

CONFORMALIZED INTERACTIVE IMITATION LEARNING: HANDLING EXPERT SHIFT & INTERMITTENT FEEDBACK

Michelle Zhao^a, Reid Simmons^a, Henny Admoni^a, Aaditya Ramdas^{†b}, Andrea Bajcsy^{†a}

^a Robotics Institute, School of Computer Science, Carnegie Mellon University

^b Departments of Statistics and Machine Learning, Carnegie Mellon University

{mzhao2, hadmoni, rsimmons, aramdas, abajcsy}@andrew.cmu.edu

ABSTRACT

In interactive imitation learning (IL), uncertainty quantification offers a way for the learner (i.e. robot) to contend with distribution shifts encountered during deployment by actively seeking additional feedback from an expert (i.e. human) online. Prior works use mechanisms like ensemble disagreement or Monte Carlo dropout to quantify when black-box IL policies are uncertain; however, these approaches can lead to overconfident estimates when faced with deployment-time distribution shifts. Instead, we contend that we need uncertainty quantification algorithms that can leverage the expert human feedback received *during deployment time* to adapt the robot’s uncertainty *online*. To tackle this, we draw upon online conformal prediction, a distribution-free method for constructing prediction intervals online given a stream of ground-truth labels. Human labels, however, are intermittent in the interactive IL setting. Thus, from the conformal prediction side, we introduce a novel uncertainty quantification algorithm called intermittent quantile tracking (IQT) that leverages a probabilistic model of intermittent labels, maintains asymptotic coverage guarantees, and empirically achieves desired coverage levels. From the interactive IL side, we develop ConformalDagger, a new approach wherein the robot uses prediction intervals calibrated by IQT as a reliable measure of deployment-time uncertainty to actively query for more expert feedback. We compare ConformalDagger to prior uncertainty-aware DAgger methods in scenarios where the distribution shift is (and isn’t) present because of changes in the expert’s policy. We find that in simulated and hardware deployments on a 7DOF robotic manipulator, ConformalDagger detects high uncertainty when the expert shifts and increases the number of interventions compared to baselines, allowing the robot to more quickly learn the new behavior. Project page at cmu-intentlab.github.io/conformalized-interactive-il/.

1 INTRODUCTION

End-to-end robot policies trained via imitation learning (IL) have proven to be an extremely powerful way to learn complex robot behaviors from expert human demonstrations (Price & Boutilier, 2003; Schaal, 1996; Argall et al., 2009; Levine et al., 2016; Jang et al., 2022; Chi et al., 2023; Kim et al., 2024). At the same time, distribution shift is a core challenge in this domain, hampering the reliability of deploying such robot policies in the real world (Chang et al., 2021).

One way to combat this is via uncertainty quantification. By training an ensemble of policies (Menda et al., 2019) or via monte-carlo dropout during training (Cui et al., 2019), the robot learner can detect uncertain states and *actively elicit* additional action labels from the human expert online via an interactive IL framework (such DAgger (Ross et al., 2011)). At their core, these prior uncertainty-aware IL methods look to the training demonstration data as a proxy for deployment-time uncertainty. Any human expert labels requested at deployment time using this uncertainty estimate are simply stored for later re-training; the uncertainty estimate itself is not adapted online to the expert data nor does it inform any subsequent queries during the deployment episode.

Instead, we contend that the human feedback requested and received *during* deployment time is a valuable uncertainty quantification signal that should be leveraged to *update* the robot’s uncertainty estimate *online*. If properly accounted for, the updated uncertainty estimate will influence when

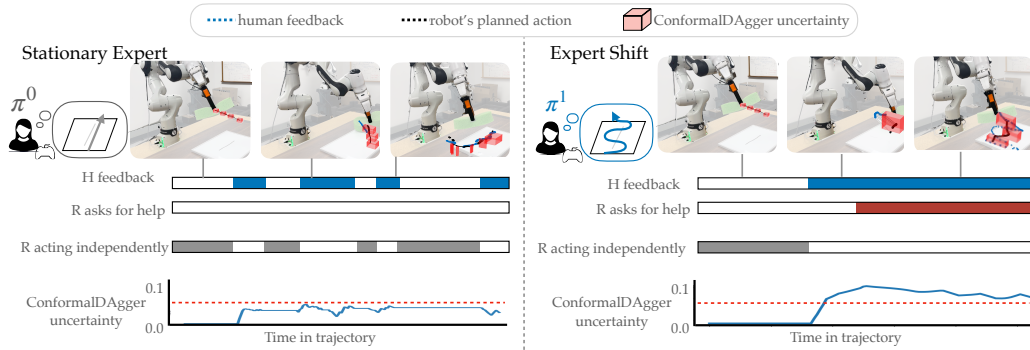


Figure 1: **Conformalized Interactive Imitation Learning.** The robot learns an initial policy via imitation learning to wipe a line drawing by following a straight path. When the initial policy is deployed, ConformalDagger calibrates the robot’s uncertainty (represented by red prediction interval boxes) based on feedback received from the human during the interactive IL loop. (left) The robot refrains from asking questions when its uncertainty is low because the expert policy at deployment time is aligned to the demonstrations on which the learner was trained. (right) When the human shifts in their task strategy from the training data distribution, the robot uncertainty increases and enables it to actively query for more labels.

the robot asks for more help, enabling it to targetedly probe the human expert to improve policy performance. The challenge is how to do this online uncertainty update in the presence of the end-to-end “black box” policies underlying modern imitation learning.

To tackle this, we take inspiration from online conformal prediction (Gibbs & Candes, 2021) which is a distribution-free way to represent uncertainty via prediction sets constructed on the output of a black-box model. Our first contribution is extending online conformal prediction to the case where labels are observed intermittently, as is the case in interactive IL with an expert. Specifically, we instantiate Intermittent Quantile Tracking (IQT), an algorithm which adjusts prediction intervals online to ensure that the true label lies within the predicted interval with high-probability despite probabilistic access to labels. On standard conformal time series datasets, we empirically find that IQT achieves empirical coverage close to the desired level by boosting the size of the calibrated intervals based on the likelihood of observing feedback.

With our intermittent conformal algorithm in hand, we develop ConformalDagger, a new interactive IL approach wherein the robot uses prediction intervals calibrated by IQT as a reliable measure of deployment-time uncertainty to actively query for more expert feedback. We instantiate ConformalDagger in a simulated 4D robot goal-reaching task and in hardware on a 7 degree-of-freedom robotic manipulator that uses a state-of-the-art Diffusion Policy (Chi et al., 2023) learned via IL to perform a sponging task (Figure 1). We study how ConformalDagger compares to prior uncertainty-aware DAgger methods when the deployment-time and training-time datasets are from the same distribution (left, Figure 1) as well as a shifted one (right, Figure 1). Specifically, we instantiate a potential source of distribution shift as *expert policy shift*: the training-time expert wipes a line-drawing with a straight-line path while the deployment-time expert refines their strategy to be a zigzag path (Figure 1). We find that ConformalDagger automatically increases uncertainty online when the expert shifts, resulting in more expert labels queries compared to EnsembleDagger and allowing our approach to rapidly learn a policy aligned with the expert’s intentions.

2 RELATED WORK

Interactive Imitation Learning (IL) with Online Experts. Interactive IL is a branch of imitation learning wherein a robot learner can query a (human) expert to receive additional labels either during or after task execution (Celemin et al., 2022). A foundational approach for interactive IL is DAgger (Dataset Aggregation) (Ross et al., 2011), which iteratively augments the training dataset by aggregating data from the expert and learner policies and assuming the expert is stationary. In the online case, the robot learner can cede control to the expert at any time to get additional state-action data (also known as *robot-gated* feedback) or the expert can actively intervene at any time (also known as *human-gated* feedback) (Kelly et al., 2019). From the learner’s perspective, a key question is when to request feedback from the expert so that it minimizes expert effort but also minimizes

negative events caused by an erroneous policy (e.g., running into a wall). On one hand, prior works focus on minimizing human effort by constraining robot requests via a limited human attention model (Hoque et al., 2023) or a budget of human interventions (Hoque et al., 2021a;b). On the other hand, prior works prioritize deployment-time safety by classifying safe versus unsafe states (SafeDagger (Zhang & Cho, 2017), Replay Estimation (Swamy et al., 2022)), estimating uncertainty via ensemble disagreement (EnsembleDagger (Menda et al., 2019)), or Monte Carlo dropout (Cui et al., 2019). We present ConformalDagger, which uses our novel uncertainty quantification method grounded in conformal prediction to adaptively increase the learner requests for help under uncertainty and decrease the number of requests when confident.

