Balancing Efficiency and Comfort in Robot-Assisted Bite Transfer

Suneel Belkhale[1], Ethan K. Gordon[2], Yuxiao Chen[1],
Siddhartha Srinivasa[2], Tapomayukh Bhattacharjee[3], Dorsa Sadigh[1]
Stanford University[1], University of Washington[2], Cornell University[3]

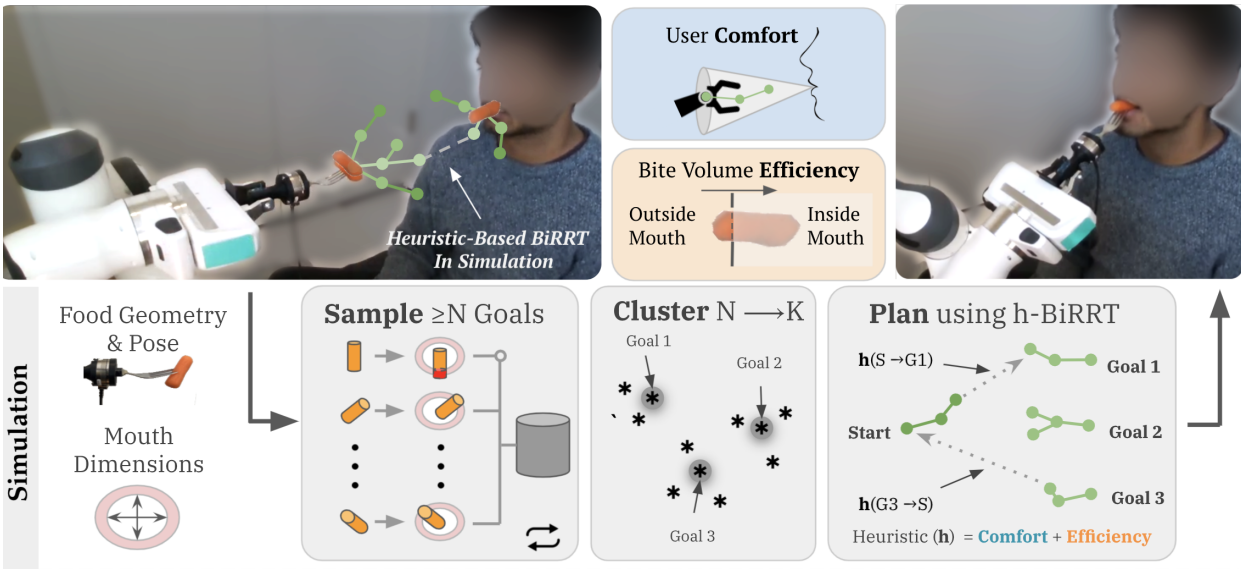

Fig. 1: Our method finds feasible bite transfer trajectories in simulation. Given the food geometry and pose on the fork, we sample at least N goal food poses that are checked for collisions with the mouth geometry using a learned constraint model. Next, we cluster the goal poses and use heuristic-guided BiRRT to reach cluster centroids with comfort (blue) and bite volume efficiency (orange) heuristics.

***Abstract*— Robot-assisted feeding in household environments is challenging because it requires robots to generate trajectories that effectively bring food items of varying shapes and sizes into the mouth while making sure the user is comfortable. Our key insight is that in order to solve this challenge, robots must balance the *efficiency* of feeding a food item with the *comfort* of each individual bite. We formalize comfort and efficiency as heuristics to incorporate in motion planning. We present an approach based on heuristics-guided bi-directional Rapidly-exploring Random Trees (h-BiRRT) that selects bite transfer trajectories of arbitrary food item geometries and shapes using our developed bite efficiency and comfort heuristics and a learned constraint model. Real-robot evaluations show that optimizing both comfort and efficiency significantly outperforms a fixed-pose based method, and users preferred our method significantly more than that of a method that maximizes only user comfort. Videos and Appendices are found on our website: https://tinyurl.com/bticra22.**

I. Introduction

Imagine a setting where you want to pick up a piece of food, e.g., a baby carrot, from a salad bowl to eat. Non-disabled people might overlook the complexity of this daily task — they might use a fork to pick up the carrot, while carrying a conversation and not paying as much attention on how the carrot is placed on the fork. Regardless of this placement, they move the fork in a manner that is not only efficient in how much food can be eaten but is also comfortable for the duration of the motion. This task presents numerous challenges for more than 12 million people with mobility-related disabilities [1]. Assistive robot arms have the potential to bridge this gap, and therefore provide care for those with disabilities. However, operating these arms can be challenging [2], [3]. In our initial surveys, people with mobility impairment mentioned the need for intelligent autonomy that optimizes comfort and adapts to the food item being fed. We envision intelligent algorithms that are aware of user comfort without the need for explicit user input. Achieving this level of autonomy presents a number of challenges which carry over to other robotics applications, including: 1) *perceiving* and choosing the next bite of food on a plate, 2) *acquiring* the food item with an appropriate tool, 3) *transferring* these items into the mouth in an efficient and comfortable manner. In recent years, there has been significant advances in food perception and acquisition [4], [5]. It turns out that the food acquisition strategy (e.g. fork skewering angle) heavily affects a user's comfort during bite transfer [5]; however, prior bite transfer methods rely on predetermined transfer trajectories for a discrete set of acquisition strategies and food geometries [6].

