

# Leveraging Post Hoc Context for Faster Learning in Bandit Settings with Applications in Robot-Assisted Feeding

Ethan K. Gordon<sup>1</sup>, Sumegh Roychowdhury<sup>2</sup>, Tapomayukh Bhattacharjee<sup>1</sup>, Kevin Jamieson<sup>1</sup>, and Siddhartha S. Srinivasa<sup>1</sup>

**Abstract**—Autonomous robot-assisted feeding requires the ability to acquire a wide variety of food items. However, it is impossible for such a system to be trained on all types of food in existence. Therefore, a key challenge is choosing a manipulation strategy for a previously unseen food item. Previous work showed that the problem can be represented as a linear bandit with visual context. However, food has a wide variety of multi-modal properties relevant to manipulation that can be hard to distinguish visually. Our key insight is that we can leverage the haptic context we collect during and after manipulation (i.e., “post hoc”) to learn some of these properties and more quickly adapt our visual model to previously unseen food. In general, we propose a modified linear contextual bandit framework augmented with post hoc context observed after action selection to empirically increase learning speed and reduce cumulative regret. Experiments on synthetic data demonstrate that this effect is more pronounced when the dimensionality of the context is large relative to the post hoc context or when the post hoc context model is particularly easy to learn. Finally, we apply this framework to the bite acquisition problem and demonstrate the acquisition of 8 previously unseen types of food with 21% fewer failures across 64 attempts.

## I. INTRODUCTION

Many of us take eating for granted, but approximately 1.0 million people in the US alone cannot eat without assistance [1]. Autonomous robot-assisted feeding could save time for caregivers and increase people’s sense of self worth [2,3]. Existing robotic feeding systems [4,5] rely on preprogrammed movements, making it difficult for them to adapt to environmental changes. In general, a robust feeding system must be able to acquire a bite of food in an uncertain environment, a task known as “bite acquisition.” This work focuses on the acquisition of food items that the robot may not have seen or manipulated before.

Previous work has identified useful manipulation strategies for a variety of food items [6] and ways to generalize those strategies to similar-looking food items [7,8], but one challenge is figuring out how to select a manipulation strategy for previously unseen food. Recent work has suggested modeling bite acquisition as a contextual bandit [9]. The robot can observe visual context for each food item, select from a set of discrete manipulation strategies, and observe partial (or bandit) feedback in the form of a binary success or failure.

However, the method presented in that work required about 10 failures over 25 attempts to learn the optimal manipulation strategy for a given new food item. Visual




Fig. 1. The Assistive Dextrous Arm (ADA) can leverage post hoc context to quickly figure out which action is best for picking up a banana slice. *Left*: Banana slices have never been seen before. Acquired visual context. *Middle*: Haptic data suggests that banana slices are soft. *Right*: Combining visual and haptic contexts, appropriate action is predicted.

context is only a single mode. Food items that look similar (but not identical), such as ripe and un-ripe banana slices, can have very different consistencies, leading to different optimal manipulation strategies. It is hard to learn this map with only visual data. Haptic feedback from physical interactions can be informative for object classification [6,10,11] and inferring object properties such as haptic adjectives [12], rigidity [13], elasticity [14], hardness [15], and compliance [16,17]. Previous work has also shown that combining visual and haptic modalities can help towards inferring global haptic mapping [18] and learning multi-modal representations [19] during manipulation. Our key insight is that *we can leverage haptic feedback collected after action selection during manipulation to more quickly learn how to map visual information to the optimal strategy for a given type of food.*

More generally, we propose augmenting the traditional contextual bandit framework with post hoc, multi-modal context collected after action selection. We show empirically that when the post hoc context is relatively simple to model, it can be used to more quickly model the regular context, leading to faster learning and lower cumulative regret. Our major contributions are (1) a proposed modification to LinUCB [20] and online linear regression to handle post hoc context, (2) experimental results on synthetic data that provide insights on the efficacy of this framework, and (3) empirical evidence that demonstrates improvement over the traditional contextual bandit setting in a real robot bite acquisition task relevant to the application of robot-assisted feeding. Our current action space is small, consisting of only 6 discrete strategies, but future work can leverage these

<sup>1</sup> Ethan K. Gordon, Tapomayukh Bhattacharjee, Kevin Jamieson, and Siddhartha S. Srinivasa are with the Department of Computer Science and Engineering, University of Washington, Seattle, WA 98195 {ekgordon, tapo, jamieson, siddh}@cs.washington.edu

<sup>2</sup>Sumegh Roychowdhury is with the Indian Institute of Technology Kharagpur, Kharagpur, India, sumegh01@iitkgp.ac.in. Work done as a UW intern.

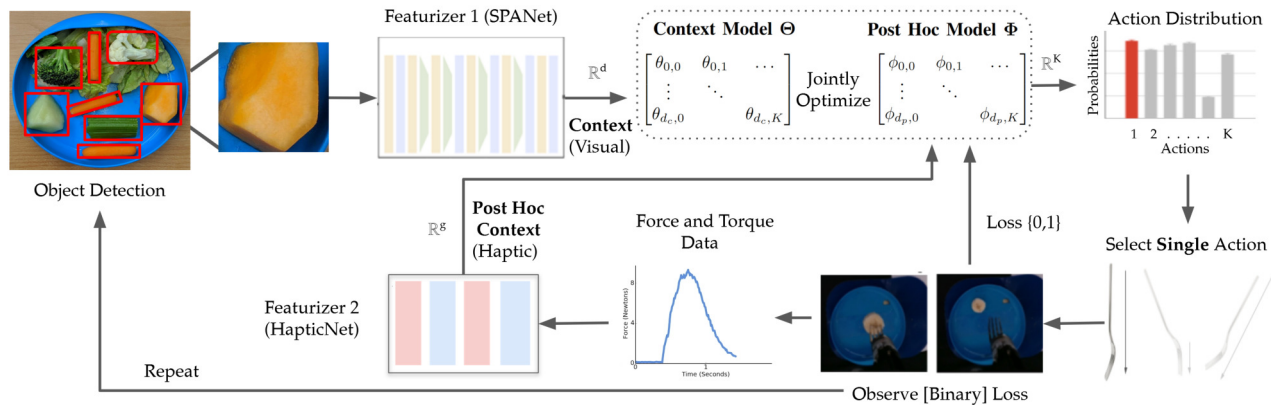


Fig. 2. Post hoc augmented contextual bandit framework. We only observe the visual context from SPANet prior to action selection, but the post hoc context from HapticNet is used with the observed loss to update the visual model.

insights to examine a larger action space that can handle an even wider variety of food items and realistic plates.

