

# Learning Sequential Acquisition Policies for Robot-Assisted Feeding

Anonymous Author(s)

Affiliation

Address

email

1       **Abstract:** A robot providing mealtime assistance must perform specialized  
2 maneuvers with various utensils in order to pick up and feed a range of food  
3 items. Beyond these dexterous low-level skills, an assistive robot must also plan  
4 these strategies in sequence over a long horizon to clear a plate and complete a meal.  
5 Previous methods in robot-assisted feeding introduce highly specialized primitives  
6 for food handling without a means to compose them together. Meanwhile, existing  
7 approaches to long-horizon manipulation lack the flexibility to embed highly  
8 specialized primitives into their frameworks. We propose Visual Action Planning  
9 Over Sequences (**VAPORS**), a framework for long-horizon food acquisition.  
10 **VAPORS** learns a policy for high-level action selection by leveraging learned  
11 latent plate dynamics in simulation. To carry out sequential plans in the real world,  
12 **VAPORS** delegates action execution to visually parameterized primitives. We  
13 validate our approach on complex real-world acquisition trials involving noodle  
14 acquisition and bimanual scooping of jelly beans. Across 38 plates, **VAPORS**  
15 acquires much more efficiently than baselines, generalizes across realistic plate  
16 variations such as toppings and sauces, and qualitatively appeals to user feeding  
17 preferences in a survey conducted across 49 individuals. Code, datasets, videos,  
18 and supplementary materials can be found on our [website](#).

19       **Keywords:** Deformable Manipulation, Dexterous Manipulation

## 20 1 Introduction

21 Millions of people are impacted logistically, socially, and physically by the inability to eat  
22 independently due to upper mobility impairments, age or health-related changes, or simply the  
23 toll of meal preparation [1, 2, 3]. Robot-assisted feeding has the potential to greatly improve the  
24 quality of life for these individuals while reducing caregiver burden. However, realizing a performant  
25 system in practice remains challenging. For instance, humans eat spaghetti noodles as shown in [Fig. 1](#)  
26 using nuanced fork-twirling motions. Other dishes, such as ramen, require even more diverse  
27 strategies—scooping soup, acquiring noodles, or acquiring meat. Thus, not only must an autonomous  
28 feeding system employ various utensils and strategies to handle different foods and quantities, but it  
29 must also operate over long horizons to finish a meal.

30 Prior assistive feeding work has focused on learning individual low-level vision-parameterized  
31 primitives for food manipulation. Examples include separate policies for skewering [4, 5, 6],  
32 scooping [7], bite transfer [8, 9, 10], cutting [11, 12, 13], and pushing food piles [14]. While  
33 highly specialized, these policies cannot reason over an extended horizon or make use of multiple  
34 strategies for more effective plate clearance. Humans, on the other hand, interleave *acquisition*  
35 and *rearrangement* actions with ease—pushing multiple peas together before scooping instead of  
36 painstakingly acquiring each individual pea or gathering noodles closer to each other before twirling  
37 with a fork. Replicating this long-horizon foresight in robotic feeding has yet to be demonstrated.

38 Recent work in skill-based reinforcement learning (RL) provides a natural way to model long-  
39 horizon manipulation sequences hierarchically. This entails first learning a high-level policy for

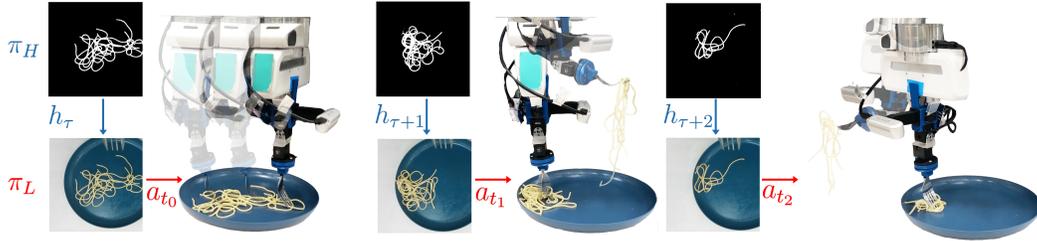


Figure 1: **Visual Action Planning Over Sequences (VAPORS)** employs a high level policy  $\pi_H$  to select amongst discrete manipulation strategies  $h$ , such as grouping and twirling, and a low-level vision-parameterized policy  $\pi_L$  to execute these actions  $a_t$  for long-horizon dexterous food acquisition.

40 composing skills [15, 16, 17], and then optionally inferring the parameters of low-level skills  
 41 separately [18, 19, 20]. These approaches tend to favor learning from simulation to scale data  
 42 collection [21], but current state-of-the-art simulators lack high-fidelity models for food deformation,  
 43 visuals, and cutlery interaction. This complicates learning food manipulation policies in simulation  
 44 and transferring them to the real world [13]. Existing hierarchical approaches also assume that  
 45 the low-level skills come from a general-purpose library of primitives such as grasping and path  
 46 planning [19, 22, 18, 15, 23], limiting their applicability to the food domain which requires highly  
 47 specialized behaviors. Thus, we seek to find an appropriate layer of abstraction for feeding, which can  
 48 leverage the benefits of (1) hierarchical planning for long-horizon manipulation and (2) vision-based  
 49 primitives for fine-grained control. Our key insight is that learning from simulated experience only  
 50 at a *high-level*, which need not capture the intricacies of food dynamics, and incorporating visual  
 51 planning to instantiate *low-level* specialized primitives, yields a powerful approach to dexterous,  
 52 multi-step food manipulation.

53 In this work, we present **VAPORS**: Visual Action Planning Over Sequences, a unified framework  
 54 for food manipulation. Our approach is decoupled into a high-level planner, which sequentially  
 55 composes low-level primitives. We first learn a policy in simulation that models latent dynamics of  
 56 plates from images. Specifically, we use segmented image observations as a representation space,  
 57 which captures the distribution of food items and is transferable between simulation and reality for  
 58 high-level plans. We train the policy using model-based RL with a reward that encourages both  
 59 acquisition and rearrangement. Separately, we instantiate a library of specialized primitives in the real  
 60 world from learned food pose estimation and segmentation. Finally, we use the learned high-level  
 61 planner on segmented real food images to plan sequences of primitives for long-horizon acquisition.

62 We experimentally validate our approach on two real food manipulation tasks: robotic noodle  
 63 acquisition and bimanual scooping. Across both real-world trials and a comprehensive user study  
 64 of 49 users, **VAPORS** achieves the highest efficiency, plate clearance, and qualitative user ratings  
 65 compared to heuristic and single-primitive baselines, all while generalizing to unseen plates.

## 66 2 Related Work

67 **Robot-Assisted Feeding.** Recently, a number of devices for mealtime assistance have become  
 68 available on the market [24, 25], but are limited in functional reach due to reliance on pre-programmed  
 69 trajectories or teleoperation by users. While *bite transfer* of a food item to a user’s mouth is the  
 70 eventual goal of autonomous feeding [9, 8, 10], we focus on *bite acquisition* as a primary initial  
 71 step for downstream feeding. Prior work in bite acquisition demonstrates the effectiveness of visual  
 72 planning for precise manipulation. Feng et al. [6], Gordon et al. [5, 26] and Sundaresan et al. [4]  
 73 leverage bounding box localization, food pose estimation, and visual servoing to geometrically plan  
 74 precise fork skewering motions. Similarly, Grannen et al. [7] and Suh and Tedrake [14] plan bimanual  
 75 scooping and grouping actions, respectively, for segmented food piles. These works focus only on  
 76 developing a specialized individual primitive for food manipulation. In isolation, this does not capture  
 77 many long-horizon real-world feeding scenarios with multiple utensils and strategies.

78 **Long-Horizon Planning and Control.** Several recent frameworks tackle long-horizon manipulation  
 79 by separating motion-level decision-making from sequential plans. Traditionally, task-and-motion-

80 planning (TAMP) approaches tend to assume extensive domain knowledge including after-effects  
 81 of actions and fixed task plans [27, 28, 29, 30, 31, 28]. In feeding, plate dynamics can be highly  
 82 uncertain, and state estimation is notoriously challenging, rendering these approaches ineffective.  
 83 An alternative approach is model-based planning and control, with recent impressive results on  
 84 complex tasks like dough manipulation [16, 32, 33]. This family of methods leverage learned  
 85 environment dynamics over visual states like images [34, 35, 36, 37, 33], keypoints [38], or particle-  
 86 based representations [16, 32] to sample and plan action sequences that maximize predicted rewards.  
 87 However, these methods do not scale well to high-dimensional continuous action spaces such as that  
 88 of food acquisition. To address this, hierarchical RL decouples policies into a high-level planner  
 89 which selects amongst discrete but parameterized low-level primitives [39]. These works have  
 90 demonstrated promising results on simulated long-horizon tabletop manipulation [18, 19, 15], but  
 91 have yet to consider (1) real-world deployment beyond carefully controlled experimental setups,  
 92 or (2) complex manipulation beyond commonplace primitives like pick-place, path-planning, and  
 93 grasping. In contrast, we consider highly diverse real plates requiring specialized primitives and  
 94 tools.

