to these solid food items, we included noodles and potato salad (in separate containers), to diversify the manipulation strategies. Figure 2(c) shows typical plates of food offered to the subjects. We compiled the data as rosbag files using ROS Indigo on Ubuntu 14.04. The system clocks were synchronized to a Network Time Protocol server. We measured the average sensor delay between the Optitrack mocap signal and force/torque signal to be 30ms from 10 repeated trials. Our dataset is available at [38].

IV. HUMAN STUDY PROCEDURE

The task of each participant was to feed the mannequin. Before each experiment, we asked the participants to sign a consent form and fill a pre-task questionnaire. We asked our participants to pick up different food items from a plate or bowl using the *Forque* and *feed* a mannequin head as if they

were actually feeding a person. The head was placed at the height of a seated average human (Figure 2(a)).

For each *session*, we provided the participant with a plate of 48 pieces of food (4 pieces per item for 12 food items), a cup of potato salad, and a bowl of noodles. We asked each participant to pick up noodles and potato salad 4 times each to maintain consistency. Before each trial, a participant held the Forque at a predefined position marked on the table by a tape. When a computerized voice said “start” the participant could pick up any food item of their choice and feed the mannequin. After the participant brought the food item near the mouth of the mannequin, they waited until the experimenter said “stop”. They then discarded the food item and repeated another trial. We define a *trial* as one instance of feeding the mannequin, from “start” to “stop”.

There were $14 \times 4 = 56$ trials per session. Each participant had 5 sessions with a 2 to 5 minute break between each session, and each session began with a new plate (Figure 2(c)), giving us $56 \times 5 = 280$ trials per participant. We had 12 participants in the range of 18 - 62 years of age. This resulted in a grand total of $280 \times 12 = 3360$ trials. However, due to a technical glitch, we missed recording data for one of the sessions, thus giving us $3360 - 56 = 3304$ trials. For a left-handed participant, we inverted the experimental setup so that they could naturally feed the mannequin with their left hand. At the end of each experiment (after 5 sessions), we gave each participant a post-task questionnaire asking about their manipulation strategies during the task. The experiments were done in accordance with our University’s Institutional Review Board (IRB) review.

V. INSIGHTS FROM HUMAN SUBJECT EXPERIMENTS

Feeding is a complex task. Creating a taxonomy of manipulation behaviors for feeding is helpful in systematically categorizing it into sub-tasks. Segmentation allows us to better understand the different strategies people use in different phases of this task. We developed a partial taxonomy (Figure 3¹) of manipulation strategies relevant to a feeding task

¹Drawings in taxonomy are derivatives of “Fork” icon by Stephanie Szemetylo, “Bowl” icon by Anna Evans, “Hand” icon by Jamie Yeo, and “Noodle” icon by Artem Kovyazin [39].

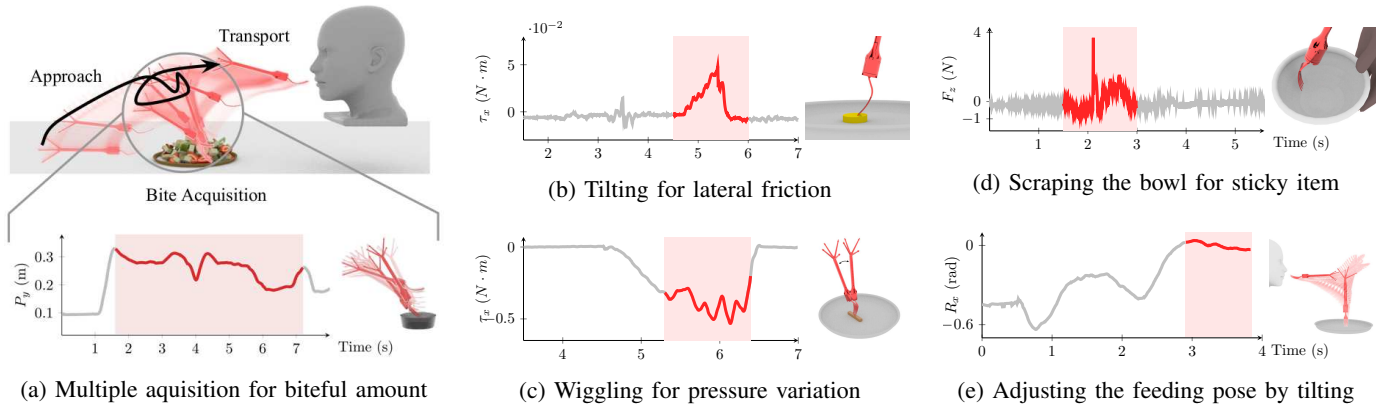


Fig. 4: Selected highlights: Different manipulation strategies in different feeding phases. F_z is the applied force on Forque’s z-axis, τ_x is the torque about Forque’s x-axis, P_y is the position of the Forque along global y-axis, and R_x is the rotation about global x-axis.

by dividing the feeding task into four primary phases: 1) rest, 2) approach, 3) bite acquisition, and 4) transport.