Interacting with Non-stationary Experts. A core kind of distribution shift we study in this work is human expert distribution shift. This can occur for a variety of reasons, from human teleoperators having biases when teaching robots (Thomaz et al., 2006; Thomaz & Breazeal, 2008), to generating different data distributions because of suboptimality (Grüne-Yanoff, 2015; Thompson, 1999) or varied risk tolerance (Kwon et al., 2020). Moreover, simply observing robot behavior can influence how humans react or teach robots over time (Hong et al., 2024; Sagheb et al., 2023; Xie et al., 2021). Despite its prevalence in the real world, the majority of interactive IL works assume a stationary expert policy (Likmeta et al., 2021; Shin et al., 2023; Zheng et al., 2022). Our conformalized imitation learning approach takes a key step towards closing this gap by developing an online conformal prediction algorithm that can account for expert distribution shift.

Online Conformal Prediction. Conformal prediction is a distribution-free uncertainty quantification method for constructing prediction intervals for both classification and regression problems (Angelopoulos & Bates, 2023; Romano et al., 2019; 2020; Zaffran et al., 2022), as well as for offline or online data. We focus on the *online* setting (e.g., timeseries) where uncertainty quantification is performed on streaming pairs of input-label data that are not necessarily i.i.d. (Gibbs & Candes, 2021). Broadly speaking, there are two predominant algorithms in this setting: adaptive conformal inference (ACI) Gibbs & Candes (2021); Gibbs & Candès (2024); Bhatnagar et al. (2023); Zaffran et al. (2022) and quantile tracking (QT) Angelopoulos et al. (2024a). Both are online gradient descent-based methods which guarantee asymptotic coverage in the online setting. Unlike ACI which is prone to infinitely sized intervals after a series of miscoverage events, QT directly estimates the value of the empirical quantile itself, ensuring coverage with finite intervals. In this work, we relax the assumption that labels must be observed at each time point in the streaming data and extend the online conformal paradigm to ensure coverage in the intermittent label regime.

Conformal Prediction for Robotics. Recently, conformal prediction has become popular in the robotics domain in part due to the distribution-free guarantees it provides for arbitrarily complex learned models present within modern robotics pipelines. Specifically, conformal prediction has been used to provide collision-avoidance assurances (Chen et al., 2021; Lindemann et al., 2023; Dixit et al., 2023; Muthali et al., 2023; Taufiq et al., 2022; Dietterich & Hostetler, 2022; Lin & Bansal, 2024), calibrate early warning systems (Luo et al., 2022), and quantify uncertainty in large language model based planners (Ren et al., 2023; Lidard et al., 2024). There are several core challenges with the input-and-label data encountered in robotics: data is non-i.i.d. (e.g., sequential decision-making), data distributions are non-stationary (e.g. changing environment conditions), and labels are intermittently observed (e.g., limited expert feedback in the IL domain). By extending online conformal prediction to the intermittent label setting, we take a step towards addressing these challenges.

3 ONLINE CONFORMAL PREDICTION WITH INTERMITTENT LABELS

From the uncertainty quantification side, our core technical contribution is extending online conformal prediction to settings where ground truth labels are intermittently observed. We present our algorithm in the context of quantile tracking for brevity and relevance to our interactive IL experiments. However, we also derive an extension of ACI to intermittent labels in the Appendix Section B.

Setting. We focus on online conformal prediction in the adversarial setting, such as time-series forecasting. This considers an *arbitrary* sequence of data




Figure 2: **Graphical Model of IQT.** Intermittent quantile tracking introduces random variable obs_t , which represent the receiving of ground truth observations at timestep t . obs_t is distributed according to p_t , which may depend on x_t or history through $t - 1$.

points $(x_t, y_t) \in \mathcal{X} \times \mathcal{Y}$, for $t = 1, 2, \dots$, that are not necessarily I.I.D. Our goal is to produce prediction sets on the output of any base prediction model such that the sets contain the true label with a specified coverage probability. Mathematically, at each time t , we observe x_t and seek to cover the true label y_t with a set $C_t(x_t)$, which depends on a base prediction model, $\hat{f} : \mathcal{X} \rightarrow \mathcal{Y}$. The base model takes as input the current x_t and outputs prediction \hat{y}_t ; in the non-intermittent case, we observe the ground-truth label y_t after each prediction.

Background: Quantile Tracking. The quantile tracking (QT) algorithm Angelopoulos et al. (2024b) implicitly seeks to track the value of the $1 - \alpha$ quantile via online gradient descent on the quantile loss (Koenker & Bassett Jr, 1978). As in Angelopoulos et al. (2024b), we leverage a bounded nonconformity score function: $s : \mathcal{X} \times \mathcal{Y} \rightarrow [0, B]$ where $0 < B < \infty$, to quantify the error made by the initial prediction of the base model. We assume the nonconformity score function $s(x_t, y_t)$ is negatively oriented (lower values indicate less nonconformity or greater prediction accuracy). Let q_t represents the estimated $1 - \alpha$ quantile of the score sequence $s_t, t \in \mathbb{N}$. Prediction intervals are constructed using the nonconformity score function:

$$C_t(x_t) = \{y \in \mathcal{Y} : s(x_t, y_t) \leq q_t\} \quad (1)$$

To expand or contract the prediction intervals, the level q_t is adjusted via the online update:

$$q_{t+1} = q_t + \gamma_t(\text{err}_t - \alpha), \quad (2)$$

where $\gamma_t > 0$ is a time-varying step size and $\text{err}_t = \mathbb{1}_{y_t \notin C_t(x_t)}$ is the empirical miscoverage at t .

Intuitively, this update increases the quantile threshold when the model continuously miscovers and decreases the quantile when the model coverage is performant. For arbitrary step size γ_t with no assumptions on the sequence of data points $(x_1, y_1), (x_2, y_2), \dots$ and an initial quantile threshold $q_1 \in [0, B]$, the quantile tracking update satisfies Equation 3 (Theorem 2 of Angelopoulos et al. (2024b)):

$$\left| \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{y_t \in C_t(x_t)} - (1 - \alpha) \right| \leq \frac{B + \max_{1 \leq t \leq T} \gamma_t}{T} \cdot \|\Delta_{1:T}\|_1 \quad (3)$$

where Δ is defined $\Delta_1 = \gamma_1^{-1}$ and $\Delta_t = \gamma_t^{-1} - \gamma_{t-1}^{-1}$ for $t \geq 2$, and $\Delta_{1:T} = (\Delta_1, \dots, \Delta_T)$. We can see that $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T 1 - \text{err}_t$ approaches $1 - \alpha$. This guarantees quantile tracking gives the $1 - \alpha$ long-term empirical coverage frequency.

Intermittent Quantile Tracking (IQT). Our paradigm lifts the assumption that the ground truth label y_t is observed constantly, but rather with some known probability at each timestep. Let the binary random variable $\text{obs}_t \in \{0, 1\}$ represents whether the robot observes label y_t at timestep t :

$$\text{obs}_t := \begin{cases} 1, & \text{if } y_t \text{ observed} \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

We introduce a probabilistic observation model, $p_t := \mathbb{P}(\text{obs}_t = 1 | x_t)$, where p_t represents the probability of receiving feedback, which can be dependent on input x_t (and past history, this possibility is demonstrated by the dotted line in Figure 2). A key observation under the paradigm of intermittent labels is that err_t may not at every timestep be accessible to the algorithm if y_t is not provided, but the value of err_t exists at every timestep even if inaccessible.

Our quantile tracking update under probabilistic observations, which we call Intermittent Quantile Tracking (IQT), is defined (Equation 5):

$$q_{t+1} = q_t + \frac{\gamma_t}{p_t}(\text{err}_t - \alpha)\text{obs}_t \quad (5)$$

Proposition 1. *Let $(X_1, Y_1), (X_2, Y_2), \dots$ be an arbitrary sequence of data points, and let $s : \mathcal{X} \times \mathcal{Y} \rightarrow [0, B]$. Let γ_t be an arbitrary positive sequence, and fix an initial threshold $q_1 \in [0, B]$. Then Intermittent Quantile Tracking (IQT) satisfies for all $T \geq 1$:*

$$\left| \frac{1}{T} \sum_{t=1}^T \text{err}_t - \alpha \right| \leq \frac{B + \max_{1 \leq t \leq T} \frac{\gamma_t}{p_t}}{T} \cdot \|\Delta_{1:T}\|_1 \quad (6)$$




Figure 3: **IQT on Amazon Stocks with AR base model: Coverage & Interval Visualization.** We set $p_t = 0.1, \forall t$ to simulate seeing the true price only 10% of the time and show the prediction interval for 1 seed and coverage averaged over 5 seeds (shaded error on coverage plots show standard deviation). **IQT-pd** is prone to large intervals when lr is large, and boosts coverage over **IQT-pi** when lr is small.

Refer to full proof in Appendix Section A. Note that when feedback is constant (i.e. $p_t = 1$), $\text{obs}_t = 1$ for all timesteps, and IQT reduces to online quantile tracking with arbitrary step sizes.