95 **Learning and Control for Manipulation in the Real World.** A large body of robotics research  
 96 focuses on learning real-world policies for manipulation either through sim-to-real transfer or  
 97 exclusively from real interactions. With sufficient domain randomization, sim-to-real transfer has  
 98 proven effective for tasks involving rigid objects or a limited set of deformable items like cloth, which  
 99 state-of-the-art simulators support [40, 41, 42]. However, adapting these simulators to modeling  
 100 food appearance and deformation is highly non-trivial. Meanwhile, learning exclusively from real  
 101 data has been shown to work well in challenging domains such as semantic grasping [43] or cable  
 102 untangling [44, 45, 46]. These approaches rely on state representations that are scalable to learn,  
 103 such as descriptors learned from self-supervised interaction [43] or keypoints learned from a small  
 104 amount of manually annotated images [47, 48, 49, 50]. In our setting, it is difficult to scale real-world  
 105 data collection across the range of food shapes, appearances, and properties a robot may encounter.  
 106 Self-supervised learning is also complicated due to resets and utensil interchange. We instead take a  
 107 hybrid approach which takes advantage of simulation for modeling high-level plate dynamics from  
 108 large-scale interactions, but leverages visual planning at the low level for precise real manipulation.

### 109 3 Problem Statement

110 We formalize the long-horizon food acquisition setting by considering an agent interacting in a  
 111 finite-horizon Partially Observable Markov Decision Process (POMDP). This is defined by the tuple  
 112  $(\mathcal{S}, \mathcal{O}, \mathcal{A}, \mathcal{T}, \mathcal{R}, T, \rho_0)$ . We assume access to plate image observations  $o_t \in \mathbb{R}_+^{W \times H \times C} = \mathcal{O}$  of  
 113 unknown plate states  $\mathcal{S}$ , with the initial state distribution given by  $\rho_0$ . Here,  $W$ ,  $H$ , and  $C$  denote  
 114 the image dimensions.  $\mathcal{A}$  denotes the action space, and  $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$  represents the unknown  
 115 transition function mapping states and actions to future states. The time horizon  $T$  denotes the  
 116 discrete budget of actions to clear the plate and  $\mathcal{R}(s, a)$  refers to the reward which measures progress  
 117 towards plate clearance. Our goal is to learn a policy  $\pi(a_t|o_t)$  that maximizes expected total return:  
 118  $\mathbb{E}_{\pi, \rho_0, \mathcal{T}}[\sum_t R(s_t, a_t)]$ , with  $t \leq T$ .

119 To do so, we decouple  $\pi$  into separate high and low-level sub-policies. We assume access to  $K$  discrete  
 120 manipulation primitives  $h^k$ ,  $k \in \{1, \dots, K\}$ , and learn a high-level policy  $\pi_H$  which selects amongst  
 121 these primitives. Additionally, we learn a low-level policy  $\pi_L$  which continuously parameterizes a  
 122 selected primitive according to visual input. The components we aim to learn are summarized below,  
 123 where  $h^k$  denotes a discrete primitive type and  $a_t$  denotes its continuous low-level instantiation:

$$\text{High-level policy : } \pi_H(h^k|o_{\leq t}, a_{\leq t-1}) \quad \text{Low-level policy : } \pi_L(a_t|o_t, h^k)$$

124 We consider low-level actions  $a_t$ , parameterized by the position of the tip of a utensil  $(x, y, z)$  and  
 125 utensil roll and pitch  $(\gamma, \beta)$ . Here,  $\beta = 0^\circ$  corresponds to an untilted fork handle, for instance, and  
 126  $\gamma = 180^\circ$  corresponds to the fork tines being horizontal when viewed top-down (Fig. 2).  
 127

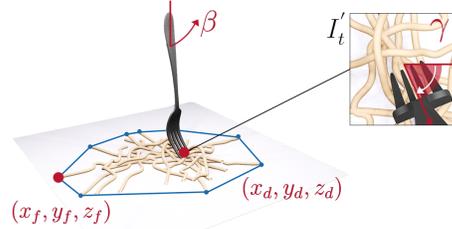
#### 128 3.1 State-Action Representations

129 In this section, we outline the visual state and action representations which are at the core of our  
 130 learning approach introduced in Section 4.

131 **Visual State Space.** Our approach makes use of RGB-D images and segmented plate observations,  
 132  $I_t \in \mathbb{R}_+^{W \times H \times 3}$ ,  $D_t \in \mathbb{R}^{W \times H}$ ,  $M_t \in \mathbb{R}_+^{W \times H}$  at different levels of abstraction. We leverage binary  
 133 segmentation masks to capture the spread of food items on a plate, informing high-level planning  
 134 with  $\pi_H$ , and RGB-D observations as input to  $\pi_L$  which better capture fine geometric details of food.

135 **Action Parameterization.** We consider an agent that may either perform *acquisition* or  
 136 *rearrangement* actions, parameterized below. Acquisition actions attempt to pick up a bite of  
 137 food, and rearrangement actions consolidate items. For example, as a plate of noodles becomes more  
 138 empty, the robot may need to employ a rearrangement action by pushing multiple strands together  
 139 before twirling (acquiring) for a satisfactory bite size.

140 In *acquisition*, a robot with a utensil-mounted end-  
 141 effector approaches the position  $(x_d, y_d, z_d)$  in the  
 142 workspace, and executes an acquisition motion  
 143 parameterized by roll  $\gamma$  and pitch  $\beta$  (i.e. twirling,  
 144 skewering, scooping, etc.). Here,  $(x_d, y_d, z_d)$  denotes  
 145 the *densest* location of the plate, where food is most  
 146 closely packed to encourage a high-volume bite.  
 147 Specifically,  $a_{t, \text{acq}} = (x_d, y_d, z_d, \gamma, \beta)$  (1).



148 The intent of *rearrangement* is to bring food items  
 149 from the sparsest plate region to the densest by  
 150 pushing from  $(x_f, y_f, z_f)$  to  $(x_d, y_d, z_d)$ , while  
 151 maintaining contact with the plate throughout. As  
 152 this is a planar push, we simply orient the tool orthogonal to the push direction, such that

Figure 2: **Action Parameterization:** We parameterize *acquisition* and *rearrangement* actions relative to the densest  $(x_d, y_d, z_d)$  and furthest  $(x_f, y_f, z_f)$  regions on the plate, as well as the utensil roll  $\gamma$  and pitch  $\beta$ .

153  $\gamma = \arctan\left(\frac{y_f - y_d}{x_f - x_d}\right)$ , and is untilted ( $\beta = 0^\circ$ ):  $a_{t, \text{rearrange}} = (x_d, y_d, z_d, x_f, y_f, z_f)$  (2).

## 154 4 VAPORS: Visual Action Planning Over Sequences

155 Within the visual state and action space outlined in Section 3.1, we present our approach VAPORS  
 156 for tackling long-horizon food acquisition. First, VAPORS learns a policy  $\pi_H$ , detailed in Section 4.1,  
 157 to select amongst high-level strategies for long-horizon plate clearance via model-based planning.  
 158 Finally, VAPORS learns a low-level policy  $\pi_L$ , which leverages visually-parameterized primitives to  
 159 carry out generated sequential plans for real-world food acquisition detailed in Section 4.2.

### 160 4.1 Learning High-Level Plans from Simulation

161 Our goal is to first learn a policy  $\pi_H$  for selecting amongst  $K$  discrete acquisition or rearrangement  
 162 strategies without concern for the low-level action parameters. To do so, we learn a latent dynamics  
 163 model of the plate from segmented image observations, and instantiate  $\pi_H(h^k | M_{\leq t}, a_{\leq t-1})$ ,  $k \in$   
 164  $\{1, \dots, K\}$  with model-based planning over this learned dynamics model. In this section,  $\tau$  denotes  
 165 the running counter of high-level primitives executed so far, and  $t$  denotes the current timestep.

166 **Simulator Overview.** We train  $\pi_H$  entirely in  
 167 simulation, where interactions can be collected  
 168 at scale as opposed to the real world where  
 169 manual plate resets and potential food waste are  
 170 prohibitively expensive. As current simulators  
 171 lack out-of-the-box support for many feeding  
 172 scenarios, we develop a custom simulated  
 173 food manipulation environment, visualized in  
 174 Figure 3 in Blender 2.92 [40], further detailed  
 175 in Appendix A.1. The simulator exposes RGB  
 176 images  $I_t$ , binary food segmentation masks  
 177  $M_t$ , and food item positional states  $s_t =$   
 178  $\{(x_i, y_i, z_i)\}_{i \in \{1, \dots, N\}}$ . Using this information, we design rewards for food acquisition in terms  
 179 of ground truth plate state and collect transitions to train  $\pi_H$ .

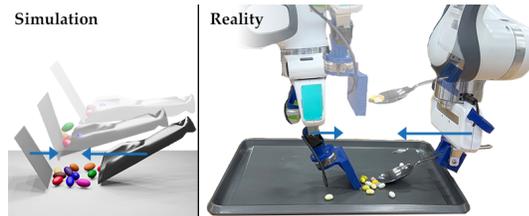


Figure 3: **Simulation vs. Real:** We visualize the task of bimanual scooping of jelly beans. Due to the sim-to-real gap, we merely leverage simulation to learn high-level food dynamics, and leave low-level action planning to real vision-parameterized primitives.

180 **Reward Design.** With access to a simulated testbed for feeding, we train  $\pi_H$  to select amongst  
 181 strategies via model-based reinforcement learning (RL). Our goal of efficient plate clearance can be  
 182 specified with a reward that incentivizes either (1) successfully picking up food, or (2) reducing the  
 183 spread of items on a plate. Optimizing for the first objective alone might lead to plate clearance, but  
 184 at a slow pace of taking low-volume bites. The second objective encourages rearrangement when the  
 185 plate is sparse to aid downstream acquisition. Concretely, we express this as a weighted reward with  
 186 tunable weight  $\alpha \in [0, 1]$ :  $r_t = \alpha(\text{PICKUP GAIN}) + (1 - \alpha)(\text{COVERAGE LOSS})$  (3). Here, PICKUP  
 187 measures the quantity of food items picked up. COVERAGE measures the spread of items on the plate,  
 188 illustrated in blue in Fig. 2). We provide the details for computing both in Appendix A.2.