A. The rest phase: choose which item to pick up

We define the *rest* phase as the phase before any feeding motion is executed. During this phase, decisions such as which item to pick up are made.

B. The approach phase: prepare for bite acquisition

After choosing which item to pick up, the subject moves the Forque to acquire the item. We define the *approach* phase to be from the moment the subject starts moving the Forque until contact is made on the item. This phase serves as a key preparation step for successful bite acquisition. During this phase, the shape and size of the food item were a key factor in deciding the manipulation strategy.

1) *Subjects re-aligned the food for easier bite acquisition:* For food items with asymmetric shapes or irregular curvatures, such as celery, strawberry, or pepper, seven subjects used their Forque at least once to reorient the food items and expose a flat surface so that they could pierce the food item normal to the surface during bite acquisition.

2) *Subjects used environment geometry to stabilize the motion of oval food items for skewering:* For food items such as grapes, tomatoes, or hard-boiled eggs resting on high curvature surface, which tended to slip or roll, some subjects used the geometry of the plate (extruded edge) or other nearby food items as a support to stabilize the items. In one of the responses to the post-task questionnaire, one of the subjects mentioned, “I would ... corner it at the edge of the plate.” Five subjects used the environment geometry at least once to stabilize food items.

3) *Subjects used bimanual manipulation strategies to access difficult-to-reach items:* For containers with little potato salad or noodles, subjects applied bimanual manipulation strategies to access the food. They used one hand to tilt or hold the container, while the other hand scraped the food with the Forque, often using the container wall as a support (Figure 4(d)). All subjects used bimanual strategies at least once to either hold or tilt the container.

C. The bite acquisition phase: control positions and forces

Subjects used various position and force control strategies to acquire a bite. We define the *bite acquisition* phase to be from the moment the Forque is in contact with the food item until the item is lifted off from the plate (*liftoff*). During this phase, the *compliance* of food items was a key factor in deciding the control strategy. While simple vertical skewering was common for medium-compliance items, a few interesting strategies emerged for the hard-skin, hard, and soft categories. Also, the strategies for acquiring food items were influenced by the feeding task itself. In the post-task questionnaire, many subjects mentioned two key factors for feeding which affected their bite acquisition strategy: (a) ease of bite and (b) appropriate amount of bite.

1) *Subjects applied wiggling motions to pierce hard and hard-skin items:* Subjects skewered the hard and hard-skin food items using wiggling. Wiggling results in tilting the fork in various directions, which leads to fewer tines in contact with forces in variable directions and increased pressure. All subjects used this strategy at least once. Eight subjects used a wiggling motion to pierce the food items (Figure 4(c)). One of the subjects mentioned, “(I) sometimes needed to wiggle the fork back and forth to concentrate the pressure at only one tine to break through the skin of tomato, grape, etc.”

2) *Subjects skewered soft items at an angle to prevent slip:* For soft items such as slices of bananas which tended to slip out of the Forque tines during *liftoff*, subjects tilted the Forque (Figure 4(b)) to prevent slip by increasing friction using gravity. All subjects used this strategy at least once. For example, one of the subjects mentioned in the post-task questionnaire, “I would try to penetrate the fork at an angle to the food to minimize slices coming out.”

3) *Subjects skewered food items at locations and orientations that would benefit the feeding task:* For long and slender items, such as carrots, some subjects skewered it in one corner so that a person may be able to easily take a bite without hitting the Forque tines. This also played a role in selecting the orientation of the Forque when skewering the food item. For example, some subjects reported that they changed the orientation of the Forque before piercing a food item for ease of feeding. Eight subjects used these strategies.

4) *Subjects acquired food multiple times to feed an appropriate amount:* Acquiring an appropriate amount of food also

influenced the bite acquisition strategy. Although we never specified any specific amount per bite, six subjects attempted multiple scoops or twirls for noodle and potato salad to acquire an appropriate amount of food for a bite (Figure 4(a)).