To handle a wide variety of food items and acquisition methods, a bite transfer strategy must optimize its trajectories on the fly by bringing food into a mouth without sacrificing user comfort. However, this is challenging with real world sources of variation (e.g. food geometries, sizes, acquisition poses on the fork, and mouth shapes). Even with one food geometry and acquisition pose, there are often many different "collision-free" paths into the mouth, so the feeding agent should filter this solution space intelligently. For instance,

consider a vertically aligned baby carrot oriented perpendicular to the fork, as shown in Fig. 1. There are a wide range of possible feeding paths; some may come too close to a person's face, affecting comfort, while others may only bring the tip of the carrot into the mouth, limiting bite volume.

Regardless of the orientation or type of food on our fork, caregivers will intuitively balance the bite volume *efficiency* for a single bite with the *comfort* of that bite. Motivated by this behavior, we present a bite transfer algorithm for selecting trajectories in a continuous space of mouth sizes, food geometries, and poses. Our approach (Section III) takes as input a food mesh and an acquisition pose on the fork from the real world, and generates an analogous simulation environment. We learn a constraint model to sample goal food poses near the mouth, and perform motion planning based on a novel set of heuristics (Section IV) to shape the perceived comfort and bite volume efficiency of each transfer. To our knowledge, our approach is the first to formulate comfort and efficiency for bite transfer, to consider non-bite sized food items, and to work for a continuum of acquisition poses and food geometries. We demonstrate our algorithm in practice through a limited user study (Section V). Our results show that while comfort alone and efficiency alone are able to outperform fixed trajectories on average, our approach of blending comfort and efficiency is the *only* method to outperform a fixed pose baseline with statistical significance. We run our method on various food items of differing geometries and scales in simulation (Appendix IV).

II. Related Work

Our work draws inspiration not only from the state-of-the-art in the robot-assisted feeding literature but also from the shared autonomy and general robot-human handovers.

Robot-assisted Feeding: Bite acquisition and transfer. Several specialized feeding devices for people with disabilities have come to market in the past decade. Although several automated feeding systems exist [7], [8], [9], [10], they lack widespread acceptance as they use minimal autonomy, demanding a time-consuming food preparation process [11], or pre-cut packaged food and cannot adapt the bite transfer strategies to large variations due to pre-programmed movements. Existing autonomous robot-assisted feeding systems such as [4], [5], [12], and [13] can acquire and feed a fixed set of food items, but it is not clear whether these systems can adapt to different food items that are either not bite-sized and require multiple bites or require other bite transfer strategies. Feng *et al.* [4] and Gordon *et al.* [14] developed an online learning framework using the SPANet network and showed *acquisition* generalization to previously-unseen food items, but did not address the *bite transfer* problem. Gallenberger *et al.* [5] showed a relationship exists between bite acquisition and transfer, but did not propose how to transfer bites for non bite-sized items in such a setting. Our paper aims to close this gap in bite transfer by developing a context-aware framework for robot-assistive feeding which generalizes to food items that are not bite-sized.

Shared Autonomy for Robotic Assistance. Adding autonomy to provide robotic assistance to tasks by inferring human intent is a well-studied field [15], [2], [16], [17], [18], [19], [20], [3]. This is especially relevant for precise manipulation tasks such as bite acquisition or bite transfer during robot-assisted feeding, while drawing parallels to other tasks such as peg-in-hole insertion [21], [22]. For example, there has been work on using the concept of shared autonomy for bite acquisition tasks such as stabbing a bite, scooping in icing, or dipping in rice [23], where the researchers combined embeddings from a learned latent action space with robotic teleoperation to provide assistance. Unlike this body of work, this paper focuses on completely autonomous bite transfer of food items by keeping in mind our end user population, which may have severe mobility limitations.

Robot-human Handovers. There are many works analyzing robot-human handovers, but most of the studies focus on objects that are handed over without using an intermediate tool [24], [25], [26] in a single attempt. In this paper, we focus on tool-mediated handover of food-items that may not be bite-sized, and thus may require multiple handover attempts. The feeding handover situation poses an additional challenge of transferring to a constrained mouth, instead of a hand [5]. Gallenberger et al. [5] explore the problem of bite-transfer by providing the insight that bite transfer depends on bite acquisition and thus the transfer trajectories are not only food-item dependent but are also based on how a food was acquired. Cakmak et al. [27] study the handover problem in an application-agnostic way, where they identify human preferences for object orientations and grasp types. Similarly, Aleotti et al. [28] confirmed that orienting items in specific ways can make handover easier. Canal et al. [29] take a step further and explore how bite transfer can change with personal preferences. In our paper, we focus on tool-mediated bite-transfer of food items that may not be bite-sized and hence may require multiple transfer attempts.

III. Context-Aware Multi-Bite Transfer

A caregiver can guide a food item into their patient's mouth agnostic to the orientation of the food on the fork — they do not spend minutes optimizing how the food should be placed on the fork for the most optimal transfer. In this section, we first formalize the goal of *food acquisition-agnostic* bite transfer, and then discuss our approach.