Practical Note: Choosing Gamma. In practice, an important decision is the choice of γ_t . Prior works in quantile tracking with constant labels (Angelopoulos et al., 2024a) choose $\gamma_t = lr\hat{B}_t$, where $\hat{B}_t := \max_{t-k \leq r \leq t-1} s(x_r, y_r)$, k represents a look-back window across the previous timesteps with observed ground truth labels, and lr is a small constant. On our domains, we find that choosing γ_t such that it contains \hat{B}_t improves empirical coverage over the choice of a constant $\gamma_t = \gamma$.

3.1 EXPERIMENTS: INTERMITTENT QUANTILE TRACKING ON STANDARD TIMESERIES DATA

We first empirically evaluate our approach for intermittent conformal quantile tracking on time series benchmarks used in the online conformal prediction literature. The goal of testing on standard conformal datasets is to evaluate (1) how different choices of step size γ_t affect the empirical coverage of IQT, and (2) validate that under probabilistically intermittent observations, IQT maintains coverage close to the desired level. We use the findings from these experiments to inform how we leverage intermittent quantile tracking in the interactive imitation learning domain.

Setup. We test on three benchmark datasets from Angelopoulos et al. (2024a): (1) Amazon stock prices, (2) Google stock prices (Nguyen, 2018), and the (3) Elec2 dataset (Harries, 1999). We test four base prediction models, \hat{f} , all trained via darts (Herzen et al., 2022) to see if IQT consistently maintains coverage close to the desired level. We present the Autoregressive (**AR**) model with 3 lags for brevity in the main text and defer the other model results to the Appendix. Our nonconformity score is the asymmetric (signed) residual score. We measure (1) marginal coverage over the time series, (2) longest miscoverage error sequence, and (3) mean prediction interval size.

Observation Models. We test three different observation frequencies: (1) *infrequent* $p_t = 0.1, \forall t$, (2) *partial* $p_t = 0.5, \forall t$, and (3) *frequent* $p_t = 0.9, \forall t$. We focus in this section on the results for *infrequent* observations, and discuss further results in the Appendix Section C.

Methods. We compare two variants of our IQT algorithm by controlling if time-varying step-size γ_t in Equation 5 depends or does not depend on p_t . For each variant, we test $lr \in [1, 0.1, 0.01]$. **IQT with a p_t -independent update (IQT-pi)** uses $\gamma_t = lr\hat{B}_t p_t$. The quantile tracking update reduces to $q_{t+1} = q_t + lr\hat{B}_t(\text{err}_t - \alpha)\text{obs}_t$. Since the p_t term cancels out in the update, we refer to this model as the *p-independent* update. **IQT with a p_t -dependent update (IQT-pd)** uses $\gamma_t = lr\hat{B}_t$. Since the p_t term remains in the update $q_{t+1} = q_t + \frac{lr\hat{B}_t}{p_t}(\text{err}_t - \alpha)\text{obs}_t$, we refer to this model as the *p-dependent* update. For both variants, we set our desired coverage to $1 - \alpha = 0.9$ for all experiments. We set lookback window $k = 100$ timesteps for the Amazon stock price data.

Results. We center our discussion on the Amazon stock price results for **IQT-pd** and **IQT-pi** with AR base model under ground truth labels observed with $p_t = 0.1$ frequency, as the intermittent setting is our focus. See Appendix Section C for further results on other base models and datasets. Inversely scaling the quantile tracking step size by p_t causes **IQT-pd** to construct larger intervals than **IQT-pi**. Figure 3 shows the prediction coverage as a moving average (window=50) for 5 random

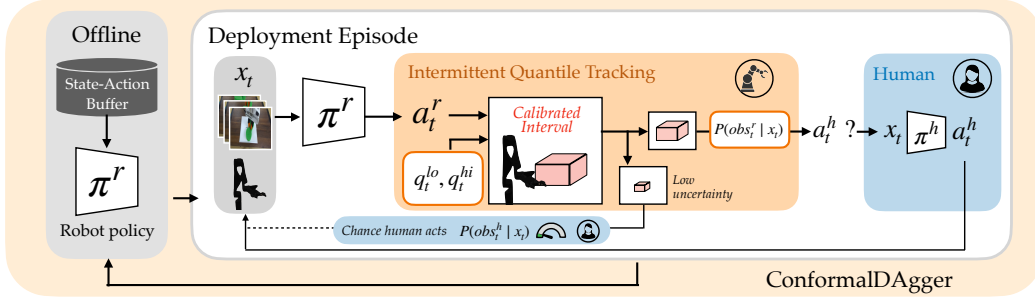


Figure 4: **ConformalDagger Framework.** After obtaining an initial learner policy π^r (left), ConformalDagger calibrates uncertainty *during* the interactive deployment episode with the expert via intermittent quantile tracking (center). When the size of the uncertainty intervals is high, the robot actively queries the user for feedback. When uncertainty is low, the robot executes its predicted action and the human may independently intervene with some low probability. After deployment episode ends, the data is aggregated and the learner retrained (arrow from right to left).

seeds and interval sizes for one seed. When lr is high (1.0), **IQT-pi** and **IQT-pd** achieve comparable coverage, but **IQT-pd** is prone to constructing much larger prediction intervals than **IQT-pi**. Smaller and mid-sized learning rates ($lr = 0.01, 0.1$) regulate the size of **IQT-pd** intervals, leading to tight prediction intervals which maintain desired coverage levels.

4 CONFORMALDAGGER: A CALIBRATED APPROACH TO ASKING FOR EXPERT FEEDBACK

Since intermittent quantile tracking enables us to rigorously quantify the learner’s uncertainty despite the fact that the labels are only revealed intermittently (i.e., when the expert intervenes), we can develop ConformalDagger: a new way for the robot to tradeoff between acting autonomously and strategically asking for help when uncertainty increases.

Setup. ConformalDagger treats the robot’s policy as the base model $\hat{f} := \pi^r$ on which we perform intermittent quantile tracking. The initial novice policy $\pi_0^r : \mathcal{X} \rightarrow \mathcal{A}$ is trained on initial demonstration data \mathcal{D}_0 of task \mathcal{T} performed by the expert π_0^h . Here, \mathcal{X} represents the policy’s inputs (e.g., image observations, proprioception), \mathcal{Y} are the labels representing the robot’s actions (e.g., future end-effector positions). During deployment, the robot policy generates a sequence of input-predicted action pairs $(x_t, a_t^r) \in \mathcal{X} \times \mathcal{Y}$ for $t = 1, 2, \dots$, that are temporally correlated. Relatedly, for each input x_t the learner observes, there is a corresponding expert action $(x_t, a_t^h) \in \mathcal{X} \times \mathcal{Y}$ that is the ground-truth label we seek to cover via our IQT prediction intervals $C_t(x_t)$.

Observation Model. A key component of IQT is the observation likelihood model, $p_t = \mathbb{P}(\text{obs}_t = 1 \mid x_t)$. A nice byproduct of this model is that we can naturally derive a feedback model that is simultaneously *human-* and *robot-*gated by decomposing the observation model into the combination of both gating functions. Intuitively, the likelihood of observing the expert feedback at t is given by the probability that the human chooses to give feedback *or* the robot asks the expert for feedback. Let $\text{obs}_t^h \in \{0, 1\}$ be a random variable representing observing a human-gated feedback (if $\text{obs}_t^h = 1$) where the expert initiates providing an action label $a_t^h = \pi^h(x_t)$ for x_t . Let $\text{obs}_t^r \in \{0, 1\}$ be a random variable representing robot-gated feedback, where the robot asks the expert for a label if $\text{obs}_t^r = 1$. We assume that if the robot elects to ask a question, the human will respond with probability 1. Our observation model takes the form:

$$p_t := \mathbb{P}(\text{obs}_t = 1 \mid x_t) = \mathbb{P}(\text{obs}_t^h = 1 \vee \text{obs}_t^r = 1 \mid x_t) \quad (7)$$

$$= \mathbb{P}(\text{obs}_t^h = 1 \mid x_t) + \mathbb{P}(\text{obs}_t^r = 1 \mid x_t) - \mathbb{P}(\text{obs}_t^h = 1 \mid x_t)\mathbb{P}(\text{obs}_t^r = 1 \mid x_t) \quad (8)$$

While our approach is not prescriptive about these models, in our experiments we model $\mathbb{P}(\text{obs}_t^h = 1 \mid x_t) = c$ as a small constant to represent infrequent human interventions. We describe the robot-gated likelihood model $\mathbb{P}(\text{obs}_t^r = 1 \mid x_t)$ below, informed by our uncertainty estimates.