D. The transport phase: feed the target

We define the *transport* phase as the phase after the food item is lifted from the plate until it is brought near the mouth of the mannequin.

1) *Subjects adapted their transport motion to prevent food from falling off:* Subjects adapted their motion (speed, angle, etc.) towards the mannequin to prevent the items from falling off. One subject mentioned, “I tried to be faster with eggs because they break apart easily and fall off the fork.” Another said, “With many softer foods (bananas specifically), I brought my arm up in a scooping motion to the mouth.” Depending on these subtle haptic cues, subjects varied the transport motion resulting in the application of either tensile forces or compressive forces on the fork and thereby kept a slippery food from falling off (Figure 3).

2) *Subjects oriented the Forque to benefit the feeding task:* While approaching the mannequin, the subjects oriented the Forque such that the item would be easy for a person to bite (Figure 4(e)). All subjects used this strategy. One of the subjects said, “I had to re-orient the fork often after picking food up in order to make it easier to bite for the humans.”

E. Subjects learned from failures

The subjects were not perfect in manipulating food items. For example, for small oval shaped food items with hard-skin, such as grapes and tomatoes, the food either slipped or rolled multiple times. When skewering halved hard-boiled eggs, the yolk was often separated from the white during liftoff. The subjects also dropped soft items multiple times. Even when the motion led to a successful bite acquisition, there were unintended results such as hitting the plate when piercing a hard-skin food item. This was probably because there was a mismatch between subjects’ initial estimations of the forces and motions required to pick up a food item and the actual physical interactions.

However, after a few failures, they changed their manipulation strategies. One subject mentioned, “The celery was harder than I was expecting. So, after a couple of times, I knew to exert more force.” Another subject mentioned, “The egg was tricky. I learned to spear it by the white part and the yolk at the same time to keep it together.” Yet another remarked, “I also learned to spear grapes by just one prong of the fork.” Out of all the trials when subjects learned from their previous failures and changed their strategy, 42.4% were for hard-skin, 29.2% for hard, 15.9% for soft, and 12.5% for medium food items. However, although there were various adaptations, subjects were never perfect in manipulating food items of varying compliance even after learning from failures.

F. Cultural influences and personal choices affected manipulation strategies

We observed interesting cultural factors that could affect the forces and motions of the feeding task. Some subjects

grasped the Forque much closer to the tines while others held it unusually high. Some subjects held the Forque at unusual rotations about its principal axis. Interestingly, subjects’ personal choices could also affect their manipulation strategies. For example, one subject mentioned, “(I) prefer [to] avoid yolk (I hate hard-boiled eggs).” We also observed that subjects picked up noodles using both clockwise and counter-clockwise twirls.

VI. HAPTIC CLASSIFICATION

One key observation from the study was that humans use compliance-based strategies for bite acquisition. To facilitate control policies based on compliance, we present haptic classification of food items into four compliance-based categories: hard-skin, hard, medium, and soft. Note, we used 12 solid food items for this experiment, thus resulting in 2832 trials (without potato salad and noodles).

A. Discriminative models using LSTM, TCN, and SVM

We use three discriminative models: Long Short Term Memory Networks (LSTM [40]), Temporal Convolutional Networks (TCN [11]), and Support Vector Machines (SVM [41]).

LSTMs are a variant of Recurrent Neural Networks (RNN) which have been shown to be capable of maintaining long-term information. At every time-step, an LSTM updates its internal states and outputs a categorical distribution across the four categories. We stacked two layers of LSTM with 50 layers, which is then connected to a rectified linear unit (ReLU) and a linear layer. We then performed a softmax operation to get the probability distribution.

Unlike an LSTM, which maintains an internal state, a Temporal Convolutional Network (TCN) takes the whole trajectory as one input. It learns kernels along the temporal dimensions and across features. We stacked four convolutional networks, each with one dimensional temporal kernels of window size 5. Between each layer, we performed one ReLU operation and max pooling of width 2. The final output is connected to a ReLU and a linear layer before performing a softmax operation. For the input of TCN, we scaled the temporal dimension of each time series feature to have 64 steps using bilinear interpolation, where 64 was chosen to approximately match the average temporal length of the data. Cross-entropy loss was used for LSTM and TCN.