## II. RELATED WORK

### A. Online Learning and Contextual Bandits

Contextual bandit algorithms have seen widespread success in health interventions [21,22], online advertising [23,24], adaptive routing [25], clinical trials [26], music recommendations [27], education [28], and financial portfolio design [29]. Adoption in robotics has included selecting trajectories for object rearrangement [30], kicking strategies in robotic soccer [31], selecting among deformable object models for acquisition tasks [32], and the aforementioned work on selecting manipulation strategies for deformable food items [9]. All of these applications leverage a single context vector observed prior to action selection and the scalar incurred loss or reward. We propose leveraging higher dimensional feedback observed after action selection to speed up learning.

Baseline exploration strategies include epoch-greedy [33], LinUCB [34], RegCB [35] and Online Cover [36]. For a recent and thorough overview, we refer the interested reader to [37,38]. Our work is distinct from bandits with delayed feedback [39] in that the post hoc context is not delayed by any time steps, but just observable after action selection. Our work could potentially be compared with the bandits-with-expert-advice setting and associated algorithms like EXP4 [40], in the sense that the context model and the post hoc context model could be thought of as competing action recommendations, though the experts in this setting generally make their predictions exclusively prior to action selection.

### B. Robot-Assisted Feeding: Food Manipulation

General food manipulation has been studied in various environments, such as the packaging industry [41–46]. These tend to focus on the design of application-specific grippers for robust sorting and pick-and-place. Other work shows the need for visual sensing for quality control [47–49] and haptic sensing for grasping deformable food items without damaging them [41–46]. Research labs have also explored meal preparation [50,51], baking cookies [52], making pancakes [53], separating Oreos [54], and preparing meals [55] with robots. Most of these studies either interacted with a specific food item with a fixed manipulation strategy [52,53] or with an unchanging set of food items and manipulation

strategies [7,8,56,57]. Some of these studies have looked at using multi-modal data [55] or online learning [9], but not a combination of the two.

Our visual context is generated using the *Skewering Position Action Network* (SPANet) from [7] while our action space and our haptic context specification are derived from human data [6].

## III. PRELIMINARY: CONTEXTUAL BANDITS

Previous work [9] showed the utility of representing the bite acquisition setting as a contextual bandit. For each attempt, the agent observes  $d_c$ -dimensional visual context (an image of the food) and selects from a discrete set of  $K$  manipulation strategies, observing a loss of 0 on success and 1 on failure. Here we cover the specifics of this setting that are relevant to our proposed augmentation.

### A. Formulation

General online supervised learning has an agent learn a map  $f : \mathbb{R}^{d_c} \rightarrow \mathbb{R}^K$  between a  $d_c$ -dimension context vector  $c$  and a  $K$ -dimension loss vector  $l$  given a sample  $(c_t, l_t)$  at each time step  $t$ . In a discrete interactive learning setting, the agent will first observe the context  $c_t$ , choose an action  $a_t \in [K]$ , and then observe the full loss vector  $l_t$  while incurring loss  $l_t[a_t]$ . The agent’s goal in this setting is to minimize *cumulative regret*

$$R_T := \max_{a'} \sum_t (l_t[a_t] - l_t[a']) \quad (1)$$

the difference between the loss incurred by the agent and the lowest loss it was possible to incur.

In a contextual bandit setting, the agent is restricted to *bandit feedback*: observing only the loss incurred ( $l_t[a_t]$ ) rather than the full loss vector  $l_t$ . This creates a trade-off between exploring actions we are unsure about and exploiting actions likely to incur little loss. In general, a contextual bandit algorithm consists of two parts: (1) an *exploration strategy* that determines which action to take at each time step given  $c_t$  and some policy  $\pi : c_t \rightarrow a_t$ , and (2) a *learner* that incorporates the bandit feedback received into the  $\pi$ .

### B. Learning: Online Linear Regression

Assume that the true map  $f^*$  exists in some function class  $\mathcal{F}$ . One method for solving the contextual bandit setting is

to reduce the problem to regular online supervised learning and create an estimate of this function  $\hat{f}$  with least squares regression. Importance weighting [38] can eliminate the bias that comes from only using partial feedback from random exploration, but as our proposed exploration strategy is deterministic, we do not need to do this here.

Previous work [7] demonstrated a model could accurately recommend the optimal action for a given food item in a full supervised learning setting. Subsequent work [9] treated all but the last layer of SPANet as a fixed featurizer, treating the final linear layer as the context model. With this motivation, we assume that  $f^*$  is linear and all observed noise is Gaussian, such that  $l_t[a] = \theta_a^\top c_t + \epsilon$  with weights  $\theta_a \in \mathbb{R}^{d_c}$  and noise  $\epsilon \sim \mathcal{N}(0, \mathbf{I}\sigma^2)$ . Applying least squares regression to this linear model allows us to produce familiar weight estimates

$$\hat{\theta}_a = \left( \mathbf{C}_a^\top \mathbf{C}_a \right)^{-1} \mathbf{C}_a^\top L_a \quad (2)$$

where  $\mathbf{C}_a \in \mathbb{R}^{T_a \times d_c}$  is the matrix contexts observed during the  $T_a$  time steps where the agent selected action  $a$ ,  $L_a \in \mathbb{R}^{T_a}$  is the vector of scalar losses observed on those same time steps.