A. Bite Transfer Problem Formulation

To begin bite transfer iteration, we are given a 3D mesh $\mathcal{M}_{\text{food}}$ of the food item, the constant pose $p_f \in \mathbb{R}^6$ of the food item on the fork, a kinematics model for the robot and fork system with corresponding mesh $\mathcal{M}_R$, and the pose estimate of the mouth $p_m \in \mathbb{R}^6$. To capture the motion of the food item into the mouth, we want to find waypoints of the food item over time, represented by poses $p \in \mathbb{R}^6$. Additionally, we assume the mouth can be represented by a simple elliptical tube $\mathcal{M}_{\text{mouth}}$, where the ellipse axes are in the face plane, and open mouth dimensions $d_m \in \mathbb{R}^2$ are specified per end user. These inputs are visualized in our PyBullet-based simulation environment in Fig. 2. We outline our method for acquiring these inputs in Appendix III. Given these inputs, we formulate the goal of bite transfer

as finding a sequence of food poses $\mathcal{T} = \{p_0, \ldots, p_{L-1}\}$ of varying length L respecting a set of physical constraints $\mathcal{C}$ and cost function $\mathcal{J}(\mathcal{T})$, shown in Eq. (1). For the task of bite transfer, $\mathcal{C}$ consists of physical constraints. $\mathcal{C}_0$ (Eq. (2)) and $\mathcal{C}_1$ (Eq. (3)) ensure no collisions between the mouth mesh $\mathcal{M}_{\text{mouth}}$ of dimensions d_m with pose p_m and the food mesh $\mathcal{M}_{\text{food}}$ for each pose $p_i \in \mathcal{T}$ as well as the robot-fork mesh $\mathcal{M}_R$ respectively. $\mathcal{C}_2$ (Eq. (4)) constrains the final food pose to be near the mouth opening, i.e., $p_G \doteq p_{L-1}$ is in the support of goal pose distribution $\mathcal{D}_g$.

$$\mathcal{T}^* = \underset{\mathcal{T}}{\operatorname{argmin}}\, \mathcal{J}(\mathcal{T}) \text{ s.t. } \mathcal{C}_i(\mathcal{T}) = 1 \;\forall \mathcal{C}_i \in \mathcal{C} \tag{1}$$

$$\mathcal{C}_0(\mathcal{T}) = \mathcal{M}_{\text{food}}(p_j) \cap \mathcal{M}_{\text{mouth}}(p_m) = \emptyset \;\forall p_j \in \mathcal{T} \tag{2}$$

$$\mathcal{C}_1(\mathcal{T}) = \mathcal{M}_R(p_j) \cap \mathcal{M}_{\text{mouth}}(p_m) = \emptyset \;\forall p_j \in \mathcal{T} \tag{3}$$

$$\mathcal{C}_2(\mathcal{T}) = p_G \in \mathcal{D}_g \tag{4}$$

B. Approach Overview

When taking a bite, a person intuitively *simulates* the physics of their mouth's interaction with a carrot on our fork, regardless of the carrot's orientation, or where their arm starts. They might initially visualize where the carrot should be in the mouth and work backwards to find the most comfortable and efficient path. Our approach captures this intuition. Our simulation environment (see Fig. 2) reflects the real world setup, where the mouth is replaced with a static elliptical mouth model, allowing us to simulate the interactions between the human mouth and the food item.

Our approach in Fig. 1 consists of three phases: sampling, clustering, and planning. Since the space of feasible goal food poses in the mouth is continuous, we outline two efficient goal sampling methods (*Projection* & *Learned Constraints*), which leverage simulation to batch sample from a set of "feasible" orientations and offsets from the mouth, defined by the distribution $\mathcal{D}_g$, and then check these samples against the constraints $\mathcal{C}$ to generate a varied set of feasible goal food poses p_G near the mouth. Next, we cluster the constraint satisfying poses into a set of K goals with broad coverage over $\mathcal{D}_g$. We use heuristic guided bi-directional rapidly-exploring random trees to search for paths to goal food poses within the mouth that respect the physical constraints $\mathcal{C}$. We guide the addition of new nodes to the h-BiRRT with a cost-to-come function h and a cost-so-far function g, where the sum of h and g defines the overall predicted cost $\mathcal{J}$ of a node in the h-BiRRT produced graph:

$$\mathcal{J}(\mathcal{T}) \approx \mathcal{J}(\{p_0 \ldots p_i\}, p_G) = g(p_0 \ldots p_i) + h(p_i, p_G) \tag{5}$$

Section IV discusses how we incorporate comfort and efficiency into h and g. Here, we first outline a method for generating goal poses $p_G \sim \mathcal{D}_g$ to satisfy the constraints $\mathcal{C}$.