For SVM, we interpolated each time series feature similar to that of TCN, concatenated the interpolated time series features to obtain a feature vector [41]–[43] and then used a linear kernel [44] to train the SVM classifier. We implemented LSTM, TCN using PyTorch [45], and SVM using scikit-learn [46].

B. Generative models using HMMs

To use hidden Markov models (HMMs) for classification, we train one HMM per food category [41], [47], [48]. We characterize an HMM model (λ) by $\lambda = (\mathbf{A}, \mathbf{B}, \pi)$ where \mathbf{A} is the state-transition matrix, \mathbf{B} defines the continuous multivariate Gaussian emissions, and π is the initial state distribution [41], [47], [48]. Let M be the number of food categories and let $\mathbf{O}_{\text{train}}$ be a training observation vector for

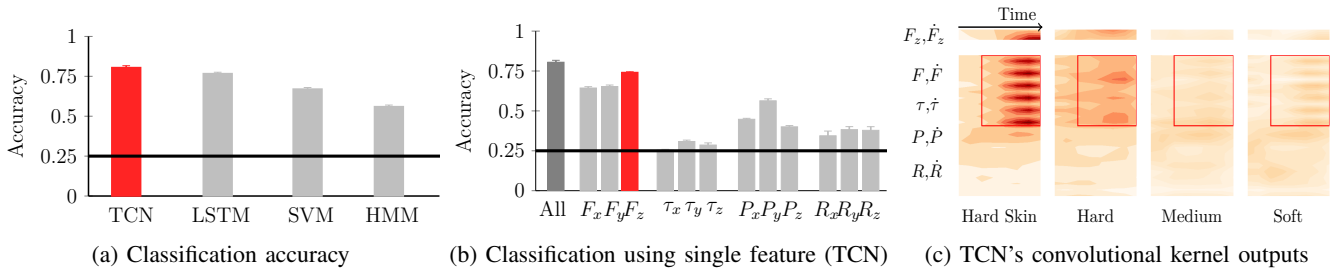


Fig. 5: Figure 5(a) compares 4 classifiers. Each classifier was trained with its best-performing feature set. TCN outperformed other classifiers. Figure 5(b) compares the predictive power of various features using TCN models. F and τ are the three forces and torques in Forque's local frame, P and R are the three positions and rotations in the global frame. Each feature includes its first-order derivative. Force along the principal axis of the Forque, F_z , is the most informative feature. Solid black lines show a random classifier's performance. Figure 5(c) shows TCN's convolutional layers' final output before its linear layers, indicating which (time, feature) pair contributes the most to classification. The most distinctive features are found in the later half of the time series in force and torque features (the red boxed regions).

True label	Predicted label			
	Hard-Skin	Hard	Medium	Soft
Hard-Skin	0.86	0.08	0.03	0.03
Hard	0.07	0.84	0.08	0.01
Medium	0.05	0.10	0.75	0.10
Soft	0.04	0.01	0.14	0.80

(a) Confusion matrix for human data

True label	Predicted label										
	Gra	Pep	Tom	App	Car	Cel	Can	Str	Wat	Ban	Egg
Grape	0.5		0.3								0.1
Pepper		0.7			0.1	0.1					
Tomato	0.3		0.6								
Apple		0.1		0.5	0.1	0.2	0.1	0.1	0.1		
Carrot		0.1		0.1	0.7	0.2					
Celery		0.1		0.2	0.2	0.5					
Cantaloupe				0.1			0.4	0.1	0.2		
Strawberry	0.1			0.1			0.1	0.5	0.1		0.1
Watermelon				0.1			0.2	0.1	0.4	0.1	0.1
Banana								0.1	0.6	0.1	0.2
Blackberry	0.1		0.1					0.1		0.1	0.7
Egg							0.1		0.1	0.2	0.1

(b) Confusion matrix of per-item recognition for human data

Fig. 6: Confusion matrices for haptic classification using TCN. Most confusion happens across nearby haptic categories, e.g. between hard-skin and hard, or medium and soft. In the per-item classification (Figure 6(b)), confusions across different categories are minimal compared to within-category confusion.

contact duration T . During training, we estimate the model parameters λ_m to locally maximize $P(\mathbf{O}_{\text{train}}|\lambda_m)$ using the iterative Baum-Welch method [41], [47], [48]. In our case, M is 4 (hard-skin, hard, medium, soft). For a test sequence \mathbf{O}_{test} , we assign the label (food category) $m^* \in M$ which maximizes the likelihood of the observation [41], [47]:

$$m^* = \arg \max_{m \in M} P(\mathbf{O}_{\text{test}}|\lambda_m)$$

We implemented HMMs using the GHMM library [49]. For each of the food category based HMMs, we optimized the number of hidden states to give maximum validation accuracy. This resulted in 3 hidden states for all the categories. These hidden states implicitly describe the Forque-food interaction once the Forque tines are inside the food item. We set a uniform prior to all the states.