### C. Exploration: LinUCB

As described in Section II-A, there are many existing exploration strategies for contextual bandit settings. We initially focus on the Linear Upper Confidence Bound (LinUCB) [20] algorithm in the work since previous work [9] suggests that it performs well empirically in the bite acquisition setting.

LinUCB implicitly balances exploration and exploitation by using the estimated linear model to construct a confidence interval around  $l_t$  for a given  $c_t$  and optimistically playing  $a_t$  with the lowest lower confidence bound on its expected loss. In this way, the algorithm prefers relatively unknown actions with larger intervals (encouraging exploration) and actions with low loss (encouraging exploitation). UCB-style algorithms like this are known to achieve cumulative regret bounded by  $\mathbb{E}[R_T] \leq \tilde{O}(d_c \sqrt{T})$  [58].

For a given confidence level, this lower bound can be calculated as [20]

$$LCB(a) = \hat{\theta}_a^\top c_t - \alpha \sqrt{c_t^\top \Sigma_a c_t} \quad (3)$$

for some constant  $\alpha > 0$  and  $\Sigma_a := \left( \mathbf{C}_a^\top \mathbf{C}_a \right)^{-1}$  the covariance of the estimator  $\hat{\theta}_a$ .

We are further motivated to use LinUCB because, through this covariance matrix, it uses information about the learning scheme to tune its exploration. This is in contrast with learner-agnostic strategies like  $\epsilon$ -greedy. Therefore, when we modify the learning scheme with post hoc context, LinUCB will seamlessly incorporate the extra knowledge into its exploration strategy.

## IV. LEVERAGING POST HOC CONTEXT

In general, we propose augmenting the conventional contextual bandit setting with an action-independent  $d_p$ -dimension post hoc context vector  $p_t$  (haptic force-torque feedback in our setting) observable after selecting the action  $a_t$ . We can justify this inclusion as follows: At each time step, assume there exists some hidden state  $z_t$  (for example,

all the information describing an unripe banana slice). In our setting, the context  $c_t$  (e.g., the picture of the food) and the post hoc context  $p_t$  (e.g., haptic parameters only available after the action is taken) are just alternate representations of this underlying state (e.g., we can name the food item by either looking at it or touching it). Therefore, if there exists some function  $h$  that maps the state onto the loss vector, then there should also exist functions  $f$  and  $g$  that map the context and post hoc context respectively onto that same loss vector. In other words, we assume  $\mathbb{E}[l_t] = h(z_t) = f(c_t) = g(p_t)$ .

This augmentation maps neatly onto the bite acquisition setting. At each round  $t = 1, \dots, T$ , the interaction protocol consists of

- 1) *Context observation.* The user selects a food item to acquire. We observe an RGBD image containing the single discrete food item (detected using RetinaNet [59]). We pass this through SPANet [9] which returns the visual context  $c_t \in \mathbb{R}^{2048}$ .
- 2) *Action selection.* The algorithm selects a discrete manipulation strategy  $a_t \in \mathcal{A} = \{1, 2, \dots, K\}$ . In our initial implementation,  $K = 6$ , matching a subset of the taxonomy of manipulation strategies from [6]. Action execution is detailed in Section VI-A.
- 3) *Bandit loss observation.* The environment provides a binary loss  $l_t(c_t)[a_t] \in \{0, 1\}$ , where  $l_t = 0$  corresponds to the robot successfully acquiring the single desired food item.
- 4) *Post hoc context observation.* During action execution, time series force and torque data is passed through HapticNet (described below) to create the haptic context  $p_t \in \mathbb{R}^4$ .

To expand on (4), HapticNet is a small multi-layer perceptron (MLP) from [8]. The first 50ms of force and torque data after contact with the food (as determined by force thresholding) are passed through two ReLU layers. The output is the softmax-ed vector  $p_t$  classifying the food as ‘‘hard’’, ‘‘medium’’, ‘‘soft’’, or ‘‘hard-skin’’. Importantly, the hardness of the food is intrinsic, independent of the action. In [8], the manipulation action used by human participants was directly affected by this categorization, and so we are motivated to use a linear model here as well.

The structure of this environment as it pertains to bite acquisition is shown in Fig. 2.

### A. Learning: Joint Model Regression

Recalling our assumption  $\mathbb{E}[l_t] = f(c_t) = g(p_t)$ , we propose jointly estimating  $f$  and  $g$  with least squares regression under the constraint that they produce the same outputs, i.e.  $f(c_t)[a] = g(p_t)[a] \forall a$ . This constraint could be a soft constraint, weighting the square difference between the outputs by hyperparameters. However, since the contextual bandit setting in general does not come with a well-defined training and validation set, it is desirable to reduce the number of hyperparameters requiring tuning. Therefore, in this work, we only consider using a hard constraint. Importantly, this constraint should be valid for all actions, allowing all time steps to factor into the estimate no matter which action was taken.