Sampling Food Objects with Projection. When sampling goal food poses, there are certain fork orientations that are impossible or unsafe for the arm to reach (e.g. the fork pointing backwards relative to the face). We restrict the orientations of the robot end effector to be within a spherical cut centered on the into-mouth axis, and position offsets from the mouth center are bounded. These bounds form the uniform goal distribution $\mathcal{D}_g$. The full sampling algorithm is outlined in the appendix in Algorithm 1. We first generate batches of food goal poses from $\mathcal{D}_g$ and check for collision,

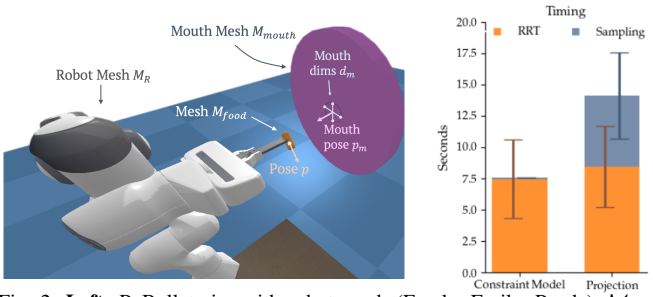

Fig. 2: **Left:** PyBullet sim with robot mesh (Franka Emika Panda) $\mathcal{M}_R$, food object mesh $\mathcal{M}_{\text{food}}$ (e.g. carrot) at pose p, and mouth mesh $\mathcal{M}_{\text{mouth}}$ (cylindrical tube, radii from d_m) at pose p_m. **Right:** End-to-end algorithm timing for learned constraint model compared to projection-based sampling (100 trajectories each).

repeating this process until reaching N collision-free samples or timing out. Since a person's true mouth cavity fits within the tube-like elliptical mouth in simulation, which has a constant cross section in the mouth plane (see Fig. 2), we accelerate the 3D collision check by slicing the food by the mouth plane and then projecting the inner food mesh vertices for each goal pose onto the mouth plane (*Projection*). The second image from the left in Fig. 7 (Appendix I) shows the slicing plane, with a sample carrot geometry. Finally we can check if the vertices are within the 2D mouth cross section to detect if the goal pose is collision-free.

Improved Sampling via Learned Constraints. While *Projection* checks samples for collision faster than a naïve 3D collision check, it still has significant and high variance lag (Fig. 2). We thus propose a sampling method, *Learned Constraints*, that learns to predict constraint values (e.g., collision prediction) from 1M sim samples, with model inputs (d_m, $\mathcal{M}_{\text{food}}$, p_f, p_G). In Fig. 2, the right plot shows that a learned collision predictor significantly reduces sampling time. In Appendix I, we provide further details and show that *Learned Constraints* maintain sample quality (e.g., predictive accuracy) and end-to-end trajectory performance (e.g., comfort & efficiency costs).

Clustering Goal Food Poses. Once we have timed out or reached N collision free samples, we consolidate these goal poses into a representative set over $\mathcal{D}_g$ for the planning step. We use a standard implementation of k-mediods, although any mediod clustering algorithm can be substituted.

Motion Planning with Heuristic-Guided BiRRT. Once collision-free goal food poses have been generated and clustered, we must find trajectories to reach these goals. We adapt Rapidly-exploring Random Trees (RRT) for our motion planning. Inspired by Lavalle et al. [30], who used bi-directional search ideas to grow two RRTs, we used one tree from the start state p_0 and the other from the goal state p_G^k. To bias the two search trees towards each other, we take a heuristics-based approach [31]. See Appendix II for details on our heuristic-guided implementation. Designing the heuristic cost functions will be discussed next.

IV. Comfort & Efficiency in Motion Planning

The solution space of feasible transfer trajectories is often large: a small strawberry can be eaten in a wide variety of fork orientations due to its size and inherent symmetry.

Our approach narrows this solution space with *comfort* and *efficiency* heuristics during motion planning. One intuitive formulation is the path cost g being the distance between food poses (Eq. (6)), with the heuristic h being the distance to the goal pose (Eq. (7)).

$$g(p_0 ... p_i) = \sum_{j=0}^{i-1} ||p_{j+1} - p_j|| \tag{6}$$

$$h(p_i, p_G) = ||p_G - p_i|| \tag{7}$$

While this cost function guides hRRT to the goal pose, finding the shortest distance path in food pose space ignores both comfort and bite volume efficiency. Consider the vertically oriented carrot in the first row of Fig. 3. If we sample a goal pose with the carrot oriented into the mouth and just past the teeth, a straight path to this goal is optimal in distance cost, but the person can barely take a bite; to add, the end effector would be close to the user's face, which could be considered uncomfortable. From our initial surveys with users with mobility limitations, we indeed conclude that comfort and efficiency are essential during bite transfer. One participant comments, "*The orientation should be comfortable for the utensil and the food ... the arm should maintain a low profile to not obstruct sight ... it should be fast...*". To this end, we develop two competing cost functions to shape the trajectories produced by hRRT: (1) bite efficiency, capturing the percentage of the food inside the mouth at the end of a trajectory, and (2) trajectory comfort, capturing the perceived user comfort along a given trajectory.

A. Modeling Efficiency

The most *efficient* goal pose brings the most food into a person's mouth, which can be measured in the real world by comparing the food mesh before and after each bite. In simulation, we approximate the new food mesh without knowing the biting physics for a user. Instead, we assume a bite slices the food mesh in the face plane (see Figure 7). Let V_i be the volume of the food geometry, and V_f be the remaining volume after the bite. We estimate the efficiency cost of goal poses with Eq. (8). The n-root ($n = 3$ in practice) of the volume ratio amplifies the cost difference between goal poses of lower final volumes (high efficiency) to more noticeably bias RRT growth towards the most efficient goal poses. The resulting costs are in Eq. (9) & (10).

$$\mathcal{J}_E(p_G) = (V_f / V_i)^{1/n} \tag{8}$$

$$g(p_0 ... p_i) = \sum_{j=0}^{i-1} ||p_{j+1} - p_j|| \tag{9}$$

$$h(p_i, p_G) = ||p_G - p_i|| + \beta_E \mathcal{J}_E(p_G) \tag{10}$$

Note that this cost only considers the goal pose, rather than the entire trajectory. We empirically found that other notions of efficiency applied to *paths*, like trajectory execution time, or the distance of the path traveled, do not vary as much between outputs of h-BiRRT, and so yield less impact on the quality of trajectories produced.