189 **Learning Latent Plate Dynamics.** With  
 190 a means of measuring task progress via  $r_t$   
 191 and access to plate observations  $M_t$ , we  
 192 propose a model-based agent that learns  
 193 plate dynamics from segmented observations  
 194 and uses the learned model to plan actions  
 195 that maximize reward. We achieve this  
 196 by training a multi-headed latent dynamics  
 197 model with the following (Fig. 4): (1)  
 198 An *encoder*  $q(z_t|M_{\leq t}, a_{\leq t-1})$  compressing  
 199 high-dimensional segmented images  $M_t$  to  
 200 compressed latent states  $z_t$ , (2) A *transition*  
 201 *function* over the latent states  $p(z_\tau|z_{\tau-1}, h_{\tau-1}^k)$   
 202 with which to imagine rollouts, and (3) A decoded  
 203 *reward model* given by  $p(r_t|z_t)$ , such that at test time, we can sample action sequences and determine  
 204 which maximize predicted rewards. We note that the transition function learns to predict high-level  
 205 plate state changes between  $\tau - 1$  and  $\tau$  as a result of executing a primitive  $h_\tau^k$ , rather than between  
 individual timesteps  $t - 1$  and  $t$  due to  $a_t$ .

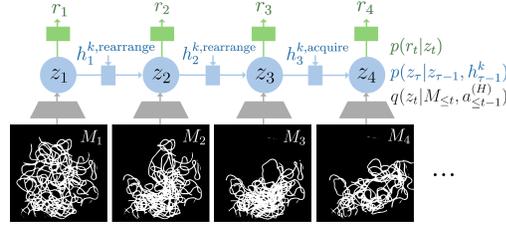


Figure 4: **Latent Plate Dynamics Model:** We learn a latent dynamics model of the plate comprised of an encoder  $q$ , transition model  $p(z_\tau|z_{\tau-1}, h_{\tau-1}^k)$ , and a reward model  $p(r_t|z_t)$ . We use this model to select action sequences that maximize future rewards.

206 During training, we collect simulated transitions consisting of the masked image, high-level primitive,  
 207 low-level action, and reward  $\{(M_t, h_\tau^k, a_t, r_r)\}$ . We train each head of this network using the  
 208 objectives detailed in Appendix B.1.

209 We note that this approach is highly related to [37] with several crucial design choices. First, we  
 210 learn plate dynamics over *segmented* image observations  $M_t$  of food items on a plate, as opposed to  
 211 raw RGB observations. This allows the dynamics model to attend to food items rather than the  
 212 whole plate, provides an easily transferable representation between simulation and reality, and eases  
 213 pressure for latent representations to capture irrelevant details in pixel space. Additionally, we learn a  
 214 policy within an action space of discrete but continuously parameterized primitives as opposed to a  
 215 high-dimensional space like joint-motor commands. This encourages actions that induce meaningful  
 216 and perceptible plate changes likely to be encountered in downstream feeding.

217 **Model-Based Planning.** Once trained, we leverage the learned encoder, transition model, and  
 218 reward model towards instantiating  $\pi_H$  as an MPC-style planner with a receding  $T$ -step horizon.  
 219 At timestep  $t$ , we enumerate all  $K^T$  future candidate action sequences for the small library of  
 220 primitives  $K$ . Conditioned on a history of observations  $M_{1:t}$  and actions  $a_{1:t-1}$ , we imagine the  
 221 future latent states  $z_{\tau:\tau+T+1}$  under each action sequence  $h_{\tau:\tau+T}^k$  via the transition function. Next,  
 222 we predict decoded rewards according to the reward model  $p(r_t|z_t)$  for each candidate sequence:  
 223  $R = \sum_{i=\tau+1}^{\tau+T+1} \mathbb{E}[p(r_i|z_i)]$ . Given the sequence of actions  $(\hat{h}_\tau^k, \hat{h}_{\tau+1}^k, \dots, \hat{h}_T^k)$  which maximizes  
 224 predicted cumulative reward  $R$ , we take  $\pi_H(M_{\leq t}, a_{\leq t-1}) = \hat{h}_\tau^k$ , the first primitive in the predicted  
 225 sequence. After executing this action, we replan with  $\pi_H$ , terminating when  $\tau = T$ . Details of the  
 226 full planning pipeline, adapted from [37], are provided in Appendix B.2.

## 227 4.2 Visual Policies for Low-Level Real Manipulation

228 Our learned simulated task dynamics model from Section 4.1 relies on segmented images  $M_t$  as an  
 229 observation space and parameterized primitives as an action space. In this section, we describe the  
 230 visual state estimation pipelines we use to instantiate our state-action representations on real data.

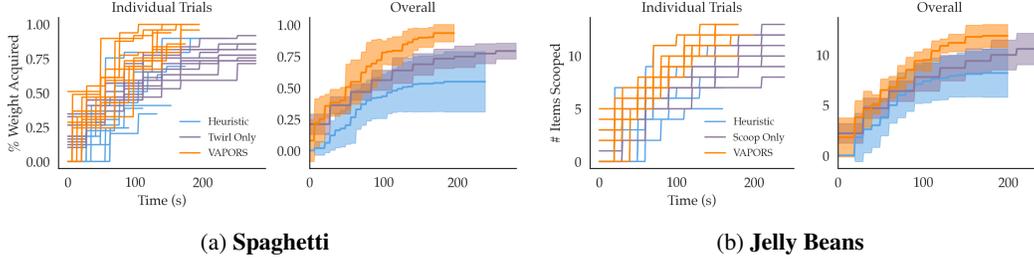


Figure 5: Across 10 trials for spaghetti (a) and jelly bean (b) acquisition, we visualize the cumulative amount acquired across individual trials (left) and averaged overall (right). Shading denotes the standard error.

231 **Food Segmentation.** To define acquisition and rearrangement actions relative to the poses of food,  
 232 we learn to segment food items on a plate as shown in Fig. 1. We learn a binary segmentation  
 233 model  $f_{\text{seg}} : \mathbb{R}_+^{W \times H \times 3} \rightarrow \mathbb{R}_+^{W \times H}$ , where for a real image  $I_t \in \mathbb{R}_+^{W \times H \times 3}$ ,  $f_{\text{seg}}(I_t)$  yields a binary  
 234 segmentation mask  $\hat{M}_t$  which serves as input to  $\pi_H$ . To train  $f_{\text{seg}}$ , we require a paired dataset of  
 235 real plate images and ground truth segmentation masks. However, manually labeling pixel-level  
 236 segmentation annotations on images is a painstaking and time-consuming process for real plates  
 237 of food. Instead, we use a self-supervised annotation process which starts by taking an image of  
 238 an empty plate, gradually adding food items to the plate, and using the absolute frame difference  
 239 between the empty plate image and current observation to obtain the food segmentation mask. We  
 240 implement  $f_{\text{seg}}$  as a fully convolutional FPN (Feature Pyramid Network) and train it according to the  
 241 procedure detailed in Appendix B.3.

242 **Food Orientation.** Although segmentation provides a means to sense global *positional* information  
 243 about food on the plate, we also care about precisely *orienting* a utensil with respect to the local  
 244 geometry of a food item. For instance, using a fork to pick up a group of noodles requires orienting  
 245 the fork tines opposite the grain of the strands. This is crucial to preventing slippage during twirling  
 246 (Fig. 11), which tends to occur when the tines and strands run parallel. To address this, we also  
 247 learn a network  $f_{\text{ori}} : \mathbb{R}_+^{W' \times H' \times 3} \rightarrow \mathbb{R}$  mapping a local RGB crop of a food item of dimensions  
 248  $W' \times H'$  to the desired roll orientation of the utensil  $\gamma$ . Prior work has shown that acquiring a food  
 249 item orthogonal to its main principal axis, such as skewering a carrot against its length-wise axis  
 250 rather than width-wise, can improve acquisition stability [4, 6]. Thus, we implement  $f_{\text{ori}}$  as a fully  
 251 convolutional network with a ResNet backbone and train it from a small amount of real food item  
 252 crops (200), manually annotated with keypoints defining the principal food item axis as in [4].

253 **Action Instantiation.** With the visual state estimation pipelines  $f_{\text{seg}}$  and  $f_{\text{ori}}$  trained offline, we can  
 254 instantiate  $\pi_L(a_t | o_t, h^k)$  for real-world manipulation. Given an RGBD image observation  $I_t, D_t$ , we  
 255 first infer the segmentation mask  $\hat{M}_t = f_{\text{seg}}(I_t)$ . Next, we query  $\pi_H$  to obtain a selected primitive  
 256  $\hat{h}^k = \pi_H(\hat{M}_{\leq t}, \hat{a}_{\leq t-1}^{(H)})$ .