To demonstrate this, we jump into the linear setting, where




Fig. 3. Experiment 1 shows the effect of dimensionality on the regret difference between the *Context Only* and *Post Hoc Augmented* model. The context vector is sampled from a uniform distribution and the post hoc vector is constructed using first  $d_p$  components. The full loss vector is defined as  $l_t = p_t^\top \phi^*$  where  $\phi^*$  is unknown to the model. We observe an improvement in regret difference as  $d_c$  is increased from 10 (*Center*) to 100 (*Right*).

$$f(c_t)[a] = \theta_a^\top c_t \text{ and } g(p_t)[a] = \phi_a^\top p_t.$$

$$\hat{\theta}_a, \hat{\phi}_a := \arg \min_{\theta_a, \phi_a} \|\mathbf{C}_a \theta_a - L_a\|_2^2 + \|\mathbf{P}_a \phi_a - L_a\|_2^2 \quad (4)$$

$$\text{s.t. } \mathbf{C} \theta_a = \mathbf{P} \phi_a \quad (5)$$

As before,  $\mathbf{C}_a \in \mathbb{R}^{T_a \times d_c}$ ,  $\mathbf{P}_a \in \mathbb{R}^{T_a \times d_p}$ , and  $L_a \in \mathbb{R}^{T_a}$  are matrices of importance weighted contexts, post hoc contexts, and losses respectively. The constraint, being valid for all actions, uses the full context matrix  $\mathbf{C} \in \mathbb{R}^{T \times d_c}$  and post hoc context data matrix  $\mathbf{P} \in \mathbb{R}^{T \times d_p}$ . Using the constraint to define a transformation matrix  $\phi_a = (\mathbf{P}^\top \mathbf{P})^{-1} \mathbf{P}^\top \mathbf{C} \theta_a := \mathbf{H} \theta_a$ , we can solve for the weight estimate.

$$\hat{\theta}_a = \left[ \mathbf{C}_a^\top \mathbf{C}_a + \mathbf{H}^\top \mathbf{P}_a^\top \mathbf{P}_a \mathbf{H} \right]^{-1} \left[ \mathbf{C}_a^\top + \mathbf{H}^\top \mathbf{P}_a^\top \right] L_a \quad (6)$$

This formulation provides a normative reason to expect empirical improvements over the context-only setting. Consider the case exemplified by Fig. 1. If the post hoc context model is known perfectly, it can recommend the correct action for a given context after only a single attempt, cutting down exploration by a factor of  $K$ . More formally, if we know  $\phi_a^* \forall a$ , then we know what the expected loss  $\mathbb{E}[L] = \mathbf{P} \phi_a^*$  would have been for action  $a$  at all time steps, *including time steps where we did not take action  $a$* . In other words, we can rewrite the hard constraint from Equation 5 as

$$\theta_a = (\mathbf{C}^\top \mathbf{C})^{-1} \mathbf{C}^\top \mathbf{P} \phi_a = (\mathbf{C}^\top \mathbf{C})^{-1} \mathbf{C}^\top \mathbb{E}[L] \quad (7)$$

This surface defined by the hard constraint is just the solution to standard linear regression on all time steps. The upshot is that knowing the post hoc context model reduces the problem from bandit feedback to full feedback regression, which is a much easier problem.

### B. Exploration: Modified LinUCB

Under the linearity assumption, i.e.,

$$L_a = \mathbf{C}_a \theta_a^* + \epsilon = \mathbf{P}_a \phi_a^* + \epsilon \quad (8)$$

We can show that  $\hat{\theta}_a$  is an unbiased estimate of  $\theta_a^*$  with a covariance bounded from above (via Cauchy-Schwartz and Jensen's inequalities) by

$$\Sigma_p := 2 \left( \mathbf{C}_a^\top \mathbf{C}_a + \mathbf{H}^\top \mathbf{P}_a^\top \mathbf{P}_a \mathbf{H} \right)^{-1} \quad (9)$$

From here, following the same logic as [20], we can construct a lower confidence bound equivalent to Equation 3 and replacing  $\Sigma_a$  with  $\Sigma_p$ . This work focuses on comparing this modified algorithm to the baseline version of LinUCB as described in Section III-C.

Further details about the environment and algorithms in this section can be found at [60].

## V. SYNTHETIC DATA EXPERIMENTS

Before implementing this framework on the robot, we conducted two experiments with synthetic post hoc context to validate the potential benefits from this setting. Experiment 1 is designed to demonstrate that a low-dimension post hoc context vector can lead to faster learning, even if it contains no new information. It does so by varying the size of the synthetic context vector while keeping the size of the synthetic post hoc context vector fixed. Experiment 2 is designed to see if a well trained post hoc context model can effectively reduce the contextual bandit to the easier problem of full feedback online learning. The synthetic post hoc context is constructed to be easy to learn, while the context vectors, derived from MNIST data, do not even adhere to the linear model assumption.

### A. Experiment 1: Low Dimension Post Hoc Context

The setup for this experiment is outlined in Fig. 3(Left). We first fix the number of actions  $K = 10$  and the dimension of the post hoc context  $d_p = 3$ . We then generate a random, hidden pseudo-invertible post hoc context model  $\phi^* \in \mathbb{R}^{d_p \times K}$ . At each time step, we sample a context vector  $c_t \sim [0, 1]^{d_c}$ , which is shown to the bandit algorithm. The first  $d_p$  components are defined to be the post hoc context, and the full loss vector  $l_t = p_t^\top \phi^*$  is computed accordingly. The algorithm incurs regret  $l_t[a_t] - \min_a l_t[a]$  and is shown  $l_t[a_t]$  and  $p_t$ .

*Results:* For  $d_c = 10$  and  $d_c = 100$ , we run the context-only and the post-hoc-augmented learners for 40 trials, 1000 attempts per trial, and record the cumulative regret. The results are shown in Fig. 3(Center, Right). With the lower dimensional context, we observe that the two learners perform comparably, with a slight advantage to the post-hoc-augmented learner. However, with the higher dimensional context, the post-hoc-augmented learner exhibits significantly better performance than the context-only learner.

The context and the post hoc context contain the exact same information about the loss vector, but the post hoc context does so with fewer dimensions, making each observation more informative. These results support the idea that lower dimensional post hoc context is beneficial. Even if it contains no new information, it can still be used to train an accurate context model more quickly.