B. Modeling Comfort and Personal Space

A trajectory that brings the arm too close within a user's *personal space* could influence the user's perceived safety, even if the transfer efficiency is high. We develop a notion of comfort that draws from proxemics literature in human-robot interaction, a well-studied field [32] for tasks such as in social robot navigation [33]. Building off the notion of "personal space", we hypothesize that trajectories should stay within a conic region stemming from the mouth, with a wide cross-sectional area further from the face that narrows towards the mouth. Prior work in human factors for social navigation has shown that a person's comfortable personal space can be different for each cardinal direction, usually being larger within a person's visual field than outside [33]. Building on this intuition, we skew the cone down, away from the visual field. We define a spatial cost function resembling an elliptical Gaussian at each cross section centered along the mouth axis (Fig. 4). We posit that the upward direction relative to a person's face, which is closer to the visual field, should penalize deviation from the mouth axis more than in the downward direction. For a distance $z \in \mathbb{R}^+$ along the mouth axis and offset from the mouth axis $x \in \mathbb{R}^2$ in the cross section plane, we define the spatial comfort cost in Eq. (11).

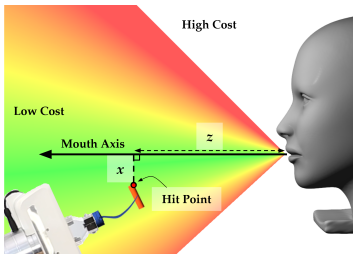

Fig. 4: Spatial comfort cost (red higher, green lower). The steeper cost gradient in the upward direction than downward ensures trajectories near the face (e.g., Fig. 3) have high "comfort" cost.

$$\mathcal{J}_C^s(x, z) = 1 - e^{-\alpha \frac{x^T \Sigma(x) x}{z^2}} \tag{11}$$

$$\mathcal{J}_C(p) = \frac{1}{NM} \sum_{h_j \in H} \mathcal{J}_C^s(h_j) \tag{12}$$

Here, $\Sigma(x)$ is a piece-wise covariance matrix in the face plane. In our experiments, we used a diagonal covariance matrix, with smaller variances above the mouth horizontal plane than below, and equal variances left and right. This cost and the mouth axis can be visualized in Fig. 4. Our comfort cost is applied on both the food item mesh and the entire simulated robot mesh and fork. In essence, we create a low resolution depth image from the perspective of the mouth and apply our cost function on the 3D location of each pixel. For a given food pose p and the corresponding simulated robot mesh, we cast rays in simulation along the mouth axis, starting at an $N \times M$ grid of points relative to the mouth center and on the face plane (e.g. $z = 0$), ending at a fixed maximum distance along the mouth axis $z_{\max}$. The set of hit points from this ray cast, denoted $H = \{h_j \in \mathbb{R}^3\}$, are passed into the cost function in Eq. (11) and normalized by the total number of points (Eq. (12)). This comfort cost is incorporated as a distance-weighted edge cost in hRRT with weight γ_C, and can be included in the heuristic as an additional goal cost using weighting β_C (Eq. (13) & (14)).

$$g(p_0 ... p_i) = \sum_{j=0}^{i-1} ||p_{j+1} - p_j|| \cdot (1 + \gamma_C \mathcal{J}_C(p_j, p_{j+1})) \tag{13}$$

$$h(p_i, p_G) = ||p_G - p_i|| + \beta_C \mathcal{J}_C(p_G) \tag{14}$$

Here, $\mathcal{J}_C(p_j, p_{j+1})$ is shorthand for the comfort cost at

Fig. 3: **Left**: The fundamental trade off between average comfort and efficiency costs for a grid of chosen relative weightings of comfort and efficiency, with costs from our h-BiRRT method averaged over 500+ initial food geometries and poses in simulation. Teal represents high ratios of comfort to efficiency, and orange the opposite. Our weights are the green dot at the elbow of this trade-off, achieving low efficiency cost and low comfort cost. **Right**: Sample trajectories for Fixed Pose (top) and Comfort+Efficiency (bottom) for the **Vertical** food geometry (see Section V). While Fixed pose (baseline) incurs high comfort cost (close to user's face), our method finds a trajectory that is both comfortable and efficient for the user.

the midpoint of these two food poses. We denote this formulation as "comfort only," since there is no consideration of efficiency here. Incorporating comfort alone can yield trajectories that keep the robot within the cone comfort region, but often this generates final goal poses that would not be easy to bite. Next, we will discuss incorporating both efficiency and comfort as costs for h-BiRRT.