257 If  $\hat{h}^k$  is an acquisition primitive, we instantiate the continuous action  $a_{t,\text{acquir}}$  according to Eq. (1)  
 258 by estimating the densest plate region  $(\hat{x}_d, \hat{y}_d, \hat{z}_d)$  and utensil orientation  $\hat{\gamma}$ . To do so, we apply a  
 259 standard 2D Gaussian kernel over  $\hat{M}_t$  yielding  $\hat{M}'_t$ . This blurs the image such that high-density  
 260 regions in the original segmentation mask remain saturated but sparse regions have lower intensity.  
 261 From this, we take the 2D argmax  $\hat{u}_d, \hat{v}_d = \arg \max_{(u,v) \in \hat{M}'_t} \hat{M}'_t[u, v]$  to be the densest pixel in  
 262 the image, deprojected to a 3D location  $(\hat{x}_d, \hat{y}_d, \hat{z}_d)$  via  $D_t$  and known camera intrinsics. Given a  
 263 food item crop centered at the densest pixel,  $I_t$  (Fig. 2) we also infer the utensil orientation with  
 264  $\hat{\gamma} = f_{\text{ori}}(I_t)$ . For a rearrangement primitive, we parameterize  $a_{t,\text{rearrange}}$  according to Eq. (2). In  
 265 addition to sensing the densest plate region, we sense the furthest region  $(\hat{x}_f, \hat{y}_f, \hat{z}_f)$  by finding the  
 266 lowest intensity pixel in  $\hat{M}'_t$ . This yields the following instantiations:

$$\pi_L(o_t, h^{k,\text{acquir}}) = (\hat{x}_d, \hat{y}_d, \hat{z}_d, \hat{\gamma}) \quad \Bigg| \quad \pi_L(o_t, h^{k,\text{rearrange}}) = (\hat{x}_d, \hat{y}_d, \hat{z}_d, \hat{x}_f, \hat{y}_f, \hat{z}_f)$$

267 Finally, VAPORS operates in a perception-action loop using  $\pi_H$  to generate sequential plans and  $\pi_L$   
 268 to execute them. The full algorithm can be found in Algorithm 1 of the Appendix.  
 269

## 270 5 Experiments

271 We seek to evaluate **VAPORS** ability to clear plates, by effectively leveraging diverse strategies  
 272 and planning over long horizons. Thus, we compare against a single-strategy baseline with no  
 273 long-horizon reasoning and a multi-strategy approach that plans long-term actions heuristically rather  
 274 than via learned plate dynamics. We consider two challenging real-world feeding scenarios to test the  
 275 capabilities of **VAPORS** compared to other approaches: noodle acquisition and bimanual scooping.

276 **Experimental Setup:** In noodle acquisition (Fig. 9), a Franka robot with a wrist-mounted  
 277 custom motorized fork and RGBD camera must decide amongst *twirling* (acquisition) or *grouping*  
 278 (rearrangement) to clear a plate of noodles. In bimanual scooping (Fig. 10), two Franka robots  
 279 operating from overhead RGBD cameras must select amongst *scooping* (acquisition) or *grouping*  
 280 (rearrangement) to clear a plate of jelly beans. For both tasks, we consider a *half-full* initial plate  
 281 distribution ( 50 g. noodles, 15 jelly beans) and a hard count of  $\tau = 10$  actions for spaghetti and  
 282  $\tau = 8$  actions for jelly beans, encouraging the acquisition of multiple items at once to finish a  
 283 plate. For both tasks, we assume access to hand-eye calibration between the RGB-D camera and  
 284 robot end-effector. In Appendix C we outline the hardware setup and control stack, low-level action  
 285 instantiations, and training details for each task.

286 **Baselines:** *Acquire-Only* is identical to **VAPORS** in terms of  $\pi_L$ , but does not perform any long-  
 287 horizon reasoning. Instead, at each timestep, this approach only acquires via twirling or scooping,  
 288 with no rearrangement in between. *Heuristic* also utilizes  $\pi_L$  in the same manner, but replaces  $\pi_H$   
 289 with a naïve group-then-acquire strategy. This method senses the COVERAGE, as defined in Eq. (3),  
 290 over  $\hat{M}_t$  to heuristically determine when acquiring or rearrangement is appropriate. When the area  
 291 exceeds a pre-defined threshold, the policy defaults to rearrangement and otherwise acquires.

292 **Plate Clearance Results:** We evaluate **VAPORS**,  
 293 *Acquire-Only*, and *Heuristic* on clearing plates across  
 294 10 trials for each task (Fig. 5). We see that  
 295 **VAPORS** achieves the most efficient and highest  
 296 cumulative plate clearance. As expected, *Acquire-Only*  
 297 optimizes only for acquisition in the current  
 298 instant, without exploiting the benefits of grouping  
 299 for a more substantial pickup of multiple items at  
 300 once. Scooping one jelly bean at a time or attempting  
 301 to twirl just a few strands of noodles repeatedly  
 302 leads to the observed slow rate of overall clearance.  
 303 *Heuristic*'s greedy group-then-acquire approach plans  
 304 based on detected coverage thresholds, which we  
 305 find is brittle in practice especially for any artifacts  
 306 in segmentation mask predictions. This naive metric  
 307 also does not encourage acquiring any bite-sized piles  
 308 that may form intermittently, but rather aims to amass  
 309 everything into one large pile before acquiring. This  
 310 delays acquisition gains and wastes the action budget.

311 **User Evaluation:** Additionally, we conducted a user  
 312 study with 49 non-disabled participants (age range  
 313  $27.0 \pm 9.5$ , 46.9% female and 53.1% male) to gauge user preferences across methods. Of this pool,  
 314 77.6% reported prior experience interacting with robots before, 75.5% reported having fed someone  
 315 before, and 28.6% reported having been fed as an adult. We hypothesized: **H1.** *Compared to*  
 316 *baselines, VAPORS use of multiple strategies and long horizon foresight will lead to more preferable*  
 317 *feeding in terms of quantitative and qualitative metrics.*

318 We used a within-subjects design where we presented each participant with videos of all 10  
 319 plate clearance trials per each of the three methods, for either noodle acquisition or bimanual  
 320 scooping. For each participant, we randomized the method order, the order of trials per

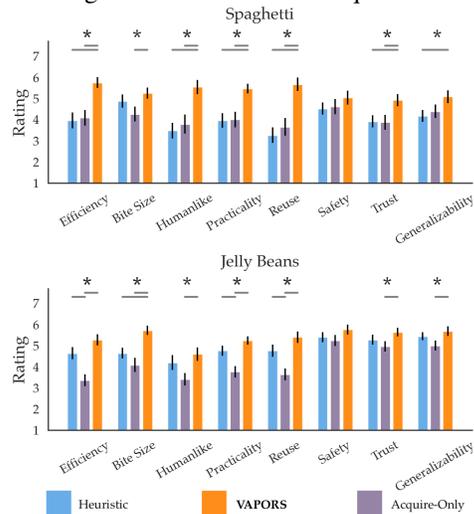


Figure 6: **Likert Ratings:** We administer a 7-point Likert survey to users after observing 10 trials per method. **VAPORS** elicits the most positive feedback across all criteria. **\*\*** indicates statistical significance ( $p < 0.05$ ).

321 method, and the food group. In the study, we ask participants to rate efficiency, bite  
 322 size, similarity to human feeding, practicality, likelihood for reuse, safety, and generalization.  
 323

324 After watching all trials, we provided users with a  
 325 7-point Likert survey to assess these criteria (Fig. 6).  
 326 VAPORS incurs the highest qualitative user ratings  
 327 across criteria, compared to the Acquire-Only and  
 328 Heuristic baselines, and with statistical significance  
 329 for certain categories ( $p < 0.05$ , denoted “\*”). Users  
 330 noted that VAPORS “mimicked natural feeding,” and  
 331 “showed a capacity for clustering as the plate got more  
 332 and more empty, which felt like a great and efficient  
 333 approach,” while Heuristic and Acquire-Only “seem  
 334 like extreme policies, where [Acquire-Only] never  
 335 tries to cluster and [Heuristic] focus too much on making big piles.” These results align with the  
 336 hypothesis that VAPORS’ use of multiple strategies and ability to reason over long horizons benefits  
 337 a user’s mealtime experience. We provide additional user study findings in Appendix C.

338 **Generalization Testing:** Finally, we stress-test VAPORS’ generalization capabilities by  
 339 experimenting with noodle dishes prepared with sauces and garnishes as well as ordered from  
 340 DoorDash (Fig. 7). We conduct 18 additional trials of plate clearance on unseen plates, separated  
 341 into three tiers of difficulty with 6 trials per tier. We summarize our findings in Table 1.

342 VAPORS achieves near full plate  
 343 clearance for Tier 1 noodles,  
 344 demonstrating generalization to  
 345 noodle shapes and sizes (Table 1).  
 346 While VAPORS is still able to  
 347 make decent progress towards  
 348 plate clearance in Tier 2, we observe the occurrence of more slip failures (D) and misplanned actions  
 349 (A, B) due to the addition of sauce and distractor food items. Somewhat surprisingly, the performance  
 350 gap between Tier 2 and Tier 3 is minimal, with VAPORS being able to clear well over half the  
 351 noodles for a fully out-of-distribution plate. The main challenges include misperceiving cabbage for  
 352 noodles in the chow mein, as well as dropping twirled noodles heavily coated in pesto or soy sauce  
 353 (D) (Fig. 11). Regardless, VAPORS demonstrates promising signs of zero-shot generalization.

## 354 6 Discussion

355 We present VAPORS, which to our knowledge is the first framework to address the multi-step food  
 356 acquisition problem in robot-assisted feeding. Our hybrid approach leverages simulation to learn  
 357 to model high-level plate dynamics at scale, and uses visual pose estimation in order to perform  
 358 dexterous maneuvers for complex low-level food pickup. We experimentally validate VAPORS on a  
 359 complex suite of real-world food acquisition tasks such as noodle acquisition and bimanual scooping  
 360 of beans. VAPORS demonstrates the ability to clear plates efficiently over non-learned baselines  
 361 while appealing to the feeding preferences of real users.