Algorithm 1 ConformalDagger (*changes from Dagger (Ross et al., 2011) highlighted*)

-
- 1: Collect initial demonstration data \mathcal{D}_0 from expert π_0^h and train initial learner policy π_0^r .
 - 2: **for** interactive deployment episode $i = 0 : M$ **do**
 - 3: Initialize q_0^{lo}, q_0^{hi} .
 - 4: **for** deployment timestep $t = 1 : H$ **do**
 - 5: Get predicted action label: $a_t^r \leftarrow \pi_i^r(x_t)$
 - 6: Construct calibrated uncertainty interval: $C_t(x_t) = [a_t^r - q_t^{lo}, a_t^r + q_t^{hi}]$.
 - 7: **if** robot queries (sample $\mathbb{P}(\text{obs}_t^r | o_t)$) or human intervenes (sample $\mathbb{P}(\text{obs}_t^h | o_t)$) **then**
 - 8: IQT update: $q_{t+1}^{lo} \leftarrow q_t^{lo} + \frac{\gamma_t}{p_t}(\text{err}_t - \alpha)\text{obs}_t$ with expert action $a_t^h \leftarrow \pi_i^h(x_t)$.
 - 9: $q_{t+1}^{hi} \leftarrow q_t^{hi} + \frac{\gamma_t}{p_t}(\text{err}_t - \alpha)\text{obs}_t$ with expert action $a_t^h \leftarrow \pi_i^h(x_t)$.
 - 10: **else**
 - 11: IQT update: $q_{t+1}^{lo} \leftarrow q_t^{lo}$.
 - 12: $q_{t+1}^{hi} \leftarrow q_t^{hi}$.
 - 13: Aggregated dataset with observed state and expert action pairs: $\mathcal{D}_{i+1} \leftarrow \mathcal{D}_i \cup \{(x, \pi_i^h(x))\}$
 - 14: Retrain learner: $\pi_{i+1}^r \leftarrow \arg \min_{\pi} \mathcal{L}(\mathcal{D}_{i+1})$
-

Quantifying Uncertainty: IQT. In the interactive IL setting, IQT begins with initial upper and lower quantiles in the action space, $q_0^{lo}, q_0^{hi} \in \mathcal{A}$. The nonconformity score is a residual on the predicted (a^r) versus expert (a^h) action. Let $s_t^{lo}(a_t^r, a_t^h) = a_t^r - a_t^h$ be the lower residual (referred to as s_t^{lo} for brevity), and $s_t^{hi}(a_t^r, a_t^h) = a_t^h - a_t^r$ be the upper residual (referred to as s_t^{hi}). If the expert action a_t^h is observed, we compute the upper and lower misscoverage of q_t^{lo}, q_t^{hi} as two indicator vectors: $\text{err}_t^{lo} = (s_t^{lo} < q_t^{lo})$ and $\text{err}_t^{hi} = (s_t^{hi} < q_t^{hi})$. IQT then updates the quantile estimates online to obtain $q_{t+1}^{lo}, q_{t+1}^{hi}$ via the update rule from equation 5. At the next timestep, the adjusted prediction interval is constructed with $C_{t+1}(x_{t+1}) = [a_{t+1}^r - q_{t+1}^{lo}, a_{t+1}^r + q_{t+1}^{hi}]$. If the expert action is not observed, $\text{obs}_t = 0$, then $q_{t+1}^{lo} = q_t^{lo}$ and $q_{t+1}^{hi} = q_t^{hi}$ and the prediction interval size remains the same. Note that although err_t is not known in the case where the expert does not provide an action label, IQT does not require it; IQT simply makes no change to the quantile estimate. In our simulated and hardware experiments, our action space, q_0^{lo} , and q_0^{hi} , are vector-valued. Our experiments instantiate IQT for continuous vectors, but the approach extends to discrete-valued action spaces.

Leveraging Uncertainty: Asking for Help. Finally, the calibrated intervals $C_{t+1}(x_{t+1})$ constructed by IQT enable us to design a new robot-gated feedback mechanism. Specifically, the model $\mathbb{P}(\text{obs}_t^r | x_t)$ is informed by the calibrated interval size, $u(x_{t+1}) := \|C_{t+1}(x_{t+1})\|_2$. In our simulated experiments, we use $\mathbb{P}(\text{obs}_t^r | x_t) = \sigma(\beta[u(x_t; \pi^r) - \tau])$ where τ is some uncertainty threshold, β is a temperature hyperparameter, and σ is the sigmoid function. In our hardware experiments, we take τ to be a hard threshold above which $\mathbb{P}(\text{obs}_t^r | x_t) = 1$, and below which $\mathbb{P}(\text{obs}_t^r | x_t) = 0$.

ConformalDagger Algorithm. We summarize ConformalDagger in Algorithm 1, highlighting the difference from traditional Dagger (Ross et al., 2011). The learner interacts with the intermittent expert during M interactive episodes. At the start of each, q^{lo}, q^{hi} and are reset. Each episode lasts H time steps and the intermittently-observed expert state-action pairs (x_t, a_t^h) are aggregated with the prior training data in a fixed size training buffer to form the updated dataset, \mathcal{D}_{i+1} , which enables us to retrain the learner π_{i+1}^r for the next deployment episode.

Two features of our setting make it difficult to prove we maintain Dagger performance guarantees: intermittent expert feedback even with a stationary expert, and the possibility of a dynamic expert even with feedback at every state along learner rollouts. Similar to approaches like (Kelly et al., 2019) and Menda et al. (2019), our approach helps to contend with these constraints that reflect real-world challenges in constrained expert capacity and potential shift.

5 SIMULATED INTERACTIVE IMITATION LEARNING EXPERIMENTS

To evaluate ConformalDagger, we run a series of simulated experiments with access to an oracle expert. We ground our experiments in a simulated robot goal-reaching task (left, Figure 5) and study scenarios where the distribution shift occurs due to the expert’s changing preference in goal location.

5.1 EXPERIMENTAL SETUP

Task & Initial Learner Policy. The robot learns a neural network policy $\pi^r : \mathcal{X} \rightarrow \mathcal{A}$ to move a cup from a start to the expert’s desired goal location. We model input $x \in \mathcal{X} \subseteq \mathbb{R}^{12}$ as the xyz position of the robot as well as its binary gripper state (open or closed) across the previous 3 timesteps. The labels are actions, $a \in \mathcal{A} \subseteq \mathbb{R}^4$, represented as next xyz position and gripper state. The robot always starts at an initial x with the gripper state closed (holding a cup) and must keep holding the cup while moving it to the unknown goal location that the expert prefers; let $\{g_0, g_1\}$ be two such goals. Before the first interactive deployment ($i = 0$), we simulate the expert as initially giving demonstrations placing the cup at g_0 , where $\pi_{i=0}^h(x) = x + \omega \frac{g_0 - x}{\max_d(|g_0 - x|)}$, where $\max_d(\cdot)$ is a maximum over dimensions. The step size is regulated to be at most $\omega = 0.01$. The initial robot policy $\pi_{i=0}^r$ is trained on a dataset \mathcal{D}_0 of 10 expert trajectories with synthetically injected noise drawn from $\mathcal{N}(1, 0.5)$ for robustness as in (Laskey et al., 2017). See Appendix D for policy implementation details.

Simulated Expert Policies. To induce controlled distribution shift, we study three expert policies (left, column of Figure 5). (1) *Stationary*: the expert has a fixed goal, g_0 , across all deployment episodes. (2) *Shift*: the expert goal shifts from g_0 to g_1 at deployment episode $i = 5$. For example, the expert may have decided that a different cup location is easier to reach. (3) *Drift*: expert’s goal slowly drifts from g_0 to g_1 over the course of deployment episodes (in Figure 5 the drift from $g_0 \rightarrow g_{1a}$ occurs at episode $i = 5$, $g_{1a} \rightarrow g_{1b}$ occurs at episode $i = 8$, and $g_{1b} \rightarrow g_1$ occurs at episode $i = 11$). For example, the expert may start with a conservative goal location initially (e.g., a goal nearby) and incrementally move the cup closer and closer to their target goal that may be further out of reach.

Interactive Deployment Episodes & Learner Re-training. We consider $M = 15$ deployment episodes before re-training the learner. Each deployment episode has two interactive task executions. The task ends when the cup has reached the correct goal position, g^* , or when the maximum timesteps (100) have been reached. The expert answers queries with optimal actions under their current policy $a^h = \pi_i^h(x)$. Following DAgger (Ross et al., 2011), after each deployment episode, the state-action pairs where the expert provided action labels are aggregated into the training dataset, forming aggregated buffer \mathcal{D}_{i+1} for the next deployment episode $i + 1$. We constrain the size of the replay buffer to 300 datapoints, dropping the old experiences.