Fig. 4. Experiment 2 is performed on the standard MNIST dataset to test learning speeds of various models. The context vector is reduced using the PCA and the post hoc context vector is constructed from the loss using  $\phi^*$  which should be easy to learn. We see from the MSE plot (Center) that the *Post Hoc Augmented* model reaches its optimal value much faster than the *Context Only* model. As soon as the post hoc model is learned the context model quickly approaches its best possible value. Also from the Regret plots (Right) we can see the *Post Hoc Augmented* model achieving lower regret value than its *Context Only* counterpart.

## B. Experiment 2: Easy-To-Learn Post Hoc Context

The setup for this experiment is outlined in Fig. 4(Left). Context vectors are derived from the MNIST dataset [61], which consists of labeled  $28 \times 28$  images of hand-written digits (0 to 9) split into a training set with 60k samples and a test set with 10k samples. At each time step, we sample an image, and then use a PCA (trained on the training set) to reduce it to a  $d_c = 200$  dimension context vector. This vector is shown to the bandit algorithm, which returns an action  $a_t \in [K = 10]$ . For a given image, we can define the full loss vector  $l_t \in \{0, 1\}^{10}$  to be 1 if the incorrect digit is guessed and 0 otherwise. As in Experiment 1, we construct a random linear post hoc context model  $\phi^* \in \mathbb{R}^{d_p \times K}$  and use it to manually construct a post hoc context vector  $p_t \in \mathbb{R}^{d_p=10}$  from the full loss vector. This makes for a post hoc context model that is extremely easy to learn perfectly. The loss  $l_t[a_t]$  and the post hoc context  $p_t$  is shown to the bandit algorithm.

*Results:* First, to test learning speed, we ran a post-hoc-augmented learner and a context-only learner on the training set with random uniform exploration. Every 10 attempts, we freeze both linear context models  $\hat{\theta}$  and record its mean square error (MSE) on 2k samples from the test set.

$$MSE := \frac{1}{|\text{Test Subset}|} \sum_{c_t, l_t \in \text{Test Subset}} \|\hat{\theta}^\top c_t - l_t\|_2^2 \quad (10)$$

For the post-hoc-augmented learner, we also recorded the MSE of the post hoc context model  $\hat{\phi}$ . All three are plotted in Fig. 4(Center). Note that the best possible context model MSE (as determined by training on the full loss vectors from the entire training set) is 0.1383. We should expect no context model  $\theta$  to beat that, even one augmented by post hoc context.

As expected, the noise-less post hoc model  $\phi$  is learned perfectly as soon as it sees  $d_p = 10$  linearly independent samples for each digit (which happens within  $\sim 500$  total samples). At this point, the post-hoc-augmented learner significantly deviates from the context-only learner, and it has reached its plateau within another  $\sim 500$  samples. Meanwhile, the context-only learner still did not reach its plateau after 2000 samples. These results support the idea that a perfect context model effectively reduces the problem from bandit feedback to full feedback, allowing for faster learning.

We then combined each learner with a LinUCB exploration strategy, ran them on the entire 10k-sample test set, and recorded the cumulative regret. These results are shown

in Fig. 4(Right), and demonstrate a significant improvement of the post-hoc-augmented bandit over the context-only bandit. This shows that an improved learning speed can translate to reduced regret, the metric that we care about in this setting.

## VI. REAL WORLD BITE ACQUISITION

### A. System Description

Our setup, the Autonomous Dexterous Arm (ADA) (Fig. 5(Left)), consists of a 6 DoF JACO2 robotic arm [62]. The arm has 2 under-actuated fingers that grab a custom-built, 3D-printed fork holder. For haptic input, we instrumented the fork with a 6-axis ATI Nano25 Force-Torque sensor [63]. For visual input, we mounted a custom built wireless perception unit on the robot’s wrist; the unit includes the Intel RealSense D415 RGBD camera and the NVidia Jetson Nano for wireless transmission. Food is placed on a plate mounted on an anti-slip mat commonly found in assisted living facilities.

Each manipulation strategy is executed by a simple impedance controller with 3 parameters: (1) the angle of the fork handle relative to vertical, (2) the roll angle of the fork relative to the the major axis of the food item, and (3) the maximum force to impart on the food during skewering. Specifically, the 6 discrete actions in this setting, motivated by [6] and used in [9], are parameterized as follows

- 1) *Vertical Skewer 0 (VSO)*: Pitch 0, Roll 0, 25N
- 2) *Vertical Skewer 90 (VS90)*: Pitch 0, Roll  $\frac{\pi}{2}$ , 25N
- 3) *Tines Vertical 0 (TV0)*: Pitch -0.5, Roll 0, 20N
- 4) *Tines Vertical 90 (TV90)*: Pitch -0.5, Roll  $\frac{\pi}{2}$ , 20N
- 5) *Tilted Angled 0 (TA0)*: Pitch 0.4, Roll 0, 10N
- 6) *Tilted Angled 90 (TA90)*: Pitch 0.4, Roll  $\frac{\pi}{2}$ , 10N

### B. Offline Results and Tuning

The first step was to ensure that our post hoc context, the haptic data, was descriptive enough to potentially benefit the visual context model. To this end, we collected 115 samples of the robot skewering 3 food items, chosen to be representative of different haptic categories and optimal action (as determined in [7]): grape is classified as “hard skin” and has the optimal action TV90, strawberry is “medium” and prefers VSO or TV0, and banana is “soft” and prefers TA0 or TA90. For each sample, we recorded the visual context  $c_t$ , post hoc context  $p_t$ , action taken  $a_t$ , loss  $l_t[a_t]$ , and food type name (e.g. “grape”).