362 **Limitations and Future Work.** Although this work highlights promising initial results toward  
 363 generalization across food variations such as shape, sauces, and toppings, we acknowledge that our  
 364 library of low-level primitives is currently limited and would benefit from a more diverse repertoire.  
 365 One actionable future direction is expanding our library with prior work on skewering, cutting, and  
 366 even toppling unstable items to tackle a more expansive set of plates. Our initial prototypes for  
 367 dexterous food acquisition, such as the motorized fork, also open up interesting possibilities for future  
 368 designs of dexterous interchangeable utensils which would enable rapid strategy switching. Currently,  
 369 the system also executes primitives in an open-loop fashion, but we hope to use reactive control in  
 370 the future to adapt online to slippage or imprecision.

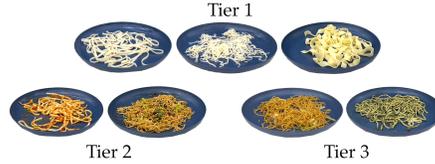


Figure 7: **Noodle Acquisition Tiers of Difficulty:** Tier 1 consists of plain noodle varieties: Dan Dan, Udon, and Pappardelle noodles. Tier 2 includes Tier 1 plates along with soy sauce, marinara sauce, and garnishes such as parsley or cilantro. Tier 3 plates include noodle dishes such as pesto pasta and chow mein ordered from DoorDash.

Tier	Description	% Cleared	Failure Categorization				Failure Rate
			A	B	C	D	
1	Plain Noodles	90% ± 6%	2	7	2	0	18%
2	Noodles w/ Sauce	68% ± 16%	2	8	1	4	25%
3	DoorDash Noodles	64% ± 13%	3	5	2	4	23%

Table 1: **OOD Results and Categorization of Failure Modes:** (A) Misperception, (B) Wrong Action, (C) Imprecision, (D) Slip.

## References

- 371
- 372 [1] M. W. Brault et al. *Americans with disabilities: 2010*. US Department of Commerce, Economics and  
373 Statistics Administration, US . . . , 2012.
- 374 [2] S. W. Brose, D. J. Weber, B. A. Salatin, G. G. Grindle, H. Wang, J. J. Vazquez, and R. A. Cooper. The  
375 role of assistive robotics in the lives of persons with disability. *American Journal of Physical Medicine &*  
376 *Rehabilitation*, 89(6):509–521, 2010.
- 377 [3] V. Maheu, P. S. Archambault, J. Frappier, and F. Routhier. Evaluation of the jaco robotic arm: Clinico-  
378 economic study for powered wheelchair users with upper-extremity disabilities. In *2011 IEEE International*  
379 *Conference on Rehabilitation Robotics*, pages 1–5. IEEE, 2011.
- 380 [4] P. Sundaresan, S. Belkhale, and D. Sadigh. Learning visuo-haptic skewering strategies for robot-assisted  
381 feeding. In *6th Annual Conference on Robot Learning*, 2022. URL [https://openreview.net/forum?](https://openreview.net/forum?id=1Lq09gVoaTE)  
382 [id=1Lq09gVoaTE](https://openreview.net/forum?id=1Lq09gVoaTE).
- 383 [5] E. K. Gordon, X. Meng, T. Bhattacharjee, M. Barnes, and S. S. Srinivasa. Adaptive robot-assisted feeding:  
384 An online learning framework for acquiring previously unseen food items. In *2020 IEEE/RSJ International*  
385 *Conference on Intelligent Robots and Systems (IROS)*, pages 9659–9666. IEEE, 2020.
- 386 [6] R. Feng, Y. Kim, G. Lee, E. K. Gordon, M. Schmittle, S. Kumar, T. Bhattacharjee, and S. S. Srinivasa.  
387 Robot-assisted feeding: Generalizing skewering strategies across food items on a plate. In *The International*  
388 *Symposium of Robotics Research*, pages 427–442. Springer, 2019.
- 389 [7] J. Grannen, Y. Wu, S. Belkhale, and D. Sadigh. Learning bimanual scooping policies for food acquisition.  
390 In *6th Annual Conference on Robot Learning*, 2022. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=qDtbMK67PJG)  
391 [qDtbMK67PJG](https://openreview.net/forum?id=qDtbMK67PJG).
- 392 [8] L. Shaikewitz, Y. Wu, S. Belkhale, J. Grannen, P. Sundaresan, and D. Sadigh. In-mouth robotic bite transfer  
393 with visual and haptic sensing. *arXiv preprint arXiv:2211.12705*, 2022.
- 394 [9] S. Belkhale, E. K. Gordon, Y. Chen, S. Srinivasa, T. Bhattacharjee, and D. Sadigh. Balancing efficiency  
395 and comfort in robot-assisted bite transfer. In *2022 International Conference on Robotics and Automation*  
396 *(ICRA)*, pages 4757–4763. IEEE, 2022.
- 397 [10] D. Gallenberger, T. Bhattacharjee, Y. Kim, and S. S. Srinivasa. Transfer depends on acquisition: Analyzing  
398 manipulation strategies for robotic feeding. In *2019 14th ACM/IEEE International Conference on Human-*  
399 *Robot Interaction (HRI)*, pages 267–276. IEEE, 2019.
- 400 [11] M. Sharma, K. Zhang, and O. Kroemer. Learning semantic embedding spaces for slicing vegetables. *arXiv*  
401 *preprint arXiv:1904.00303*, 2019.
- 402 [12] K. Zhang, M. Sharma, M. Veloso, and O. Kroemer. Leveraging multimodal haptic sensory data for robust  
403 cutting. In *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*, pages  
404 409–416. IEEE, 2019.
- 405 [13] E. Heiden, M. Macklin, Y. Narang, D. Fox, A. Garg, and F. Ramos. Disect: A differentiable simulator for  
406 parameter inference and control in robotic cutting. *arXiv preprint arXiv:2203.10263*, 2022.
- 407 [14] H. Suh and R. Tedrake. The surprising effectiveness of linear models for visual foresight in object pile  
408 manipulation. In *International Workshop on the Algorithmic Foundations of Robotics*, pages 347–363.  
409 Springer, 2020.
- 410 [15] L. X. Shi, J. J. Lim, and Y. Lee. Skill-based model-based reinforcement learning. *arXiv preprint*  
411 *arXiv:2207.07560*, 2022.
- 412 [16] X. Lin, C. Qi, Y. Zhang, Z. Huang, K. Fragkiadaki, Y. Li, C. Gan, and D. Held. Planning with spatial-  
413 temporal abstraction from point clouds for deformable object manipulation. In *6th Annual Conference on*  
414 *Robot Learning*, 2022. URL <https://openreview.net/forum?id=tyxyBj2w4vw>.
- 415 [17] Z. Cao, E. Bıyık, W. Z. Wang, A. Raventos, A. Gaidon, G. Rosman, and D. Sadigh. Reinforcement learning  
416 based control of imitative policies for near-accident driving. *arXiv preprint arXiv:2007.00178*, 2020.
- 417 [18] M. Dalal, D. Pathak, and R. R. Salakhutdinov. Accelerating robotic reinforcement learning via  
418 parameterized action primitives. *Advances in Neural Information Processing Systems*, 34:21847–21859,  
419 2021.

- 420 [19] S. Nasiriany, H. Liu, and Y. Zhu. Augmenting reinforcement learning with behavior primitives for  
421 diverse manipulation tasks. In *2022 International Conference on Robotics and Automation (ICRA)*, pages  
422 7477–7484. IEEE, 2022.
- 423 [20] C. Agia, T. Migimatsu, J. Wu, and J. Bohg. Taps: Task-agnostic policy sequencing. *arXiv preprint*  
424 *arXiv:2210.12250*, 2022.
- 425 [21] A. Kumar, Z. Fu, D. Pathak, and J. Malik. Rma: Rapid motor adaptation for legged robots. *arXiv preprint*  
426 *arXiv:2107.04034*, 2021.
- 427 [22] S. Bahl, A. Gupta, and D. Pathak. Hierarchical neural dynamic policies. *arXiv preprint arXiv:2107.05627*,  
428 2021.
- 429 [23] L. Shao, T. Migimatsu, Q. Zhang, K. Yang, and J. Bohg. Concept2robot: Learning manipulation concepts  
430 from instructions and human demonstrations. *The International Journal of Robotics Research*, 40(12-14):  
431 1419–1434, 2021.
- 432 [24] Meet Obi. <https://meetobi.com/>. [Online; accessed 6-June-2022].
- 433 [25] Meal-Mate - Made2Aid. [https://www.made2aid.co.uk/productprofile?productId=8&](https://www.made2aid.co.uk/productprofile?productId=8&company=RBF%20Healthcare&product=Meal-Mate)  
434 [company=RBF%20Healthcare&product=Meal-Mate](https://www.made2aid.co.uk/productprofile?productId=8&company=RBF%20Healthcare&product=Meal-Mate). [Online; accessed 6-June-2022].
- 435 [26] E. K. Gordon, S. Roychowdhury, T. Bhattacharjee, K. Jamieson, and S. S. Srinivasa. Leveraging post hoc  
436 context for faster learning in bandit settings with applications in robot-assisted feeding. In *2021 IEEE*  
437 *International Conference on Robotics and Automation (ICRA)*, pages 10528–10535. IEEE, 2021.
- 438 [27] L. P. Kaelbling and T. Lozano-Pérez. Hierarchical planning in the now. In *Workshops at the Twenty-Fourth*  
439 *AAAI Conference on Artificial Intelligence*, 2010.
- 440 [28] S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, and P. Abbeel. Combined task and motion planning  
441 through an extensible planner-independent interface layer. In *2014 IEEE international conference on*  
442 *robotics and automation (ICRA)*, pages 639–646. IEEE, 2014.
- 443 [29] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez. Integrated  
444 task and motion planning. *Annual review of control, robotics, and autonomous systems*, 4:265–293, 2021.
- 445 [30] R. Chitnis, D. Hadfield-Menell, A. Gupta, S. Srivastava, E. Groshev, C. Lin, and P. Abbeel. Guided search  
446 for task and motion plans using learned heuristics. In *2016 IEEE International Conference on Robotics*  
447 *and Automation (ICRA)*, pages 447–454. IEEE, 2016.
- 448 [31] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling. Ffrob: An efficient heuristic for task and motion  
449 planning. In *Algorithmic Foundations of Robotics XI*, pages 179–195. Springer, 2015.
- 450 [32] H. Shi, H. Xu, Z. Huang, Y. Li, and J. Wu. Robocraft: Learning to see, simulate, and shape elasto-plastic  
451 objects with graph networks. *arXiv preprint arXiv:2205.02909*, 2022.
- 452 [33] X. Lin, Z. Huang, Y. Li, J. B. Tenenbaum, D. Held, and C. Gan. Diffskill: Skill abstraction from  
453 differentiable physics for deformable object manipulations with tools. *arXiv preprint arXiv:2203.17275*,  
454 2022.
- 455 [34] C. Finn and S. Levine. Deep visual foresight for planning robot motion. In *2017 IEEE International*  
456 *Conference on Robotics and Automation (ICRA)*, pages 2786–2793. IEEE, 2017.
- 457 [35] F. Ebert, C. Finn, S. Dasari, A. Xie, A. Lee, and S. Levine. Visual foresight: Model-based deep  
458 reinforcement learning for vision-based robotic control. *arXiv preprint arXiv:1812.00568*, 2018.
- 459 [36] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi. Dream to control: Learning behaviors by latent imagination.  
460 *arXiv preprint arXiv:1912.01603*, 2019.
- 461 [37] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning latent dynamics  
462 for planning from pixels. In *International conference on machine learning*, pages 2555–2565. PMLR,  
463 2019.
- 464 [38] Y. Li, A. Torralba, A. Anandkumar, D. Fox, and A. Garg. Causal discovery in physical systems from  
465 videos. *Advances in Neural Information Processing Systems*, 33:9180–9192, 2020.
- 466 [39] S. Pateria, B. Subagdja, A.-h. Tan, and C. Quek. Hierarchical reinforcement learning: A comprehensive  
467 survey. *ACM Computing Surveys (CSUR)*, 54(5):1–35, 2021.