Methods. We compare **ConformalDagger** to **EnsembleDagger** (Menda et al., 2019). ConformalDagger uses an uncertainty threshold $\tau = 0.09$, temperature $\beta = 100$, lookback window $k = 100$, lr = 0.6, and initial $q_0^{lo,hi} = 0.01$. The uncertainty threshold τ is heuristically tuned to ask few, but infrequent questions in the first interactive deployment episode. EnsembleDagger queries an expert online when there is high action prediction variance across an ensemble of learner policies, and when a safety classifier detects dissimilarity between expert and robot actions. We use 3 ensemble members and an uncertainty threshold of $\tau = 0.06$ for the ensemble disagreement (selected in a similar manner to ConformalDagger) and a safety classifier threshold of 0.03.

Metrics. We measure the quality of our uncertainty quantification via the **miscoverage rate** and human effort via the **intervention percentage** of the deployment trajectory. We compute the miscoverage for EnsembleDagger using three times the standard deviation as the prediction interval. We measure the quality of the learned policy via two metrics. **Decision deviation** simulates the learner’s behavior under its current policy and queries the expert at each learner state to obtain an expert action label. We measure the average L2 distance between the predicted and expert action. **Trajectory deviation** forward simulates both the expert and the learner acting independently under their policies, starting from the same initial x . We compare the L2 distance between the trajectories.

5.2 EXPERIMENTAL RESULTS

We focus results on infrequent human-gated feedback where $\mathbb{P}(\text{obs}_t^h = 1 \mid x_t) = 0.2, \forall t$. In Appendix D, we present experiments with frequent (=0.9) and partial (=0.5) feedback. Irrespective of the human-gated likelihood, the learner can always *actively* ask for feedback based on $\mathbb{P}(\text{obs}_t^r \mid x_t)$.

Takeaway 1: When the expert shifts, ConformalDagger asks for more help immediately compared to EnsembleDagger, enabling the algorithm to more quickly align to the expert. Consider deployment episode $i = 5$ where the expert shifts from goal 0 to goal 1 (center row, Figure 5). ConformalDagger immediately increases the number of expert feedback requests: before shift, the expert intervened $\sim 20\%$ of the time but in the 5th episode they intervene $\sim 60\%$ of the time. In

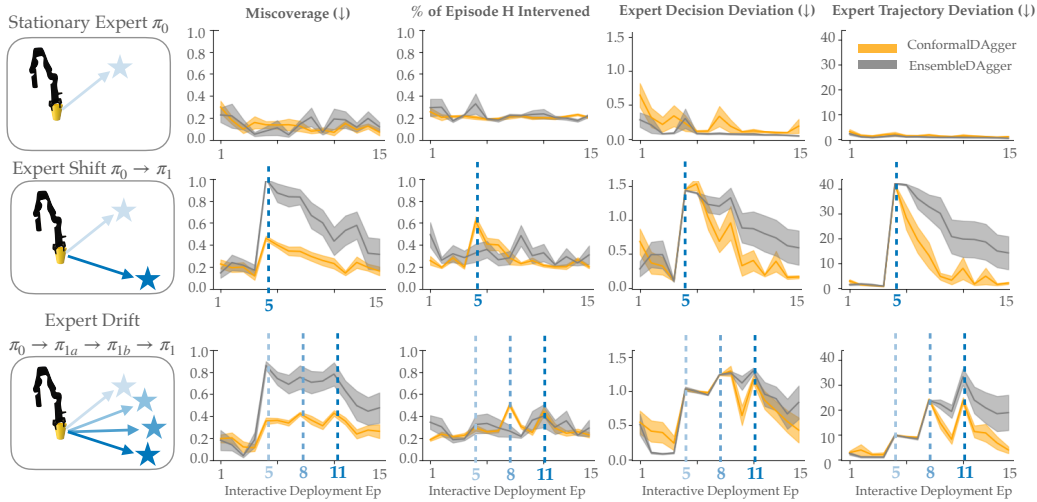


Figure 5: **Simulated Robot Results.** When the expert shifts (middle row) or drifts (bottom), **ConformalDagger** increases the number of requests for expert intervention compared to **EnsembleDagger**, decreasing miscoverage and minimizing expert deviation after re-training. With a stationary expert (top row), both algorithms are similar. Shading represents std. error across 10 seeds.

contrast, EnsembleDagger remains close to $\sim 30\%$. Relatedly, ConformalDagger has consistently lower miscoverage rate (max = 0.4 at shift; converges to 0.2) compared to the baseline (max = 1.0 at shift, converges to 0.4). Due to the extra solicited feedback, ConformalDagger’s ultimate decision and trajectory deviations are minimized in the subsequent retrained policies.

Takeaway 2: ConformalDagger automatically asks for more help each time the expert drifts. In the bottom row of Figure 5, we see ConformalDagger maintain a similar level of queries as EnsembleDagger during the first shift (at episode 5) and last shift (episode 11) but increases feedback during the intermediate shift (at episode 8). Despite these similarities, EnsembleDagger’s miscoverage rate is consistently higher than ConformalDagger’s and EnsembleDagger’s re-trained policy on average does not adapt as quickly (with higher expert decision and trajectory deviation). We hypothesize that this is because ConformalDagger asks frequent questions consistently across all seeds compared to EnsembleDagger (i.e. lower variance in Fig 5).

Takeaway 3: With a stationary expert, ConformalDagger and EnsembleDagger are similar. We find that without distribution shift (top row, Figure 5), both methods ask for minimal help (staying near the 20% human-gated probability), and achieve similar performance in alignment to the expert.

6 HARDWARE EXPERIMENTS

Finally, we deployed **ConformalDagger** in hardware on a 7 degree-of-freedom robotic manipulator that uses a state-of-the-art Diffusion Policy (Chi et al., 2023) trained via IL to perform a sponging task (Figure 1). The goal of our hardware experiments is to demonstrate how our approach can scale to a high-dimensional, real-world policy and understand how ConformalDagger enables the robot learner to query a real human teleoperator. Our goal is to train the robot to perform a real-world cleaning task where it wipes up a line drawn by an Expo marker on a whiteboard with sponge.

Human Expert. The human expert teleoperates the robot via a Meta Quest 3 remote controller. They initially provide 100 demonstrated trajectories moving in a straight-line path along the marker line to construct the initial training dataset, \mathcal{D}_0 . We model human-gated interventions with probability $\mathbb{P}(\text{obs}_t^h = 1 \mid x_t) = 0.2, \forall t$. To control this rate, our interface queries the user with 20% probability at the any timepoint along deployment to simulate that consistent independent intervention rate. We induce a potential distribution shift during interactive deployment episodes by changing the expert’s wiping strategy from a straight line to a zig-zag pattern.

Learner Policy. We represent the robot’s policy as a CNN-based diffusion policy (Chi et al., 2023). The policy predicts 16 future actions, where each action is an end-effector position and quaternion orientation. The policy takes as input the current and previous image observations from the wrist and third-person camera. Image observations are encoded using a ResNet-18 visual encoder (trained

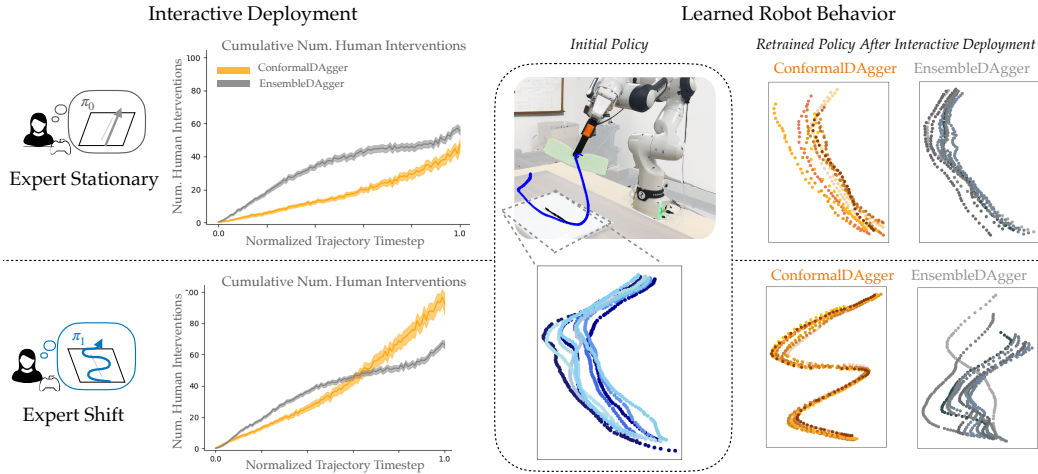


Figure 6: **Hardware Robot Results.** (left) ConformalDagger increases the number of questions asked when the expert shifts in their strategy while Ensemble maintains a very similar question asking profile. (right) Qualitatively, ConformalDagger can more effectively learn the new zigzag strategy.

end-to-end with the diffusion policy) and the action-generating process is conditioned on encoded observation features with FiLM (Perez et al., 2018). The initial policy π_0^r is trained for 60K iterations (training parameters in Appendix Sec E.1).