Since this data, by necessity, was collected with bandit feedback, we impute the full loss vector  $\hat{l}$  of each attempt




Fig. 5. Results of the bite acquisition experiment using the Autonomous Dexterous Arm (ADA) (Left). The MSE Difference plot (Center) shows that there is an early benefit in learning with the post hoc context which reduces over time. Also from the Cumulative Loss plot (Right) it’s evident that with *Post hoc Context Augmented* the cumulative loss incurred with increasing attempts is lower than it’s *Context Only* counterpart.

by averaging the loss for a given action across all samples of the same food type (e.g., average the loss of  $VS0$  across all bananas). While simple, this can introduce a herding bias relative to the real world. We can eliminate this bias (at the cost of increased variance) by using a doubly robust [64] estimator

$$\hat{l}_{DR}[a] = \hat{l}[a] + (l[i] - \hat{l}[a]) \frac{\mathbf{1}(i = a)}{\mathbb{P}[i = a]} \quad (11)$$

where  $\hat{l}[a]$  is the biased estimate from herding,  $\mathbb{P}[i = a]$  is the probability that we took action  $a$  during data collection ( $\frac{1}{6}$  in our case), and  $l[i]$  is the actual binary loss associated with that sample (if available).

Similarly to Experiment 2 (Section V-B), we divided this data set into 80 training examples and 35 test examples. We then ran a context-only learner and a post-hoc-augmented learner on the training set with uniform exploration, freezing the context model  $\theta_a$  after each time step to measure its MSE on the test set. The difference in MSE is shown in Fig. 5(Center). While the size and variance of the dataset makes it hard to show significance, the mean suggests that there may be a learning benefit to the post hoc context early on that fades over time.

We also used the full 115 samples to tune the exploration hyperparameter  $\alpha$  for both LinUCB implementations. The optimal value for both the context-only and the post-hoc-augmented learners was  $\alpha = 0.01$ , which we used for the online experiment.

### C. Online Experiment

For the online experiment, we ran both the post-hoc-augmented LinUCB and the context-only LinUCB algorithms on 8 previously unseen food types, 8 attempts per food type, for a total of 64 attempts per trial. The 8 food types included 2 “hard” foods (carrot and celery), 2 “medium” foods (strawberry and cantaloupe), 2 “soft” foods (banana and kiwi), and 2 “hard skin” foods (cherry tomato and grape). In the absence of noise, we would expect the robot to take about 48 attempts to figure out the optimal action for each food item: 6 actions for each of the 8 food items. Given that we expect the optimal action to be knowable from the haptic category, we would hope to decrease that convergence time by a factor of 2 with the post hoc context: 6 actions for each of the 4 haptic categories.

The results of the experiment are shown in Fig. 5(Right). In this setting, where it is impossible to observe the full loss vector  $l_t$ , we record cumulative loss  $\sum_t l_t[a_t]$  instead of cumulative regret. By this metric, we do see some improvement of the post-hoc-augmented agent over the context-only agent. Over 3 trials, the former experienced fewer failures,

accumulating an average total loss of 22 compared with 27.667 for the context-only learner.

## VII. DISCUSSION

One key takeaway from this work is the success of post-hoc-augmented bandits in a variety of empirical settings that potentially deviate significantly from the assumed linear model. This suggests that it may be fruitful to pair this augmentation with other empirically competitive exploration strategies such as Thompson Sampling.

With that said, a limitation of this work is the requirement that the post hoc context be independent of the action selected. While this happens to work in robot-world physical interactions like bite acquisition, it is harder to imagine it working in a domain like online recommendation systems, where a user’s experience (and thus, the feedback they can give as post hoc context) will be affected by, for example, the set of links they were shown. Future work should address this setting and potentially relax our strong assumption that  $\mathbb{E}[l_t] = f(c_t) = g(p_t) ; \forall a$ .

Finally, we note that our current action space is small and imperfect. We cannot expect to converge to 0-loss on all possible food items that a user might want to eat. In the future, we intend to broaden our scope to larger classes of food items by investigating continuous and expanded action and post hoc context spaces rather than limiting ourselves to expert-defined discrete options. It is unclear how to apply post hoc context to this environment. One possibility is to treat it as a function of the gradient of the loss function with respect to the action space  $\nabla_a l_t$ . We could also expand the action space by considering compound (or slate [65]) actions. The space of possible food items and acquisition strategies is massive, and users require a robust system if they need to use it daily [66], so a lot more research can be done here.

Overall, these results suggest that multi-modal feedback can be leveraged in interactive learning to allow for more data-efficient adaptive bite acquisition.

## ACKNOWLEDGMENTS

Research reported in this publication was supported by the Eunice Kennedy Shriver National Institute Of Child Health & Human Development of the National Institutes of Health under Award Number F32HD101192. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. This work was also (partially) funded by the National Science Foundation IIS (#2007011), National Science Foundation DMS (#1839371), the Office of Naval Research, US Army Research Laboratory CCDC, Amazon, and Honda Research Institute USA.