C. Results

Figure 5(a) compares the performance of our four classifiers using 3-fold cross validation. For each classifier, we tested various combinations of feature sets and displayed the

one with the best performance. We tested with local forces, torques, global pose (positions and orientations) of the Forque, and their first-order derivatives as the features. For classifiers trained with multiple features of different magnitude scales, we normalized their values. TCN and LSTM performed the best with all features, while SVM and HMMs achieved the best performance with a combination of forces and positions. The best performing classifier was TCN with $80.47 \pm 1.17\%$ accuracy. Note that the HMM is a generative model unlike the other classifiers presented here and thus, it classifies by modeling the distributions of these 4 categories individually. The models are not optimized to maximize the discriminative aspects of these different categories. Using ANOVA and Tukey's HSD post-hoc analysis, we found significant differences between each classifier with $p < 0.0001$ at 95% CI. To analyze the importance of various features in classification, we compared the performance of TCN (the best performing classifier) when trained with different feature sets (Figure 5). It is evident that forces and positions are critical in the classification. In fact, the z -directional force, along the principal axis of the Forque, alone can correctly identify $74.22 \pm 0.29\%$ of the samples. Using ANOVA and Tukey's HSD for post-hoc analysis, we found significant differences between each feature with $p < 0.0001$ at 95% CI.

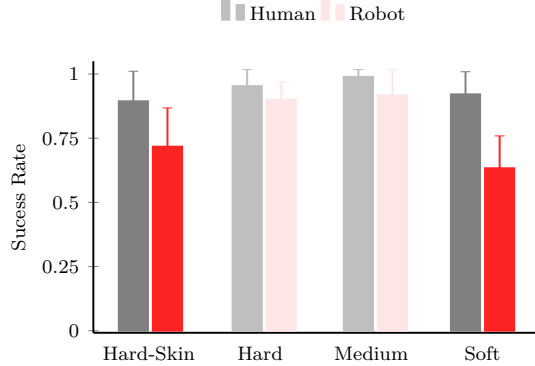
The confusion matrix in Figure 6(a) provides insights on where the classifier fails. The most confusion happens between nearby categories, e.g. between medium and soft, and hard-skin and hard which have similar haptic properties. The per-item classification (Figure 6(b)) further shows that items are most likely to be misclassified as items within the same class, which validates our compliance categories.

VII. ROBOT EXPERIMENTS

Human subjects used different forces and motions to acquire food items of varying compliance. Thus, a robot may benefit from choosing its manipulation strategy based on a compliance-based categorization, and learn to force-control as humans would. While we delegate the force-control policy learning as our future work, we performed the robot experiments to see if the robot could successfully feed the target using a fixed manipulation strategy with a position-control scheme and a vertical skewering motion. We used a Fetch robot with a 7 DOF arm. We modified the handle of the Forque so that it could be grasped by the robot's gripper. Our robot experimental setup was otherwise identical to the human setup.

	Hard-Skin	Hard	Medium	Soft
Hard-Skin	0.66	0.10	0.19	0.04
Hard	0.02	0.91	0.07	0.00
Medium	0.16	0.17	0.63	0.03
Soft	0.05	0.02	0.07	0.87

(a) Confusion matrix for robot data (TCN)



(b) Human and robot success rates

Fig. 7: Confusion matrix of the robot experiments show similar trends as that of human experiments. The robot success rate using a position-control scheme is lower than that of humans who controlled forces and motions to acquire food items of varying compliance.

A. Experimental Procedure

We programmed the robot using a programming by demonstration (PbD) technique [50] by saving a series of waypoints (joint configurations) of the arm through human demonstrations. We performed a total of 240 trials (4 categories \times 3 food items \times 4 pieces per food item \times 5 sessions). In each trial, the robot used a vertical skewering motion to pick up a food item from a pre-determined location on the plate. We randomly selected 4 such locations on the plate. After each trial, we discarded the skewered food item and manually placed another food item from the plate in that location for the next trial. After one session, we replaced the entire plate with a new plate and repeated this procedure for 5 sessions. We did not program the scooping and twirling motion, and thus, did not use noodles and potato salad for these experiments. We collected same modalities of data as during the human subject experiments.