C. Trading off Comfort and Efficiency

Ideally, an assistive robot would be able to feed bites of food with both comfort *and* efficiency in mind. In order to maximize both comfort and efficiency, we can put together the comfort costs (Eq. (13) & (14)) and efficiency costs (Eq. (9) & (10)), yielding the cost functions for h-BiRRT in Eq. (15) & (16), where the weightings β_E, β_C, and γ_C emphasize the efficiency at the goal, comfort at the goal, and comfort along the trajectory, respectively.

$$g(p_0 ... p_i) = \sum_{j=0}^{i-1} ||p_{j+1} - p_j|| \cdot (1 + \gamma_C \mathcal{J}_C(p_j, p_{j+1})) \quad (15)$$

$$h(p_i, p_G) = ||p_G - p_i|| + \beta_C \mathcal{J}_C(p_G) + \beta_E \mathcal{J}_E(p_G) \quad (16)$$

In Fig. 3, we plot the average comfort and efficiency scores for our h-BiRRT pipeline over a grid of weight values for β_E, β_C, and γ_C and over a large number of initial food poses and geometries (e.g., carrots, strawberries, celery, cantaloupes) in simulation. Refer to Appendix IV for quantitative results and example trajectories for each food type. Optimizing for efficiency only finds trajectories with the highest comfort costs but lowest efficiency costs, and vice versa for comfort only. This demonstrates that there is in fact a trade-off in comfort and efficiency costs when running h-BiRRT with our heuristic functions. Our approach will choose the "elbow" of this trade-off, balancing both efficiency and comfort.

V. User Study

A. Experimental Setup

We conducted a user study with six non-disabled participants to evaluate the perceived comfort and efficiency with our real world setup [1] In Appendix III, we discuss our real world system design, and how we ensure user safety during our user studies. Key parameter choices are shown in Table II in the Appendix. We consider carrots of varying sizes and fixed acquisition poses, visualized in the first row of Fig. 5: **Vertical**, **Horizontal**, **Roll & Pitch**, and **Yaw**. Users were instructed to sit still facing the robot, and to take a bite of each food item after each trajectory if they felt comfortable to do so. In addition, an emergency stop button was placed next to them for added assurance. See Appendix V for more user study details. We evaluated the following methods:

1) **FixedPose** (F): We fix the final orientation of the food item independent of the pose of the food item on the fork and the food size. This final orientation is hard-coded for a specific type of food and is inspired by the taxonomy of food manipulation strategies developed in [6].
2) **EfficiencyOnly** (E): Our approach with h-BiRRT and only efficiency costs, Eq. (9) and (10).
3) **ComfortOnly** (C): Our approach with h-BiRRT and only comfort costs, Eq. (13) and (14).
4) **Comfort+Efficiency** (CE): Our approach with both efficiency and comfort in mind: the h-BiRRT cost functions use both efficiency and reward, Eq. (15) and (16).

For each food pose and method, we evaluate two trajectories end-to-end with each user. After two trajectories for a given method, we ask a series of questions to gauge the user's perceived comfort of each trajectory, and the ease with which they were able to take a bite. We compare responses to these questions, in terms of **Comfort** (the average user rating of *comfort* for each evaluated trajectory, from 1 to 5, with 5 being the most comfortable), **Ease** (average rating of their *ease* of taking a bite for each evaluated trajectory, normalized from 1 to 5, with 5 being the best), **Rank** (Average relative rank of each method, from 1 to 4), and **Safety** (Average user rating of safety, from 1 to 5).

B. Results

The ease rating, comfort rating, and approach rank are summarized in Fig. 5. Importantly, our real world evaluation pipeline (Appendix III) was perceived as safe to the user regardless of food geometry or method, achieving an average **Safety** rating of 4/5. Despite the limited sample size, our method (CE) significantly outperforms the fixed baseline (F) for all three metrics (example in Fig. 3), and consistently outperforms comfort-only (C) and efficiency-only (E), significantly so in Rank ratings. Additionally, efficiency-only (E) did not perform as well as comfort-only in Comfort ratings. This supports our hypothesized connection between

[1] We decided to recruit non-disabled participants due to Covid-19 & safety concerns. Please see Appendix V for further discussion.