The robot is deployed to interactively execute the task 50 times after which its policy is retrained. This is an interactive deployment where the robot queries the user via our ConformalDagger algorithm, or the user can independently intervene intermittently. During expert feedback, the human teleoperates the robot for a sequence of 16 timesteps, after which, if the robot doesn’t ask for feedback, control is handed back over to the robot. After the interactive deployment, the initial learner policy is fine-tuned for an additional 60K iterations. We reset the learning-rate schedule, aggregate datapoints at which the expert gave feedback, and cap the training buffer size at 12k most recent input-action pairs.

ConformalDagger Hyperparameters. We apply IQT on only the predicted end-effector position, because we reason about nonconformity via the signed residual in Euclidean space. We initialize $q_0^{lo}, q_0^{hi} = 0.01$, our desired coverage level is $\alpha = 0.1$, and $\gamma_t = 0.15\hat{B}_t$, where the lookback window for \hat{B}_t is 20. We use a simple model for $\mathbb{P}(\text{obs}^r | x_t) = 1$ that asks for help if the L2 norm of the uncertainty interval $C_t(x_t)$ exceeds threshold $\tau = 0.07$.

Baseline. Because the diffusion policy implicitly represents the action distribution present in the training data, we sample from the policy $N=3$ times and evaluate variance over the predicted action in order to capture the notion of epistemic uncertainty. We consider this an *implicit* EnsembleDagger (Wolleb et al., 2022). Our implicit **EnsembleDagger** implementation does not use the safety classifier. EnsembleDagger’s uncertainty threshold is 0.0005, tuned such that the learner asks occasional but not excessive questions. EnsembleDagger and ConformalDagger use the same initial learner policy.

Results. The left of Figure 6 shows quantitative results for both methods. **ConformalDagger** increases the cumulative number of interventions up to about 100 timesteps of human feedback when interacting with the *shifted* expert (who switches to a zig-zag pattern) compared to the *stationary* expert (who always wipes in a straight line) which reaches a maximum of 40 timesteps. On the other hand, **EnsembleDagger** maintains a very similar question-asking profile (Figure 6) for both the *stationary* and *shifted* expert, near 60 timesteps. To better understand this, we analyzed Figure 1, which shows **ConformalDagger** engaging with each type of expert in one interactive deployment episode. Here, we see how ConformalDagger does have uncertainty when interacting with the stationary expert (due to inherent noise in the VR-teleoperated intermittent human interventions) but the robot’s uncertainty never rises above the threshold where it *initiates* expert feedback. This is in stark contrast to the expert shift scenario (right of Figure 1) where half-way through the deployment episode the uncertainty intervals become large enough to trigger robot asking for help. On the other hand, **EnsembleDagger** requests more human interventions than **ConformalDagger** at the *start* of the trajectory (as seen by the slope of cumulative interventions in Figure 6) but then plateaus. We

hypothesize that this may be because the training demonstrations have low variance while the sponge is in contact with the whiteboard, but higher variance as the robot approaches the table. Because **EnsembleDagger**'s uncertainty is uncalibrated to the deployment-time expert's data, it is unable to identify the need for additional feedback when it reaches the whiteboard.

Qualitative results are shown on the right of Figure 6. As expected, with **ConformalDagger** and **EnsembleDagger** rollouts from the policies remain aligned with the *stationary* expert when retrained. However, under the *shifted* expert, rollouts from the retrained **ConformalDagger** policy exhibit a more distinct zig-zag pattern, compared to **EnsembleDagger**, likely because the EnsembleDagger dataset has less datapoints because it does not ask for as much help.

7 CONCLUSION

In this work, we study uncertainty quantification for imitation learned policies that encounter distribution shift. We first extend uncertainty quantification via online conformal prediction to handle intermittent labels, such as those observed in interactive imitation learning. We then propose ConformalDagger, a unification of our online conformal prediction algorithm with interactive imitation learning. Our approach provides asymptotic coverage guarantees for deployed end-to-end policies, uses the calibrated uncertainty measure to detect expert distribution shifts and actively query for more feedback, and empirically enables the robot learner update its policy to better align with the shifted expert distribution.

ACKNOWLEDGMENTS

The authors would like to thank Gokul Swamy for insightful conversations and the detailed review, Yilin Wu for help with diffusion policy and robot hardware setup, and Pat Callaghan for the detailed review. MZ is supported by an NDSEG fellowship.

REFERENCES

- Anastasios Angelopoulos and Stephen Bates. Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4):494–591, 2023.
- Anastasios Angelopoulos, Emmanuel Candes, and Ryan J Tibshirani. Conformal PID control for time series prediction. *Advances in Neural Information Processing Systems*, 36, 2024a.
- Anastasios N Angelopoulos, Rina Foygel Barber, and Stephen Bates. Online conformal prediction with decaying step sizes. In *Proceedings of the 41st International Conference on Machine Learning*, 2024b.
- Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- Vassilis Assimakopoulos and Konstantinos Nikolopoulos. The theta model: a decomposition approach to forecasting. *International Journal of Forecasting*, 16(4):521–530, 2000.
- Aadyot Bhatnagar, Huan Wang, Caiming Xiong, and Yu Bai. Improved online conformal prediction via strongly adaptive online learning. In *International Conference on Machine Learning*, pp. 2337–2363. PMLR, 2023.
- Carlos Celemin, Rodrigo Pérez-Dattari, Eugenio Chisari, Giovanni Franzese, Leandro de Souza Rosa, Ravi Prakash, Zlatan Ajanović, Marta Ferraz, Abhinav Valada, Jens Kober, et al. Interactive imitation learning in robotics: A survey. *Foundations and Trends® in Robotics*, 10(1-2):1–197, 2022.
- Jonathan Chang, Masatoshi Uehara, Dhruv Sreenivas, Rahul Kidambi, and Wen Sun. Mitigating covariate shift in imitation learning via offline data with partial coverage. *Advances in Neural Information Processing Systems*, 34:965–979, 2021.
- Yuxiao Chen, Ugo Rosolia, Chuchu Fan, Aaron Ames, and Richard Murray. Reactive motion planning with probabilistic safety guarantees. In *Conference on Robot Learning*, pp. 1958–1970. PMLR, 2021.

- Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Robotics: Science and Systems (RSS)*, 2023.
- Yuchen Cui, David Isele, Scott Niekum, and Kikuo Fujimura. Uncertainty-aware data aggregation for deep imitation learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 761–767. IEEE, 2019.
- Thomas G Dietterich and Jesse Hostetler. Conformal prediction intervals for Markov decision process trajectories. *arXiv preprint arXiv:2206.04860*, 2022.
- Anushri Dixit, Lars Lindemann, Skylar X Wei, Matthew Cleaveland, George J Pappas, and Joel W Burdick. Adaptive conformal prediction for motion planning among dynamic agents. In *Learning for Dynamics and Control Conference*, pp. 300–314. PMLR, 2023.
- Isaac Gibbs and Emmanuel Candes. Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems*, 34:1660–1672, 2021.
- Isaac Gibbs and Emmanuel J Candès. Conformal inference for online prediction with arbitrary distribution shifts. *Journal of Machine Learning Research*, 25(162):1–36, 2024.
- Till Grüne-Yanoff. Models of temporal discounting 1937–2000: An interdisciplinary exchange between economics and psychology. *Science in context*, 28(4):675–713, 2015.
- Michael Harries. Splice-2 comparative evaluation: Electricity pricing. 1999.
- Julien Herzen, Francesco Lässig, Samuele Giuliano Piazzetta, Thomas Neuer, Léo Tafti, Guillaume Raille, Tomas Van Pottelbergh, Marek Pasięka, Andrzej Skrodzki, and Nicolas Huguenin. Darts: User-friendly modern machine learning for time series. *Journal of Machine Learning Research*, 23(124):1–6, 2022.
- Joey Hong, Sergey Levine, and Anca Dragan. Learning to influence human behavior with offline reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Ryan Hoque, Ashwin Balakrishna, Ellen Novoseller, Albert Wilcox, Daniel S Brown, and Ken Goldberg. Thriftydagger: Budget-aware novelty and risk gating for interactive imitation learning. *arXiv preprint arXiv:2109.08273*, 2021a.
- Ryan Hoque, Ashwin Balakrishna, Carl Putterman, Michael Luo, Daniel S Brown, Daniel Seita, Brijen Thananjeyan, Ellen Novoseller, and Ken Goldberg. Lazydagger: Reducing context switching in interactive imitation learning. In *2021 IEEE 17th international conference on automation science and engineering (case)*, pp. 502–509. IEEE, 2021b.
- Ryan Hoque, Lawrence Yunliang Chen, Satvik Sharma, Karthik Dharmarajan, Brijen Thananjeyan, Pieter Abbeel, and Ken Goldberg. Fleet-dagger: Interactive robot fleet learning with scalable human supervision. In *Conference on Robot Learning*, pp. 368–380. PMLR, 2023.
- Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. Bc-z: Zero-shot task generalization with robotic imitation learning. In *Conference on Robot Learning*, pp. 991–1002. PMLR, 2022.
- Michael Kelly, Chelsea Sidrane, Katherine Driggs-Campbell, and Mykel J Kochenderfer. Hg-dagger: Interactive imitation learning with human experts. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8077–8083. IEEE, 2019.
- Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. OpenVLA: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- Roger Koenker and Gilbert Bassett Jr. Regression quantiles. *Econometrica: journal of the Econometric Society*, pp. 33–50, 1978.
- Minae Kwon, Erdem Biyik, Aditi Talati, Karan Bhasin, Dylan P Losey, and Dorsa Sadigh. When humans aren’t optimal: Robots that collaborate with risk-aware humans. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 43–52, 2020.