B. Results

Figure 7(a) shows the confusion matrix using a 4-fold cross validation of robot experiments. When trained with a TCN on the robot data, we get $83.1 \pm 4.8\%^2$ accuracy, which shows that even with the position-control scheme, the observations from each category are different enough. However, the robot experiments and human experiments led to very different forces. Thus, the classifier trained on human subject data resulted in only $50.6 \pm 3.7\%$ accuracy when tested on robot data.

We also compared the bite acquisition success rates of humans and robots (Figure 7(b)). Subjects found it the most difficult to acquire hard-skin food items, whereas the robot with position-control scheme struggled with both hard-skin and soft food items. Using ANOVA and Tukey’s HSD post-hoc analysis for human studies, we found significant differences in success rates between hard-skin and hard categories ($p = 0.0006$), hard-skin and medium categories ($p < 0.0001$),

and medium and soft categories ($p = 0.0001$). For robot experiments, we found significant differences in success rates between hard and soft categories ($p = 0.0108$) and medium and soft categories ($p = 0.0067$). Using a different control policy affected the bite acquisition success rate. Figure 7(b) shows that the robot’s success rate in bite acquisition was lower than that of humans. One of the reasons could be because humans used varied forces and motions to pick up food items of different compliance (See Section V). Using a 2-tailed t-test, we found significant differences for hard-skin, medium, and soft categories ($p < 0.0025$) at 95% CI. This further shows the need for different manipulation strategies for different compliance-based categories, which we delegate as our future work for robot manipulation.

VIII. DISCUSSION

We performed two additional analyses to investigate the effect of speed on the choice of different manipulation strategies and different classes of food items. Using ANOVA and Tukey’s HSD post-hoc analysis, we found significant differences of speed between wiggling and every other manipulation strategy (skewering, scooping, twirling) with $p < 0.0167$ at 95% CI. This could be because of faster penetration during wiggling due to increased pressure. Similarly, we found significant differences of speed between all food categories except hard and hard skin categories with $p < 0.0001$ at 95% CI.

Note, bite timing is another important factor for feeding. The correct bite timing would depend on various factors such as if the care-recipient has finished chewing, if the care-recipient has finished talking to someone else etc. Since this paper does not focus on the eater interaction, this is outside the scope of this paper but is a subject of interest for our future work.

Haptics in the context of food manipulation is much less explored and hence, one of the focuses of this paper was to analyze the role of haptic modality. We envision our future robotic system to be multimodal using both vision and haptics with complementary capabilities. Relying only on visual modality may result in a suboptimal choice of manipulation strategy if two items look similar but have different compliance. A food item in a cluttered plate may not have clear line of sight or may have noisy depth image due to moisture content such as in watermelons. The presence of haptic modality can potentially alleviate these concerns by identifying a food item’s compliance class and thus reducing the uncertainty in choosing a manipulation strategy.

For haptic classification, a fork needs to be in contact with a food item. A prolonged penetration, as needed in the majority of haptic perception literature [27], [30], [34], makes it difficult to change the manipulation strategy on the fly. Our classification scheme is opportunistic and requires data only for the first 0.82s of skewering when the fork is going into the food. A robot could use vision to choose food-item dependent fork approach angles before contact based on our developed taxonomy and then use the haptic modality to refine its bite acquisition motion in case of anomalies or uncertainty. A future autonomous robotic system would use the data and taxonomy from the human experiment, methods from the haptic classification, and insights from the controlled

²Note, using a 3-fold crossvalidation scheme, we get lower accuracy of $65.8 \pm 5.3\%$ probably because of lack of data (20 trials per food item).

robot experiment to devise various manipulation strategies for feeding people food items of varying physical characteristics.