## REFERENCES

- [1] M. W. Brault, "Americans with disabilities: 2010," Current population reports, vol. 7, pp. 70–131, 2012.
- [2] S. D. Prior, "An electric wheelchair mounted robotic arm—a survey of potential users," Journal of medical engineering & technology, vol. 14, no. 4, pp. 143–154, 1990.
- [3] C. A. Stanger, C. Anglin, W. S. Harwin, and D. P. Romilly, "Devices for assisting manipulation: a summary of user task priorities," IEEE Transactions on rehabilitation Engineering, vol. 2, no. 4, pp. 256–265, 1994.
- [4] MySpoon, 2018, <https://www.secom.co.jp/english/myspoon/food.html>.
- [5] Obi, 2018, <https://meetobi.com/>.
- [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] 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," in *International Symposium on Robotics Research*, 2019.
- [8] 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.
- [9] E. K. Gordon, X. Meng, M. Barnes, T. Bhattacharjee, and S. S. Srinivasa, "Adaptive Robot-Assisted feeding: An online learning framework for acquiring Previously-Unseen food items," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.
- [10] A. Schneider, J. Sturm, C. Stachniss, M. Reisert, H. Burkhardt, and W. Burgard, "Object identification with tactile sensors using bag-of-features," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 243–248.
- [11] P. K. Allen and K. S. Roberts, "Haptic object recognition using a multi-fingered dexterous hand," in *IEEE International Conference on Robotics and Automation*, 1989, pp. 342–347.
- [12] V. Chu, I. McMahon, L. Riano, C. G. McDonald, Q. He, J. M. Perez-Tejada, M. Arrigo, T. Darrell, and K. J. Kuchenbecker, "Robotic learning of haptic adjectives through physical interaction," *Robotics and Autonomous Systems*, vol. 63, pp. 279–292, 2015.
- [13] A. Drimus, G. Kootstra, A. Bilberg, and D. Kragic, "Classification of rigid and deformable objects using a novel tactile sensor," in *International Conference on Advanced Robotics*, 2011, pp. 427–434.
- [14] B. Frank, R. Schmedding, C. Stachniss, M. Teschner, and W. Burgard, "Learning the elasticity parameters of deformable objects with a manipulation robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 1877–1883.
- [15] S. Takamuku, G. Gomez, K. Hosoda, and R. Pfeifer, "Haptic discrimination of material properties by a robotic hand," in *IEEE 6th International Conference on Development and Learning (ICDL)*, 2007, pp. 1–6.
- [16] M. Kaboli, P. Mittendorf, V. Hügel, and G. Cheng, "Humanoids learn object properties from robust tactile feature descriptors via multi-modal artificial skin," in *IEEE-RAS International Conference on Humanoid Robots*, 2014, pp. 187–192.
- [17] T. Bhattacharjee, J. M. Rehg, and C. C. Kemp, "Inferring object properties with a tactile-sensing array given varying joint stiffness and velocity," *International Journal of Humanoid Robotics*, pp. 1–32, 2017.
- [18] A. A. Sheno, T. Bhattacharjee, and C. C. Kemp, "A crf that combines touch and vision for haptic mapping," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*. IEEE, 2016, pp. 2255–2262.
- [19] 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.
- [20] C. Wei, L. Li, L. Reyzin, and R. E. Schapire, "Contextual bandits with linear payoff functions," in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 2011, pp. 208–214.
- [21] P. Klasnja, E. B. Hekler, S. Shiffman, A. Boruvka, D. Almirall, A. Tewari, and S. A. Murphy, "Micro-randomized trials: An experimental design for developing just-in-time adaptive interventions." *Health Psychology*, vol. 34, no. S, p. 1220, 2015.
- [22] I. Hochberg, G. Feraru, M. Kozdoba, S. Mannor, M. Tennenholtz, and E. Yom-Tov, "Encouraging physical activity in patients with diabetes through automatic personalized feedback via reinforcement learning improves glycemic control," *Diabetes care*, vol. 39, no. 4, pp. e59–e60, 2016.
- [23] L. Tang, R. Rosales, A. Singh, and D. Agarwal, "Automatic ad format selection via contextual bandits," in *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. ACM, 2013, pp. 1587–1594.
- [24] L. Bottou, J. Peters, J. Quiñero-Candela, D. X. Charles, D. M. Chickering, E. Portugaly, D. Ray, P. Simard, and E. Snelson, "Counterfactual reasoning and learning systems: The example of computational advertising," *The Journal of Machine Learning Research*, vol. 14, no. 1, pp. 3207–3260, 2013.
- [25] B. Awerbuch and R. D. Kleinberg, "Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches," in *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. ACM, 2004, pp. 45–53.
- [26] S. M. Shortreed, E. Laber, D. J. Lizotte, T. S. Stroup, J. Pineau, and S. A. Murphy, "Informing sequential clinical decision-making through reinforcement learning: an empirical study," *Machine learning*, vol. 84, no. 1-2, pp. 109–136, 2011.
- [27] X. Wang, Y. Wang, D. Hsu, and Y. Wang, "Exploration in interactive personalized music recommendation: a reinforcement learning approach," *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 11, no. 1, p. 7, 2014.
- [28] T. Mandel, Y.-E. Liu, S. Levine, E. Brunskill, and Z. Popovic, "Offline policy evaluation across representations with applications to educational games," in *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2014, pp. 1077–1084.
- [29] W. Shen, J. Wang, Y.-G. Jiang, and H. Zha, "Portfolio choices with orthogonal bandit learning," in *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [30] M. C. Koval, J. E. King, N. S. Pollard, and S. S. Srinivasa, "Robust trajectory selection for rearrangement planning as a multi-armed bandit problem," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 2678–2685.
- [31] J. P. Mendoza, R. Simmons, and M. Veloso, "Online learning of robot soccer free kick plans using a bandit approach," in *Twenty-Sixth International Conference on Automated Planning and Scheduling*, 2016.
- [32] D. McConachie and D. Berenson, "Bandit-based model selection for deformable object manipulation," *arXiv preprint arXiv:1703.10254*, 2017.
- [33] J. Langford and T. Zhang, "The epoch-greedy algorithm for multi-armed bandits with side information," in *Advances in neural information processing systems*, 2008, pp. 817–824.
- [34] L. Li, W. Chu, J. Langford, and R. E. Schapire, "A Contextual-Bandit approach to personalized news article recommendation," *arXiv preprint arXiv:1003.0146*, Feb. 2010.
- [35] D. J. Foster, A. Agarwal, M. Dudík, H. Luo, and R. E. Schapire, "Practical contextual bandits with regression oracles," *arXiv preprint arXiv:1803.01088*, Mar. 2018.
- [36] A. Agarwal, D. Hsu, S. Kale, J. Langford, L. Li, and R. E. Schapire, "Taming the monster: A fast and simple algorithm for contextual bandits," *arXiv preprint arXiv:1402.0555*, Feb. 2014.
- [37] T. Lattimore and C. Szepesvari, *Bandit Algorithms*. Cambridge University Press, 2019.
- [38] A. Bietti, A. Agarwal, and J. Langford, "A contextual bandit bake-off," *arXiv preprint 1802.04064*, Feb. 2018.
- [39] C. Vernade, A. Carpentier, T. Lattimore, G. Zappella, B. Ermis, and M. Brueckner, "Linear bandits with stochastic delayed feedback," *arXiv preprint arXiv:1807.02089*, July 2018.
- [40] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "The non-stochastic multiarmed bandit problem," *SIAM journal on computing*, vol. 32, no. 1, pp. 48–77, 2002.
- [41] P. Chua, T. Ilshner, and D. Caldwell, "Robotic manipulation of food products—a review," *Industrial Robot: An International Journal*, vol. 30, no. 4, pp. 345–354, 2003.
- [42] F. Erzincanlı and J. Sharp, "Meeting the need for robotic handling of food products," *Food Control*, vol. 8, no. 4, pp. 185–190, 1997.
- [43] R. Morales, F. Badesa, N. Garcia-Aracil, J. Sabater, and L. Zollo, "Soft robotic manipulation of onions and artichokes in the food industry," *Advances in Mechanical Engineering*, vol. 6, p. 345291, 2014.
- [44] P. Brett, A. Shacklock, and K. Khodabendehtloo, "Research towards generalised robotic systems for handling non-rigid products," in *International Conference on Advanced Robotics*. IEEE, 1991, pp. 1530–1533.
- [45] T. G. Williams, J. J. Rowland, and M. H. Lee, "Teaching from examples in assembly and manipulation of snack food ingredients by robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 4. IEEE, 2001, pp. 2300–2305.
- [46] C. Blanes, M. Mellado, C. Ortiz, and A. Valera, "Technologies for robot grippers in pick and place operations for fresh fruits and