- Michael Laskey, Jonathan Lee, Roy Fox, Anca Dragan, and Ken Goldberg. Dart: Noise injection for robust imitation learning. In *Conference on robot learning*, pp. 143–156. PMLR, 2017.
- Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(39):1–40, 2016.
- Justin Lidard, Hang Pham, Ariel Bachman, Bryan Boateng, and Anirudha Majumdar. Risk-calibrated human-robot interaction via set-valued intent prediction. In *Robotics: Science and Systems*, 2024.
- Amarildo Likmeta, Alberto Maria Metelli, Giorgia Ramponi, Andrea Tirinzoni, Matteo Giuliani, and Marcello Restelli. Dealing with multiple experts and non-stationarity in inverse reinforcement learning: an application to real-life problems. *Machine Learning*, 110:2541–2576, 2021.
- Albert Lin and Somil Bansal. Verification of neural reachable tubes via scenario optimization and conformal prediction. In *6th Annual Learning for Dynamics & Control Conference*, pp. 719–731. PMLR, 2024.
- Lars Lindemann, Matthew Cleaveland, Gihyun Shim, and George J Pappas. Safe planning in dynamic environments using conformal prediction. *IEEE Robotics and Automation Letters*, 2023.
- Rachel Luo, Shengjia Zhao, Jonathan Kuck, Boris Ivanovic, Silvio Savarese, Edward Schmerling, and Marco Pavone. Sample-efficient safety assurances using conformal prediction. In *International Workshop on the Algorithmic Foundations of Robotics*, pp. 149–169. Springer, 2022.
- Kunal Menda, Katherine Driggs-Campbell, and Mykel J Kochenderfer. EnsembleDagger: A Bayesian approach to safe imitation learning. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5041–5048. IEEE, 2019.
- Anish Muthali, Haotian Shen, Sampada Deglurkar, Michael H Lim, Rebecca Roelofs, Aleksandra Faust, and Claire Tomlin. Multi-agent reachability calibration with conformal prediction. *arXiv preprint arXiv:2304.00432*, 2023.
- Cam Nguyen. S&P 500 stock data, 2018. URL <https://www.kaggle.com/datasets/camnugent/sandp500>.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pp. 8162–8171. PMLR, 2021.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- Bob Price and Craig Boutilier. Accelerating reinforcement learning through implicit imitation. *Journal of Artificial Intelligence Research*, 19:569–629, 2003.
- Allen Z Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Pen Xu, Leila Takayama Takayama, Fei Xia, Jake Varley, et al. Robots that ask for help: Uncertainty alignment for large language model planners. In *Conference on Robot Learning (CoRL)*. Proceedings of the Conference on Robot Learning (CoRL), 2023.
- Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized quantile regression. *Advances in neural information processing systems*, 32, 2019.
- Yaniv Romano, Matteo Sesia, and Emmanuel Candes. Classification with valid and adaptive coverage. *Advances in Neural Information Processing Systems*, 33:3581–3591, 2020.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Shahabedin Sagheb, Ye-Ji Mun, Neema Ahmadian, Benjamin A Christie, Andrea Bajcsy, Katherine Driggs-Campbell, and Dylan P Losey. Towards robots that influence humans over long-term interaction. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7490–7496. IEEE, 2023.

- Stefan Schaal. Learning from demonstration. *Advances in Neural Information Processing Systems*, 9, 1996.
- Sangwoo Shin, Daehee Lee, Minjong Yoo, Woo Kyung Kim, and Honguk Woo. One-shot imitation in a non-stationary environment via multi-modal skill. In *International Conference on Machine Learning*, pp. 31562–31578. PMLR, 2023.
- Gokul Swamy, Nived Rajaraman, Matt Peng, Sanjiban Choudhury, J Bagnell, Steven Z Wu, Jiantao Jiao, and Kannan Ramchandran. Minimax optimal online imitation learning via replay estimation. *Advances in Neural Information Processing Systems*, 35:7077–7088, 2022.
- Muhammad Faaiz Taufiq, Jean-Francois Ton, Rob Cornish, Yee Whye Teh, and Arnaud Doucet. Conformal off-policy prediction in contextual bandits. *Advances in Neural Information Processing Systems*, 35:31512–31524, 2022.
- Sean J Taylor and Benjamin Letham. Forecasting at scale. *The American Statistician*, 72(1):37–45, 2018.
- Andrea L Thomaz and Cynthia Breazeal. Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence*, 172(6-7):716–737, 2008.
- Andrea L Thomaz, Guy Hoffman, and Cynthia Breazeal. Reinforcement learning with human teachers: Understanding how people want to teach robots. In *ROMAN 2006-The 15th IEEE International Symposium on Robot and Human Interactive Communication*, pp. 352–357. IEEE, 2006.
- Suzanne C Thompson. Illusions of control: How we overestimate our personal influence. *Current Directions in Psychological Science*, 8(6):187–190, 1999.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.
- Julia Wolleb, Robin Sandkühler, Florentin Bieder, Philippe Valmaggia, and Philippe C Cattin. Diffusion models for implicit image segmentation ensembles. In *International Conference on Medical Imaging with Deep Learning*, pp. 1336–1348. PMLR, 2022.
- Annie Xie, Dylan Losey, Ryan Tolsma, Chelsea Finn, and Dorsa Sadigh. Learning latent representations to influence multi-agent interaction. In *Conference on robot learning*, pp. 575–588. PMLR, 2021.
- Margaux Zaffran, Olivier Féron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In *International Conference on Machine Learning*, pp. 25834–25866. PMLR, 2022.
- Jiakai Zhang and Kyunghyun Cho. Query-efficient imitation learning for end-to-end autonomous driving. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Boyuan Zheng, Sunny Verma, Jianlong Zhou, Ivor W Tsang, and Fang Chen. Imitation learning: Progress, taxonomies and challenges. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.

APPENDIX

A PROOF OF PROPOSITION 1

We will start by proving the following lemma. Recall that B is the upper bound on $q_1 \in [0, B]$, and $s_t \in [0, B]$ by definition.

Lemma 1. *For all t , we have $-\alpha N_{t-1} \leq q_t \leq B + (1 - \alpha)N_{t-1}$, where $N_t = \max_{1 \leq r \leq t} \frac{\gamma_r}{p_r}$. B is the upper bound on $q_1 \in [0, B]$ and $s_t \in [0, B]$.*

Proof. $q_1 \in [0, B]$ by assumption, so the lemma is satisfied at $t = 1$. Assume $q_t \in [-\alpha N_{t-1}, B + (1 - \alpha)N_{t-1}]$. Now, for q_{t+1} , we consider Case 1, where $\text{obs}_t = 0$, which means that $q_{t+1} = q_t$, so q_{t+1} lies within the range $[-\alpha N_{t-1}, B + (1 - \alpha)N_{t-1}]$. Since $N_t \geq N_{t-1}$, q_{t+1} also lies within the larger range $[-\alpha N_t, B + (1 - \alpha)N_t]$, as desired. In Case 2 where $\text{obs}_t = 1$, $q_{t+1} = q_t + \frac{\gamma_t}{p_t}(\text{err}_t - \alpha)$. If we represent $\eta_t := \frac{\gamma_t}{p_t}$, we obtain the constant-feedback quantile tracking update: $q_{t+1} = q_t + \eta_t(\text{err}_t - \alpha)$, with variable η_t instead of γ_t . Then, Lemma 1 in Angelopoulos et al. (2024b), with $N_t = \max_{1 \leq r \leq t} \eta_r$ bounds q_{t+1} within the range $[-\alpha N_t, B + (1 - \alpha)N_t]$. \square

Next, we proceed to show how we can leverage the existing quantile tracking results to derive our results with intermittent feedback.