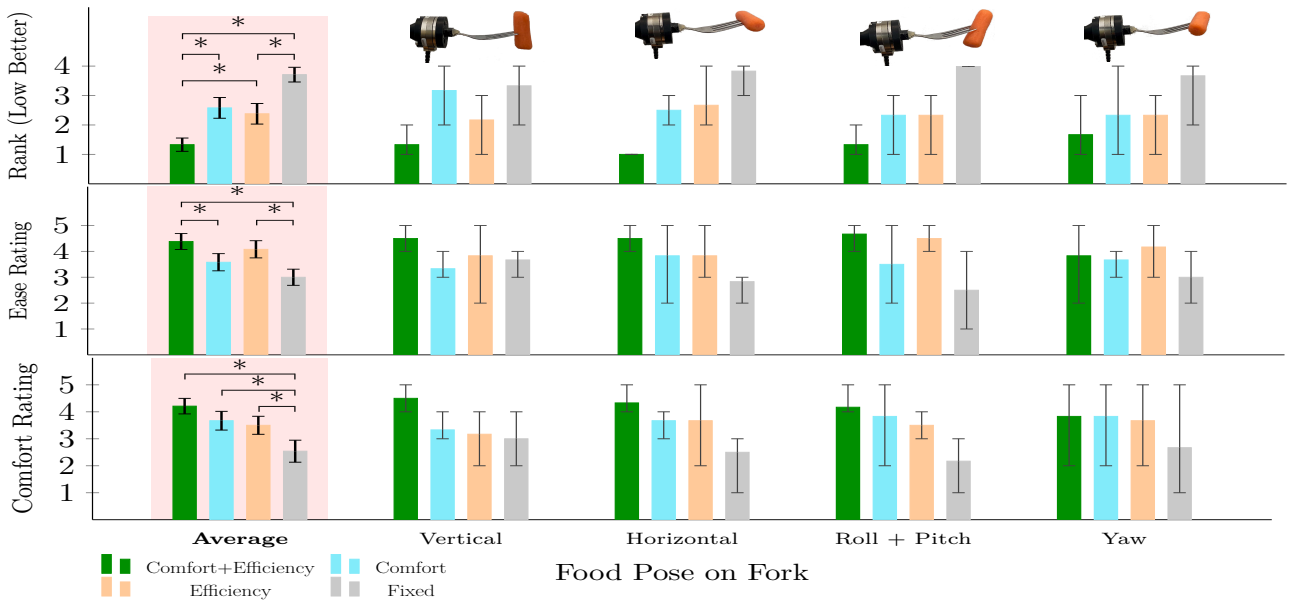

Fig. 5: User study quantitative results. Each non-highlighted plot shows the average comfort rating, ease rating, and rank between trajectory types across 4 different food poses, with the range across all 6 users plotted as error bars. The highlighted plots on the left show the average across all food poses, with error bars representing 95% confidence. Significant results, as determined by two-way ANOVA with repeated measures, Tukey HSD test, and Bonferroni correction ($P < 0.01$), are marked with an asterisk. We do not treat multiple ratings as independent. Despite the limited sample size ($N = 6$), trajectories from the combined comfort and efficiency method perform significantly better than the baseline fixed pose approach across all three metrics. Notably, the efficiency-only method often performs worse than comfort-only in comfort ratings. See Appendix V for significance testing details and more analysis.

user comfort perception and our comfort model (Eq. (11)).

The data is consistent with our hypothesis that, while optimizing over individual metrics (C and E) provides some improvement over the baseline, joint optimization performs even better in creating trajectories robust to real-world variation. We suspect that this is due to the large space of possible trajectories, where maximizing for only comfort puts no guarantee on efficiency, and vice versa.

Qualitatively, our comfort model's sensitivity to objects above mouth level fits with user expectations. When asked about low-ranked trajectories, users stated that they believed "*the robot should have approached from underneath,*" or that they "*didn't like when [the robot] came up close to [their] face.*" Users were more likely to instinctively move backwards when approached from above, near the face, and lean in when approached from below. In Fig. 6, we visualize a sample real world trajectory produced by each method for a **Vertical** pose. **FixedPose** is neither maximally efficient (carrot only partially fits in mouth) nor comfortable (robot is too close to face). Common quantitative metrics like time and path length are not as informative in gauging comfort, so we limit our evaluation to these qualitative metrics. Appendix V elaborates on quantitative and qualitative metrics and outlines how our approach naturally extends to the multi-bite setting with sample real world evaluations.

VI. Discussion

Summary. We present an approach based on motion planning for bite transfer under a continuous space of possible acquisition angles. During planning, we narrow down the solution space of possible trajectories into the mouth with an awareness of both bite *efficiency* and user *comfort*. Our user study demonstrates that considering comfort and efficiency jointly provides *significantly* more preferable trajectories compared to a fixed pose baseline. Furthermore, our method with comfort and efficiency consistently outperforms considering only comfort or only efficiency.

Limitations. One limitation of our method is the assumption that the mouth can be represented by a rigid elliptical tube, and that the food item is also rigid. In reality, the human mouth and the food item can both be deformable, which expands the set of "collision-free" paths into the mouth.

Furthermore, our user study only involved six non-disabled users due to Covid-19 related policies. In future work we plan to evaluate with more users, including users with mobility-impairment disabilities. However, we are excited that even with the given sample size, our method improves on the state-of-the-art with statistical significance.

Acknowledgements

This work is funded by NSF Award Numbers 2132847 and 2006388, and by the Office of Naval Research.

(a) Fixed Pose (b) Efficiency Only (c) Comfort Only (d) Comfort + Efficiency

Fig. 6: Example trajectories optimized for each metric. The efficiency metric (b) rotates the carrot sideways so as much can be consumed in one bite as possible. The comfort metric (c) penalizes more complicated trajectories where the robot body is likely to encroach on the face. The combined metric (d) results in a fairly straight trajectory that still ends with a sideways carrot.

REFERENCES

[1] M. W. Brault, "Americans with disabilities: 2010," *Current population reports*, vol. 7, pp. 70–131, 2012.

[2] D. P. Losey, K. Srinivasan, A. Mandlekar, A. Garg, and D. Sadigh, "Controlling assistive robots with learned latent actions," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 378–384.

[3] M. Li, D. P. Losey, J. Bohg, and D. Sadigh, "Learning user-preferred mappings for intuitive robot control," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, October 2020.