- vegetables,” *Spanish Journal of Agricultural Research*, vol. 9, no. 4, pp. 1130–1141, 2011.
- [47] T. Brosnan and D.-W. Sun, “Inspection and grading of agricultural and food products by computer vision systems—a review,” *Computers and Electronics in Agriculture*, vol. 36, no. 2, pp. 193–213, 2002.
- [48] C.-J. Du and D.-W. Sun, “Learning techniques used in computer vision for food quality evaluation: a review,” *Journal of food engineering*, vol. 72, no. 1, pp. 39–55, 2006.
- [49] K. Ding and S. Gunasekaran, “Shape feature extraction and classification of food material using computer vision,” *Transactions of the ASAE*, vol. 37, no. 5, pp. 1537–1545, 1994.
- [50] W.-T. Ma, W.-X. Yan, Z. Fu, and Y.-Z. Zhao, “A chinese cooking robot for elderly and disabled people,” *Robotica*, vol. 29, no. 6, pp. 843–852, 2011.
- [51] Y. Sugiura, D. Sakamoto, A. Withana, M. Inami, and T. Igarashi, “Cooking with robots: designing a household system working in open environments,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2010, pp. 2427–2430.
- [52] M. Bollini, J. Barry, and D. Rus, “Bakebot: Baking cookies with the pr2,” in *The PR2 workshop: results, challenges and lessons learned in advancing robots with a common platform, IROS*, 2011.
- [53] M. Beetz, U. Klank, I. Kresse, A. Maldonado, L. Mösenlechner, D. Pangercic, T. Rühr, and M. Tenorth, “Robotic roommates making pancakes,” in *IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2011, pp. 529–536.
- [54] “Oreo separator machines,” <https://vimeo.com/63347829>, [Online; Retrieved on 1st February, 2018].
- [55] M. C. Gemicic 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.
- [56] D. Park, Y. K. Kim, Z. M. Erickson, and C. C. Kemp, “Towards assistive feeding with a General-Purpose mobile manipulator,” *arXiv preprint arXiv:1605.07996*, May 2016.
- [57] L. V. Herlant, “Algorithms, Implementation, and Studies on Eating with a Shared Control Robot Arm,” Ph.D. dissertation, The Robotics Institute Carnegie Mellon University, 2016.
- [58] Y. Abbasi-yadkori, D. Pál, and C. Szepesvári, “Improved algorithms for linear stochastic bandits,” in *Advances in Neural Information Processing Systems 24*, J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2011, pp. 2312–2320.
- [59] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, Oct. 2017, pp. 2999–3007.
- [60] S. Materials, 2020, [https://drive.google.com/drive/folders/132j9tdFn3ZGCKE\\_xMGgqmpOkHqtc-7?usp=sharing](https://drive.google.com/drive/folders/132j9tdFn3ZGCKE_xMGgqmpOkHqtc-7?usp=sharing).
- [61] Y. LeCun, C. Cortes, and C. Burges, “Mnist handwritten digit database,” *ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist*, vol. 2, 2010.
- [62] K. JACO, 2018, <https://www.kinovarobotics.com/en/products/robotic-armseries>.
- [63] A.-I. F.-T. Sensor, 2018, [https://www.ati-ia.com/products/ft/ft\\_models.aspx?id=Nano25](https://www.ati-ia.com/products/ft/ft_models.aspx?id=Nano25).
- [64] M. Dudik, J. Langford, and L. Li, “Doubly robust policy evaluation and learning,” *arXiv preprint arXiv:1103.4601*, Mar. 2011.
- [65] M. Dimakopoulou, N. Vlassis, and T. Jebara, “Marginal posterior sampling for slate bandits,” in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2019, pp. 2223–2229.
- [66] T. Bhattacharjee, E. Gordon, R. Scalise, M. Cabrera, A. Caspi, M. Cakmak, and S. Srinivasa, “Is more autonomy always better? exploring preferences of users with mobility impairments in robot-assisted feeding,” in *ACM/IEEE International Conference on Human-Robot Interaction*, 2020.