REFERENCES

- [1] M. W. Brault, "Americans with disabilities: 2010," *Current population reports*, vol. 7, pp. 70–131, 2012.
- [2] L. Perry, "Assisted feeding," *Journal of advanced nursing*, vol. 62, no. 5, pp. 511–511, 2008.
- [3] "Obi," <https://meetobi.com/>, [Online; Retrieved on 25th January, 2018].
- [4] "My spoon," <https://www.secom.co.jp/english/myspoon/food.html>, [Online; Retrieved on 25th January, 2018].
- [5] "Meal-mate," <https://www.made2aid.co.uk/productprofile?productId=8&company=RBF%20Healthcare&product=Meal-Mate>, [Online; Retrieved on 25th January, 2018].
- [6] "Meal buddy," <https://www.performancehealth.com/meal-buddy-system>, [Online; Retrieved on 25th January, 2018].
- [7] M. C. Gemici and A. Saxena, "Learning haptic representation for manipulating deformable food objects," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 638–645.
- [8] T. Brosnan and D.-W. Sun, "Improving quality inspection of food products by computer vision—a review," *Journal of food engineering*, vol. 61, no. 1, pp. 3–16, 2004.
- [9] S. Gunasekaran, "Computer vision technology for food quality assurance," *Trends in Food Science & Technology*, vol. 7, no. 8, pp. 245–256, 1996.
- [10] D. G. Savakar and B. S. Anami, "Recognition and classification of food grains, fruits and flowers using machine vision," *International Journal of Food Engineering*, vol. 5, no. 4, pp. 1–25, 2009.
- [11] C. Lea, R. Vidal, A. Reiter, and G. D. Hager, "Temporal convolutional networks: A unified approach to action segmentation," in *Computer Vision—ECCV 2016 Workshops*. Springer, 2016, pp. 47–54.
- [12] P. Chua, T. Ilshner, and D. Caldwell, "Robotic manipulation of food products—a review," *Industrial Robot: An International Journal*, vol. 30, no. 4, pp. 345–354, 2003.
- [13] F. Erzincanli and J. Sharp, "Meeting the need for robotic handling of food products," *Food Control*, vol. 8, no. 4, pp. 185–190, 1997.
- [14] R. Morales, F. Badesa, N. Garcia-Aracil, J. Sabater, and L. Zollo, "Soft robotic manipulation of onions and artichokes in the food industry," *Advances in Mechanical Engineering*, vol. 6, p. 345291, 2014.
- [15] P. Brett, A. Shacklock, and K. Khodabendehtloo, "Research towards generalised robotic systems for handling non-rigid products," in *ICAR International Conference on Advanced Robotics*. IEEE, 1991, pp. 1530–1533.
- [16] M. Bollini, J. Barry, and D. Rus, "Bakebot: Baking cookies with the pr2," in *The PR2 workshop: results, challenges and lessons learned in advancing robots with a common platform, IROS*, 2011, pp. 1–7.
- [17] M. Beetz, U. Klank, I. Kresse, A. Maldonado, L. Mösenlechner, D. Pangercic, T. Rühr, and M. Tenorth, "Robotic roommates making pancakes," in *IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2011, pp. 529–536.
- [18] "Oreo separator machines," <https://vimeo.com/63347829>, [Online; Retrieved on 1st February, 2018].
- [19] M. R. Cutkosky, "On grasp choice, grasp models, and the design of hands for manufacturing tasks," *IEEE Transactions on robotics and automation*, vol. 5, no. 3, pp. 269–279, 1989.
- [20] T. Feix, J. Romero, H.-B. Schmiedmayer, A. M. Dollar, and D. Kragic, "The grasp taxonomy of human grasp types," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 66–77, 2016.
- [21] J. R. Napier, "The prehensile movements of the human hand," *Bone & Joint Journal*, vol. 38, no. 4, pp. 902–913, 1956.
- [22] I. M. Bullock, R. R. Ma, and A. M. Dollar, "A hand-centric classification of human and robot dexterous manipulation," *IEEE Transactions on Haptics (TOH)*, vol. 6, no. 2, pp. 129–144, 2013.
- [23] M. T. Ciocarlie and P. K. Allen, "Hand posture subspaces for dexterous robotic grasping," *The International Journal of Robotics Research*, vol. 28, no. 7, pp. 851–867, 2009.
- [24] E. C. Grigore, K. Eder, A. G. Pipe, C. Melhuish, and U. Leonards, "Joint action understanding improves robot-to-human object handover," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 4622–4629.
- [25] K. W. Strabala, M. K. Lee, A. D. Dragan, J. L. Forlizzi, S. S. Srinivasa, M. Cakmak, and V. Micelli, "Towards seamless human-robot handovers," *Journal of Human-Robot Interaction*, vol. 2, no. 1, pp. 112–132, 2013.