[4] R. Feng, Y. Kim, G. Lee, E. K. Gordon, M. Schmittle, S. Kumar, T. Bhattacharjee, and S. S. Srinivasa, "Robot-assisted feeding: Generalizing skewering strategies across food items on a realistic plate," *ISSR*, June 2019.

[5] D. Gallenberger, T. Bhattacharjee, Y. Kim, and S. Srinivasa, "Transfer depends on acquisition: Analyzing manipulation strategies for robotic feeding," ACM/IEEE International Conference on Human-Robot Interaction, 2019.

[6] T. Bhattacharjee, G. Lee, H. Song, and S. S. Srinivasa, "Towards robotic feeding: Role of haptics in fork-based food manipulation," *IEEE Robotics and Automation Letters*, 2019.

[7] https://meetobi.com/,[Online; Retrieved on 25th January, 2018].

[8] 2018, https://www.secom.co.jp/english/myspoon/food.html,[Online; Retrieved on 25th January, 2018].

[9] 2018, https://www.made2aid.co.uk/productprofile?productId=8&company=RBF%20Healthcare&product=Meal-Mate.

[10] 2018, https://www.performancehealth.com/meal-buddy-system,[Online; Retrieved on 25th January, 2018].

[11] 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.

[12] D. Park, Y. K. Kim, Z. M. Erickson, and C. C. Kemp, "Towards assistive feeding with a General-Purpose mobile manipulator," May 2016.

[13] L. V. Herlant, "Algorithms, Implementation, and Studies on Eating with a Shared Control Robot Arm," Ph.D. dissertation, 2016.

[14] E. K. Gordon, X. Meng, T. Bhattacharjee, M. Barnes, and S. S. Srinivasa, "Adaptive robot-assisted feeding: An online learning framework for acquiring previously unseen food items," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.

[15] A. D. Dragan and S. S. Srinivasa, "A policy-blending formalism for shared control," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 790–805, 2013.

[16] R. M. Aronson, T. Santini, T. C. Kübler, E. Kasneci, S. Srinivasa, and H. Admoni, "Eye-hand behavior in human-robot shared manipulation," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, pp. 4–13.

[17] D. Gopinath, S. Jain, and B. D. Argall, "Human-in-the-loop optimization of shared autonomy in assistive robotics," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 247–254, 2016.

[18] C.-M. Huang and B. Mutlu, "Anticipatory robot control for efficient human-robot collaboration," in *2016 11th ACM/IEEE international conference on human-robot interaction (HRI)*. IEEE, 2016, pp. 83–90.

[19] S. Javdani, H. Admoni, S. Pellegrinelli, S. S. Srinivasa, and J. A. Bagnell, "Shared autonomy via hindsight optimization for teleoperation and teaming," *The International Journal of Robotics Research*, vol. 37, no. 7, pp. 717–742, 2018.

[20] K. Muelling, A. Venkatraman, J.-S. Valois, J. E. Downey, J. Weiss, S. Javdani, M. Hebert, A. B. Schwartz, J. L. Collinger, and J. A. Bagnell, "Autonomy infused teleoperation with application to brain computer interface controlled manipulation," *Autonomous Robots*, vol. 41, no. 6, pp. 1401–1422, 2017.

[21] M. A. Lee, Y. Zhu, K. Srinivasan, P. Shah, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, "Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8943–8950.

[22] V. Unhelkar, C. Guan, N. Roy, and J. Shah, "Enabling robot teammates to learn latent states of human collaborators."

[23] H. J. Jeon, D. P. Losey, and D. Sadigh, "Shared autonomy with learned latent actions," *arXiv preprint arXiv:2005.03210*, 2020.

[24] A. Agah and K. Tanie, "Human interaction with a service robot: mobile-manipulator handing over an object to a human," in *Proceedings of International Conference on Robotics and Automation*, vol. 1, April 1997, pp. 575–580 vol.1.

[25] A. Edsinger and C. C. Kemp, "Human-robot interaction for cooperative manipulation: Handing objects to one another," in *RO-MAN 2007 - The 16th IEEE International Symposium on Robot and Human Interactive Communication*, Aug 2007, pp. 1167–1172.

[26] M. Huber, M. Rickert, A. Knoll, T. Brandt, and S. Glasauer, "Human-robot interaction in handing-over tasks," in *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, Aug 2008, pp. 107–112.

[27] 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.

[28] J. Aleotti, V. Micelli, and S. Caselli, "Comfortable robot to human object hand-over," in *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*, Sept 2012, pp. 771–776.

[29] G. Canal, G. Alenyà, and C. Torras, "Personalization framework for adaptive robotic feeding assistance," in *Social Robotics*, A. Agah, J.-J. Cabibihan, A. M. Howard, M. A. Salichs, and H. He, Eds. Cham: Springer International Publishing, 2016, pp. 22–31.

[30] J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 2. IEEE, 2000, pp. 995–1001.

[31] C. Urmson and R. Simmons, "Approaches for heuristically biasing rrt growth," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, vol. 2. IEEE, 2003, pp. 1178–1183.

[32] L. Takayama and C. Pantofaru, "Influences on proxemic behaviors in human-robot interaction," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2009, pp. 5495–5502.

[33] R. Kirby, R. Simmons, and J. Forlizzi, "Social robot navigation," Ph.D. dissertation, 2010.