- 468 [40] T. Roosendaal. Blender foundation. *The essential Blender: guide to 3D creation with the open source*  
469 *suite Blender*, 2007.
- 470 [41] E. Coumans and Y. Bai. 2019. pybullet, a python module for physics simulation for games, robotics and  
471 machine learning, 2016.
- 472 [42] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire,  
473 A. Handa, et al. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv*  
474 *preprint arXiv:2108.10470*, 2021.
- 475 [43] P. R. Florence, L. Manuelli, and R. Tedrake. Dense object nets: Learning dense visual object descriptors  
476 by and for robotic manipulation. *arXiv preprint arXiv:1806.08756*, 2018.
- 477 [44] V. Viswanath, K. Shivakumar, J. Kerr, B. Thananjeyan, E. Novoseller, J. Ichnowski, A. Escontrela,  
478 M. Laskey, J. E. Gonzalez, and K. Goldberg. Autonomously untangling long cables. *arXiv preprint*  
479 *arXiv:2207.07813*, 2022.
- 480 [45] P. Sundaresan, J. Grannen, B. Thananjeyan, A. Balakrishna, J. Ichnowski, E. Novoseller, M. Hwang,  
481 M. Laskey, J. E. Gonzalez, and K. Goldberg. Untangling dense non-planar knots by learning manipulation  
482 features and recovery policies. *arXiv preprint arXiv:2107.08942*, 2021.
- 483 [46] J. Grannen, P. Sundaresan, B. Thananjeyan, J. Ichnowski, A. Balakrishna, M. Hwang, V. Viswanath,  
484 M. Laskey, J. E. Gonzalez, and K. Goldberg. Untangling dense knots by learning task-relevant keypoints.  
485 *arXiv preprint arXiv:2011.04999*, 2020.
- 486 [47] L. Manuelli, Y. Li, P. Florence, and R. Tedrake. Keypoints into the future: Self-supervised correspondence  
487 in model-based reinforcement learning. *arXiv preprint arXiv:2009.05085*, 2020.
- 488 [48] W. Gao and R. Tedrake. kpm 2.0: Feedback control for category-level robotic manipulation. *IEEE*  
489 *Robotics and Automation Letters*, 6(2):2962–2969, 2021.
- 490 [49] M. Shridhar, L. Manuelli, and D. Fox. Cliport: What and where pathways for robotic manipulation. In  
491 *Conference on Robot Learning*, pages 894–906. PMLR, 2022.
- 492 [50] A. Zeng, P. Florence, J. Tompson, S. Welker, J. Chien, M. Attarian, T. Armstrong, I. Krasin, D. Duong,  
493 V. Sindhwani, et al. Transporter networks: Rearranging the visual world for robotic manipulation. In  
494 *Conference on Robot Learning*, pages 726–747. PMLR, 2021.
- 495 [51] D. Seita, A. Ganapathi, R. Hoque, M. Hwang, E. Cen, A. K. Tanwani, A. Balakrishna, B. Thananjeyan,  
496 J. Ichnowski, N. Jamali, et al. Deep imitation learning of sequential fabric smoothing from an algorithmic  
497 supervisor. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages  
498 9651–9658. IEEE, 2020.
- 499 [52] H. Ha and S. Song. Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding.  
500 In *Conference on Robot Learning*, pages 24–33. PMLR, 2022.

# Learning Sequential Acquisition Policies for Robot-Assisted Feeding

Please refer to our [website](#) for videos, code, and supplementary material, as well as the 'Additional Experiments' page for supplementary ablations and a comparison to additional baselines. In this section, we provide an overview of the main design choices behind VAPORS and a thorough overview at implementation and experimental details.

## A Simulator Details

### A.1 Simulator Design

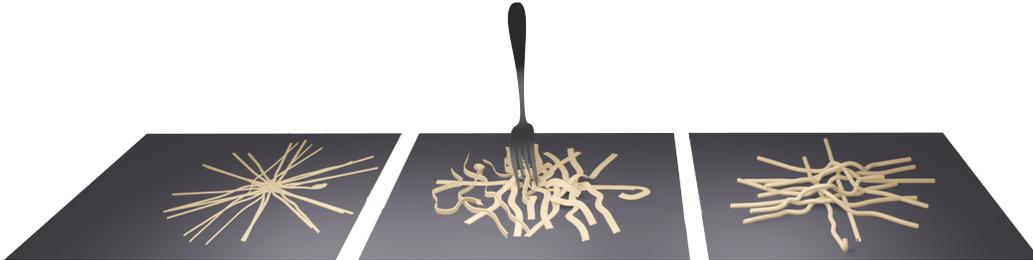


Figure 8: **Blender Food Simulation Environment:** We implement a custom food manipulation simulator in Blender 2.92 with an Open AI gym-style environment. The simulator supports softbody objects, such as noodles in different shape variations, as well as rigid, granular piles of items. We implement cutlery with arbitrary utensil meshes such as forks and spoons, and implement actions using the *keyframing* feature of Blender to control the position and orientation of a tool across frames.

We use Blender 2.92, a physics and rendering engine, to develop a custom feeding environment supporting deformable items, rigid items, and cutlery interactions. To instantiate deformable items like noodles (Fig. 8), we represent each item as a group of particles simulated with soft body physics. We treat granular piles of food such as jelly beans as separate rigid bodies. Additionally, we provide support for mesh-based utensils including a fork, spoon, and pusher tool, where we programatically keyframe the position and orientation of the tool across simulation frames to implement actions.

### A.2 Reward Design

In this section, we describe the implementation of the reward function given in Eq. (3). For a set of known food item states in simulation  $s_t = \{(x_i, y_i, z_i)\}_{i \in (1, \dots, N)}$ , PICKUP measures the quantity of food items picked up out of  $N$  total items. We detect a picked up food item in simulation by thresholding the  $z$  position of all items before and after an action, relative to plate height. Analogous to task progress metrics in cloth smoothing work [51, 52], we use COVERAGE to measure of spread of items on the plate. We compute this via the area of the convex hull of  $\{(x_i, y_i)\}_{i \in (1, \dots, N)}$ , depicted in Fig. 2, via the Scipy Python library.

## B Details of Learning-Based Methods

### B.1 Latent Dynamics Training Details

We implement the latent plate dynamics model using the recurrent state space model from [37], with  $64 \times 64$  input images and 30-dimensional diagonal Gaussian latent variables. This is a multi-headed deep recurrent network comprised of a learned encoder, transition model, and reward model. We supervise each head of the network with the following objectives:

- For the *encoder*  $q(z_t | M_{\leq t}, a_{\leq t-1})$ , we use an auxiliary decoder head that upsamples latent variables  $z_t$  to predicted images  $\hat{M}_t$  and take the mean-squared error between  $(\hat{M}_t, M_t)$  as a standard reconstruction objective. This encourages the learned latent representations to preserve the notion of food spread captured in segmented image observations.
- We supervise the *transition function*  $p(z_\tau | z_{\tau-1}, h_{\tau-1}^k)$  head using the KL-divergence for multi-step predictions as defined in [37].

535 • Finally, for the *reward model* given by  $p(r_t|z_t)$ , we take the mean-squared error between  
 536 predicted rewards and ground truth rewards  $(\hat{r}_t, r_t)$ . This objective promotes accurately  
 537 decoding rewards of future states to inform planning at test-time.