- [26] M. Cakmak, S. S. Srinivasa, M. K. Lee, J. Forlizzi, and S. Kiesler, "Human preferences for robot-human hand-over configurations," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*. IEEE, 2011, pp. 1986–1993.
- [27] V. Chu, I. McMahon, L. Riano, C. G. McDonald, Q. He, J. M. Perez-Tejada, M. Arrigo, T. Darrell, and K. J. Kuchenbecker, "Robotic learning of haptic adjectives through physical interaction," *Robotics and Autonomous Systems*, vol. 63, pp. 279–292, 2015.
- [28] A. Drimus, G. Kootstra, A. Bilberg, and D. Kragic, "Classification of rigid and deformable objects using a novel tactile sensor," in *ICAR International Conference on Advanced Robotics*, 2011, pp. 427–434.
- [29] A. Schneider, J. Sturm, C. Stachniss, M. Reisert, H. Burkhardt, and W. Burgard, "Object identification with tactile sensors using bag-of-features," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 243–248.
- [30] P. K. Allen and K. S. Roberts, "Haptic object recognition using a multi-fingered dexterous hand," in *IEEE International Conference on Robotics and Automation*, 1989, pp. 342–347.
- [31] B. Frank, R. Schmiedding, C. Stachniss, M. Teschner, and W. Burgard, "Learning the elasticity parameters of deformable objects with a manipulation robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 1877–1883.
- [32] S. Takamuku, G. Gomez, K. Hosoda, and R. Pfeifer, "Haptic discrimination of material properties by a robotic hand," in *IEEE 6th International Conference on Development and Learning (ICDL)*, 2007, pp. 1–6.
- [33] M. Kaboli, P. Mittendorf, V. Hügel, and G. Cheng, "Humanoids learn object properties from robust tactile feature descriptors via multi-modal artificial skin," in *IEEE-RAS International Conference on Humanoid Robots*, 2014, pp. 187–192.
- [34] T. Bhattacharjee, J. M. Rehg, and C. C. Kemp, "Inferring object properties with a tactile-sensing array given varying joint stiffness and velocity," *International Journal of Humanoid Robotics*, pp. 1–32, 2017.
- [35] "Optitrak markers," <http://optitrack.com/products/motion-capture-markers/#mcm-12.7-m4-10>, [Online; Retrieved on 1st February, 2018].
- [36] "Optitrak flex 13 cameras," <http://optitrack.com/products/flex-13/>, [Online; Retrieved on 1st February, 2018].
- [37] "Orbbec astra," <https://orbbec3d.com/product-astra/>, [Online; Retrieved on 1st February, 2018].
- [38] T. Bhattacharjee, H. Song, G. Lee, and S. S. Srinivasa, "A dataset of food manipulation strategies," 2018. [Online]. Available: <https://doi.org/10.7910/DVN/8TTXZ7>
- [39] "Noun project - icons for everything," <http://thenounproject.com>, [Online; Retrieved on 1st February, 2018].
- [40] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [41] J. Wiens, E. Horvitz, and J. V. Gutttag, "Patient risk stratification for hospital-associated c. diff as a time-series classification task," in *Advances in Neural Information Processing Systems*, 2012, pp. 467–475.
- [42] M. Hoai, Z.-Z. Lan, and F. De la Torre, "Joint segmentation and classification of human actions in video," in *CVPR IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2011, pp. 3265–3272.
- [43] A. Bagnall, J. Lines, A. Bostrom, J. Large, and E. Keogh, "The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances," *Data Mining and Knowledge Discovery*, vol. 31, no. 3, pp. 606–660, 2017.
- [44] C.-W. Hsu, C.-C. Chang, C.-J. Lin, et al., "A practical guide to support vector classification," *Technical Report, Department of Computer Science, National Taiwan University*, pp. 1–16, 2003.
- [45] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," in *NIPS 2017 Autodiff Workshop: The Future of Gradient-based Machine Learning Software and Techniques, Long Beach, CA, US, December 9, 2017*.
- [46] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [47] M. W. Kadous et al., *Temporal classification: Extending the classification paradigm to multivariate time series*. PhD thesis, School of Computer Science Engineering, University of New South Wales, 2002.
- [48] 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.
- [49] "General hidden markov model library," <http://ghmm.org/>, [Online; Retrieved on 12th January, 2018].
- [50] S. Elliott, R. Toris, and M. Cakmak, "Efficient programming of manipulation tasks by demonstration and adaptation," in *IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 2017.