## 538 B.2 Planning with Learned Dynamics Model

539 Once trained, we use an MPC-style loop to sample and plan actions that maximize predicted rewards  
 540 under the learned reward model.

541 At time  $\tau$ , we can enumerate all  $K^T$  future candidate action sequences for the small library of  
 542 primitives  $K$ , where  $T$  is the planning horizon. Conditioned on a history of observations  $M_{1:t}$  and  
 543 actions  $a_{1:t-1}$ , we imagine the future latent states  $z_{\tau:\tau+T+1}$  under each action sequence  $h_{\tau:\tau+T}^k$ :

$$z_{t:t+T+1} \sim q(z_\tau | M_{1:t}, a_{1:t-1}) \prod_{i=\tau+1}^{\tau+T+1} p(z_i | z_{i-1}, h_{i-1}^k), \quad (4)$$

544 where  $q(z_t | M_{\leq t}, a_{< t-1})$  is the learned encoder and  $p(z_\tau | z_{\tau-1}, h_{\tau-1}^k)$  is the learned transition model.  
 545 Next, we predict decoded rewards according to the reward model  $p(r_t | z_t)$  for each candidate sequence:

$$R = \sum_{i=\tau+1}^{i+T+1} \mathbb{E} [p(r_i | z_i)]. \quad (5)$$

546 Next, we select the sequence of actions  $(\hat{h}_\tau^k, \hat{h}_{\tau+1}^k, \dots, \hat{h}_T^k)$  which maximizes predicted cumulative  
 547 reward  $R$ . The final step of the MPC planning loop is we take  $\pi_H(M_{\leq t}, a_{\leq t-1}) = \hat{h}_\tau^k$ , which is  
 548 simply the first primitive in the predicted sequence. After executing this action, we replan with  $\pi_H$ ,  
 549 thus obtaining a second action and so on until  $\tau = T$  (Algorithm 1).

---

### Algorithm 1 Planning with VAPORS

---

- 1: **for**  $\tau \in \{1, \dots, T\}$  **do**
  - 2:    $I_t, D_t \leftarrow$  Get current RGBD image observation
  - 3:    $\hat{M}_t = f_{\text{seg}}(I_t)$  // Infer segmentation mask
  - 4:    $\hat{h}_\tau^k = \pi_H(\hat{M}_{1:t}, a_{1:t-1})$  // Select high-level action
  - 5:   Execute  $\pi_L(\hat{M}_t, \hat{h}_\tau^k)$
- 

## 550 B.3 Food Segmentation Training Details

551 **Self-Supervised Dataset Generation.** To circumvent the painstaking process of pixel-level  
 552 segmentation annotation for real food images, we design a self-supervised annotation procedure.  
 553 First, we record a grayscale RGB image of an empty plate,  $I_{\text{empty}} \in \mathbb{R}_+^{W \times H}$ . Next, we manually  
 554 place food items on the plate at random without changing the position of the plate, yielding a new  
 555 grayscale observation  $I_t$ . Let  $I_{\text{diff}} = |I_t - I_{\text{empty}}|$ , the framewise absolute difference between the  
 556 full and empty plate. We initialize the ground truth segmentation mask  $M_t$  corresponding to  $I_t$  as a  
 557 2D array of zeros, and then assign  $M_t[I_{\text{diff}} > \text{THRESH}] = 1$ . In practice, we find that  $\text{THRESH} = 20$   
 558 reasonably separates the foreground from the background to detect food. With this procedure, we  
 559 can scalably collect 280 paired RGB food images and segmentation masks in real within an hour  
 560 and a half of data collection. This includes plate resets, food placement, image capture, and offline  
 561 background subtraction post-processing.

562 **Augmentation.** We augment this dataset 8X by randomizing the linear contrast, gamma contrast,  
 563 Gaussian blur amount, saturation, additive Gaussian noise, translation, and rotation of each  
 564 RGB image, applying only the affine component of these same transformations to the associated  
 565 segmentation masks.

566 **Training Objective.** We train  $f_{\text{seg}}$ , implemented as a fully convolutional FPN (Feature Pyramid  
 567 Network) using Dice loss:

$$\mathcal{L}_{\text{dice}} = 1 - \frac{2 \times \text{TP}}{2 \times \text{TP} + \text{FN} + \text{FP}} \quad (6)$$

568 This objective encourages high overlap between predicted and ground truth masks, as TP, FN, FP  
 569 denote the number of pixel-level true positives, false negatives, and false positives in a prediction  $\hat{M}_t$   
 570 compared to ground truth  $M_t$ .

## 571 C Experimental Details

### 572 C.1 Noodle Acquisition Hardware Setup

573 Using a Franka Panda 7DoF robot, we aim to clear a plate of cooked noodles within a horizon of  
 574  $T = 10$  actions. We fit the end-effector with a custom 3D-printed mount consisting of a RealSense  
 575 D435 camera and a fork. To enable autonomous twirling and scooping capabilities, we extend the  
 576 fork’s range of motion via two servo motors (Dynamixel XC330-M288-T). We control the robot  
 577 with a Cartesian impedance controller, where the programmable servos are integrated in the forward  
 578 kinematics chain for positional control of the fork tip. The action space consists of either *group*  
 579 (rearrangement) or *twirl* (acquisition) actions, instantiated according to the learned segmentation and  
 580 pose estimation models detailed in Section 4.2.

581 A group action consolidates a sparsely distributed plate by sensing the furthest and densest points,  
 582  $(\hat{x}_f, \hat{y}_f, \hat{z}_f)$  and  $(\hat{x}_d, \hat{y}_d, \hat{z}_d)$ , and executing a planar push from the furthest to densest point. In a twirl  
 583 action, we infer the densest point and appropriate insertion angle  $\hat{\gamma}$ , roughly orthogonal to the grain  
 584 of majority of the noodles. We use positional control to insert the fork into the densest noodle pile,  
 585 and execute a fixed twirling motion by making two rotations at 6 radians per second. Finally, the  
 586 fork scoops upward until nearly horizontal ( $\beta = 80^\circ$ ) and the robot brings the acquired noodles to a  
 587 neutral position in the workspace.

588 For all trials, we use a non-slip plastic dinner plate, and mimick a bite successfully taken by a user  
 589 after twirling by autonomously untwirling onto a discard plate.

### 590 C.2 Bimanual Scooping Hardware Setup

591 We assume access to two Franka Panda robots, equipped with a pusher tool and a metal spoon,  
 592 respectively, and an external RealSense D435 camera for perception. With this setup, we aim to  
 593 acquire granular items on a plate using either *group* (rearrangement) or *scoop* (acquisition) actions,  
 594 with a total action budget of  $T = 8$  actions. In particular, we evaluate our system on the task of  
 595 scooping jelly beans, but VAPORS is agnostic to the exact choice of food. Following the experimental  
 596 setup of Grannen et al. [7], the spoon is mounted at an angle to the robot end-effector ( $\beta = 30^\circ$ ).  
 597 The pusher is a concave 3D-printed tool intended to push piles of items into the spoon and maintain  
 598 contact during lifting so as to prevent spillage.

599 Grouping actions are unimanual and use the pusher tool to push the sensed furthest item to the densest  
 600 region on the tray. In a scoop action, we sense the densest pile and execute a parameterized motion in  
 601 which the pusher and spoon move towards each other synchronously at a fixed  $\gamma = 180^\circ$ . Once they  
 602 arms are within a fixed threshold apart, the spoon scoops by tilting to  $\beta = 80^\circ$  and lifting to a neutral  
 603 workspace position.

604 We conduct all trials on a standard cooking tray due to the enlarged manipulation workspace for two  
 605 arms. To simulate a user’s bite between actions, we manually discard the spoon contents after a scoop  
 606 action.

### 607 C.3 Implementation Details

608 For each task, we use the following training procedures. We train  $\pi_H$  on simulated segmentation  
 609 observations of size  $64 \times 64$  for 2, 250 update steps, where we collect 1 episode every 150 update  
 610 steps. We instantiate the reward as per Eq. (3) with  $\alpha = 0.66$ , and train each model using the Adam  
 611 optimizer with with a learning rate of  $10^{-3}$ ,  $\epsilon = 10^{-4}$ , and gradient clipping norm of 1000 with  
 612 batch size  $B = 32$ , based on the training procedure from [37]. Each model takes approximately 1  
 613 hour to train on an Nvidia RTX A4000 GPU. To instantiate  $\pi_L$ , we train  $f_{\text{seg}}$  and  $f_{\text{ori}}$  from real data.  
 614 For segmentation, we collect 280 paired examples of images and binary segmentation masks using  
 615 the self-supervised annotation process from Section 4.2, where we use cooked noodles of randomized

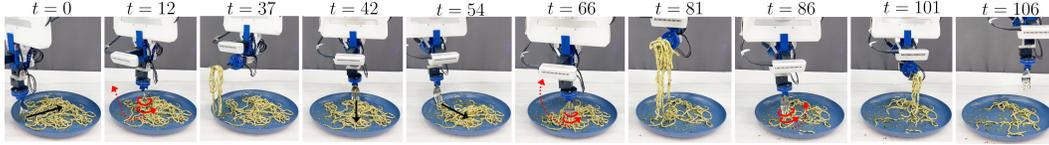


Figure 9: **Noodle Acquisition Rollout:** We visualize 6 actions performed by VAPORS on the task of clearing an initially half-full plate of Tier 3 noodles. As the distribution of noodles on the plate becomes sparse ( $t = 0, 42, 54$ ), VAPORS employs grouping strategies (black) to push noodles in close proximity. Once consolidated, VAPORS employs twirling ( $t = 12, 66, 86$ ), as shown in red, for efficient plate clearance, where  $t$  denotes the clock time in seconds.

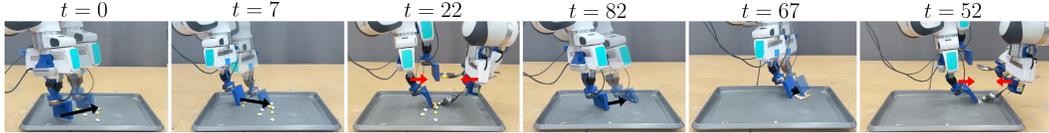


Figure 10: **Bimanual Scooping Rollout:** Using a bimanual setup with two Franka Emika Panda robots, VAPORS performs 6 actions consisting of grouping (black arrows) and scooping (red arrows) to acquire jelly beans on a tray. By grouping when the tray is sparse and acquiring when a bite-sized clump forms, VAPORS demonstrates efficient acquisition. The annotated timestamps denote clock time in seconds.

616 shape and sauce variations as well as jelly beans of randomized colors. We augment each dataset 10X  
 617 and train for 50 epochs, which takes approximately 3 hours on an NVIDIA GeForce RTX 2080 GPU.  
 618 In order to instantiate the twirl primitive for noodle acquisition, we additionally train  $f_{ori}$  to predict  
 619 fork tine orientation  $\gamma$  from 280 manually annotated crops of noodles as per Section 4.2, augmented  
 620 8X. The train time for  $f_{ori}$  is approximately 1 hour on an NVIDIA GeForce RTX 2080 GPU. For  
 621 deployment, we use an Intel NUC 7 for inference and robot control via a ROS 2-based control stack.

#### 622 C.4 Additional Experimental Results

623 In this section, we supplement the experimental findings from Section 5 with additional results.

624 **Plate Clearance:** Fig. 9 and Fig. 10 visualize two rollouts of VAPORS on plate clearance. We note  
 625 that visually, VAPORS tends to favor grouping as the plates become sparser and otherwise acquires  
 626 when there is a reasonably sized bite available.

627 **VAPORS Failure Mode Categorization:** In addition to evaluating the percentage of the plate  
 628 cleared, we observe the occurrence of a few failure modes, as depicted in Fig. 11. A *misplanned*  
 629 action (A) can occur due to a perception error, such as accidentally perceiving sauce, a garnish,  
 630 a vegetable, or plate specularities for noodles and erroneously grouping or twirling in that region.

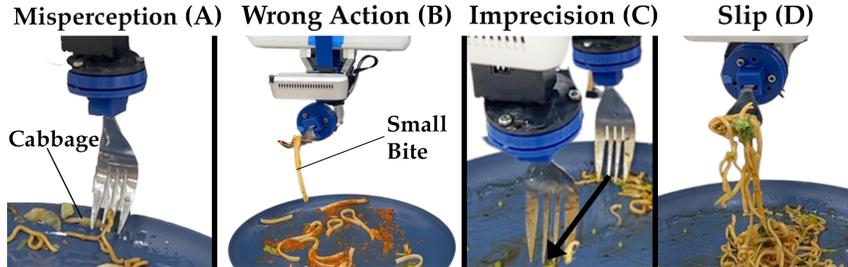


Figure 11: **VAPORS Failure Modes:** We illustrate the 4 most commonly observed failure modes with VAPORS on noodle acquisition. Misperception (A) occurs when  $\pi_L$  erroneously senses vegetables, sauce, or plate glare as a noodle due to false positives with  $f_{ori}$ , leading to a misplanned action such as grouping in that region. Occasionally,  $\pi_H$  may acquire when rearrangement is more appropriate, leading to a low-volume bite (B). In terms of action execution, food acquisition requires care so as to not miss food (C), as seen in the grouping motion which fails to group singular noodle strands due to system imprecision. Finally, slippage (D) can happen during acquisition with highly adversarial items such as those coated in sauce.

631 Alternatively, this can happen when (B) the robot twirls when grouping is more appropriate or vice  
 632 versa. A *mis-executed* action failure occurs when (C) the fork fails to group or acquire due to system  
 633 imprecision or (D) the noodles slip during acquisition due to sauce. In Table 1, we also report the  
 634 per-action failure rate, computed as the total number of failures over the total number of actions (60  
 635 = 6 trials  $\times$   $T = 10$ ).

636 **Qualitative User Study:** In the Likert survey administered to gauge user preferences across  
 637 methods, we report in Fig. 6 the statistical findings which are significant. In Table 2, we indicate the  
 638 specific margin of significance for each of the criteria, obtained via 1-way ANOVA testing.

Table 2: 1-Way ANOVA Statistically-Significant Findings ( $p$ -value  $< 0.05$ )

Criterion	Method 1	Method 2	$p$ -value	Criterion	Method 1	Method 2	$p$ -value
Efficiency	Heuristic	VAPORS	0.0004	Efficiency	Heuristic	Acquire Only	0.0029
Efficiency	Acquire Only	VAPORS	0.0010	Practicality	Heuristic	Acquire Only	0.0091
Bite Size	Acquire Only	VAPORS	0.0318	Reuse	Heuristic	Acquire Only	0.0093
Humanlike	Heuristic	VAPORS	0.0003	Efficiency	Heuristic	Acquire Only	0.0029
Humanlike	Acquire Only	VAPORS	0.0044	Efficiency	Acquire Only	VAPORS	0.0000
Practicality	Heuristic	VAPORS	0.0012	Bite Size	Heuristic	VAPORS	0.0029
Practicality	Acquire Only	VAPORS	0.0025	Bite Size	Acquire Only	VAPORS	0.0002
Reuse	Heuristic	VAPORS	0.0000	Humanlike	Acquire Only	VAPORS	0.0094
Reuse	Acquire Only	VAPORS	0.0008	Practicality	Heuristic	Acquire Only	0.0091
Trust	Heuristic	VAPORS	0.0124	Practicality	Acquire Only	VAPORS	0.0001
Trust	Acquire Only	VAPORS	0.0198	Reuse	Heuristic	Acquire Only	0.0093
Generalizability	Heuristic	VAPORS	0.0478	Reuse	Acquire Only	VAPORS	0.0001
				Trust	Acquire Only	VAPORS	0.0438
				Generalizability	Acquire Only	VAPORS	0.0481

Table 3: Noodle Acquisition

Table 4: Bimanual Scooping

639 In addition to the user study outlined in Section 5, we administered a second part of the study, in  
 640 randomized order to the first, in which users were asked to pick a preferred method for feeding in  
 641 side-by-side comparisons of jelly bean acquisition trials. To control for the initial state of the jelly  
 642 beans, we purposely arrange 16 beans into a  $4 \times 4$  grid initially, and conduct two trials per method  
 643 which are randomly selected for the comparisons. Although we would like to include an analogous  
 644 side-by-side comparisons survey for noodle acquisition for completeness, we find in practice that  
 645 controlling for the initial state of noodles is nontrivial due to their highly deformable nature and  
 646 vast set of feasible initial configurations. This makes it difficult to present users with unbiased  
 647 comparisons across methods.

648 Thus, for bimanual scooping, we presented all permutations of pairs of the three methods, for a total  
 649 of six comparisons overall. Empirically, we find that **VAPORS** is the preferred method by a large  
 650 margin compared to both baselines (Fig. 12).

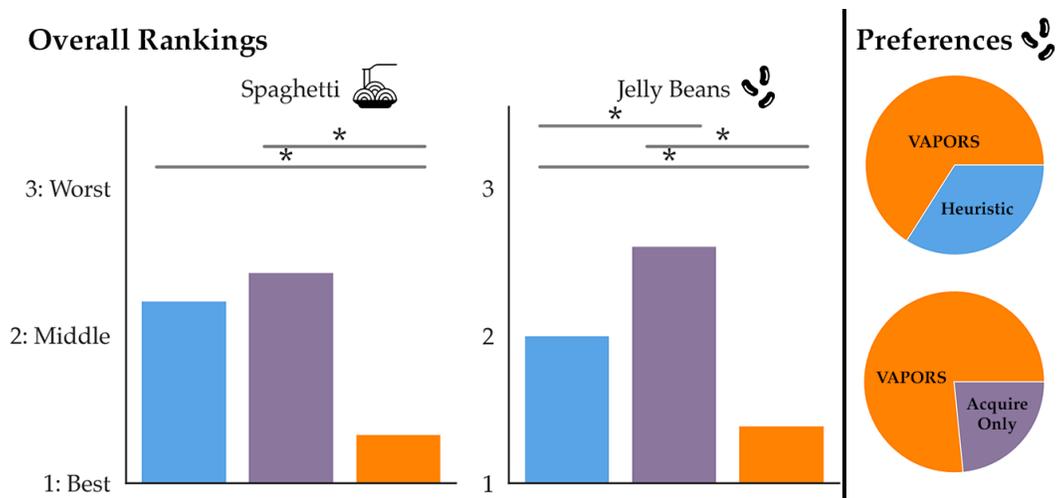


Figure 12: **Overall Ratings** Left: After observing all methods perform acquisition across 10 trials, we ask users to rank all three methods from most to least preferable. We find the **VAPORS** is most consistently ranked the best by a statistically significant margin ( $p < 0.05$ , denoted ‘\*’) compared to the baselines. Right: For jelly bean acquisition, we control for the initial state of the plate by arranging the beans in a  $4 \times 4$  grid, and ask users to select their preferred method across 6 side by side acquisition videos of different methods. **VAPORS** is the preferred method by a large margin compared to Heuristic and Acquire-Only.