Proof. First, we take expectations only with respect to obs_t conditional on x_t , err_t , and all other randomness, noting in particular that obs_t is independent of everything else given x_t and has a Bernoulli distribution with mean p_t .

We will abbreviate the left-hand side by $\mathbb{E}[q_{T+1}|D_{T+1}]$, where $D_{T+1} := \{\text{err}_t, x_t\}_{t \leq T+1}$, and use $\mathbb{E}_{\text{obs}_{1:t}} = \mathbb{E}_{\text{obs}_1 \sim p_1, \dots, \text{obs}_t \sim p_t}$ to denote this conditional expectation for brevity below:

$$\mathbb{E}_{\text{obs}_{1:t}}[q_{T+1} - q_r | D_{T+1}] = \mathbb{E}_{\text{obs}_{1:t}}[q_{T+1} | D_{T+1}] - \mathbb{E}_{\text{obs}_{1:r}}[q_r | D_r] \quad (9)$$

$$= \mathbb{E}_{\text{obs}_{1:t}}[q_1 + \sum_{t=1}^T \frac{\gamma_t}{p_t}(\text{err}_t - \alpha)\text{obs}_t | D_{T+1}] - \mathbb{E}_{\text{obs}_{1:r}}[q_1 + \sum_{t=1}^r \frac{\gamma_t}{p_t}(\text{err}_t - \alpha)\text{obs}_t | D_r] \quad (10)$$

$$= q_1 + \sum_{t=1}^T \frac{\gamma_t}{p_t}(\text{err}_t - \alpha)p_t - (q_1 + \sum_{t=1}^r \frac{\gamma_t}{p_t}(\text{err}_t - \alpha)p_t) \quad (11)$$

$$= \sum_{t=r}^T \gamma_t(\text{err}_t - \alpha). \quad (12)$$

Given the definition of Δ , we have that $\gamma_t^{-1} = \sum_{r=1}^t \Delta_r$ for all $t \geq 1$. So

$$\left| \frac{1}{T} \sum_{t=1}^T (\text{err}_t - \alpha) \right| = \left| \frac{1}{T} \sum_{t=1}^T \left(\sum_{r=1}^t \Delta_r \right) \gamma_t (\text{err}_t - \alpha) \right| \quad (13)$$

$$= \left| \frac{1}{T} \sum_{r=1}^T \Delta_r \left(\sum_{t=r}^T \gamma_t (\text{err}_t - \alpha) \right) \right| \quad (14)$$

$$= \left| \frac{1}{T} \sum_{r=1}^T \Delta_r \left(\mathbb{E}_{\text{obs}_{1:t}}[q_{T+1} - q_r | D_{T+1}] \right) \right|. \quad (15)$$

By Lemma 1, this expected difference is bounded by $B + (1 - \alpha)N_T - (-\alpha N_T)$:

$$\mathbb{E}_{\text{obs}_{1:T+1}}[q_{T+1} - q_r | D_{T+1}] \leq B + \max_{1 \leq t \leq T} \frac{\gamma_t}{p_t}. \quad (16)$$

In other words, we can drop the expectation via Lemma 1 and consider the worst case bound on $q_{T+1} - q_r$. Thus,

$$\left| \frac{1}{T} \sum_{t=1}^T \text{err}_t - \alpha \right| \leq \frac{1}{T} \sum_{r=1}^T |\Delta_r| \left(B + \max_{1 \leq t \leq T} \frac{\gamma_t}{p_t} \right) \quad (17)$$

$$= \frac{1}{T} \|\Delta_{1:T}\|_1 \left(B + \max_{1 \leq t \leq T} \frac{\gamma_t}{p_t} \right). \quad (18)$$

This completes the proof. \square

A.1 SPECIAL CASE: WHEN $\gamma_t = p_t$

Next, given the quantile tracking update with an intermittent observation model, we consider what would happen if γ_t was set as p_t . When $\gamma_t = p_t$, then quantile tracking update becomes $q_{t+1} = q_t + (\text{err}_t - \alpha) \text{obs}_t$ and $\max_{1 \leq t \leq T} \frac{\gamma_t}{p_t} = 1$.

Under Proposition 1, IQT with $\gamma_t = p_t$ gives the following finite time coverage bound:

$$\left| \frac{1}{T} \sum_{t=1}^T \text{err}_t - \alpha \right| \leq \frac{B+1}{T} \|\Delta_{1:T}\|_1 \quad (19)$$

where the sequence Δ is defined with values $\Delta_1 = \gamma_1^{-1}$ and $\Delta_t = \gamma_t^{-1} - \gamma_{t-1}^{-1}$ for all $t \geq 2$.

B INTERMITTENT ADAPTIVE CONFORMAL INFERENCE

We show in this section that intermittent observation of ground truth labels can be extended to Adaptive Conformal Inference (ACI) (Gibbs & Candès, 2021). To facilitate understanding, we briefly summarize ACI and discuss our extension Intermittent Adaptive Conformal Inference (IACI).

Setup: Quantile Regression (QR). Similar to IQT, consider an *arbitrary* sequence of data points $(x_t, y_t) \in \mathcal{X} \times \mathcal{Y}$, for $t = 1, 2, \dots$, that are not necessarily I.I.D. Our goal in ACI is to also produce prediction sets on the output of any base prediction model such that the sets contain the true label with a specified miscoverage rate α . Mathematically, at each time t , we observe x_t and seek to cover the true label y_t with a set $C_t(x_t)$, which depends on a base prediction model, $\hat{f} : \mathcal{X} \rightarrow \mathcal{Y}$. We will discuss ACI with a conformal quantile regression Romano et al. (2019) backbone. The base model takes as input the current x_t and outputs prediction \hat{y}_t as well the *estimated upper and lower conditional quantiles*:

$$\{\hat{q}_{\alpha_{lo}}(x_t), \hat{y}_t, \hat{q}_{\alpha_{hi}}(x_t)\} \leftarrow \hat{f}(x_t), \quad \forall (x_t, y_t), \quad (20)$$

where $\hat{q}_{\alpha_{lo}}(x_t)$ is an estimate of the α_{lo} -th conditional quantile and $\hat{q}_{\alpha_{hi}}(x_t)$ is the α_{hi} -th quantile estimate. During training, $\hat{q}_{\alpha_{lo}}(x_t)$ is learned with an additional Pinball loss Koenker & Bassett Jr (1978); Romano et al. (2019).

Adaptive Conformal Inference (ACI). At each time t , we compute the nonconformity score s_t as:

$$s(x_t, y_t; f) = \max\{\hat{q}_{\alpha_{lo}}(x_t) - y_t, y_t - \hat{q}_{\alpha_{hi}}(x_t)\}, \quad (21)$$

which is the coverage error induced by the regressor's quantile estimates.

Let S_t be the set of conformity scores for all data points through time t in $\mathcal{D}_{\text{calib}}$. In general, the magnitude of $s(x_t, y_t; \hat{f})$ is determined by the miscoverage error and its sign is determined by if the true value of y_t lies outside or inside the estimated interval.

Mathematically, the calibrated prediction interval for Y_t is:

$$C_t(x_t) = [\hat{q}_{\alpha_{lo}}(x_t) - \hat{Q}_{S_t}(1 - \alpha_t), \quad (22)$$

$$\hat{q}_{\alpha_{hi}}(x_t) + \hat{Q}_{S_t}(1 - \alpha_t)], \quad (23)$$

where the $\hat{Q}_{S_t}(1 - \alpha_t) := (1 - \alpha_t)(1 + \frac{1}{|\mathcal{D}_{\text{calib}}|})$ -th adaptive empirical quantile of $S_t = S_{t-1} \cup s(x_t, y_t; \hat{f})$. The empirical quantile is defined as the following (where k is the lookback window):

$$\hat{Q}_{S_t}(c) := \inf \left\{ m : \left(\frac{1}{|D_{t-k:t}|} \sum_{(x_i, y_i) \in D_{t-k:t}} \mathbb{1}_{\{s(X_i, Y_i) \leq m\}} \right) \geq c \right\} \quad (24)$$

Given the non-stationarity of the data distribution, ACI examines the empirical miscoverage frequency of the previous interval, and then decreases or increases a time-dependent α_t . Fixing step size parameter $\gamma > 0$, ACI updates

$$\alpha_{t+1} := \alpha_t + \gamma(\alpha - \text{err}_t) \quad (25)$$

ACI Coverage Guarantee. The adaptive quantile adjustments made in ACI provide the following coverage guarantee:

$$\left| \frac{1}{T} \sum_{t=1}^T \text{err}_t - \alpha \right| \leq \frac{\max\{\alpha_1, 1 - \alpha_1\} + \gamma}{T\gamma} \quad (26)$$

obtaining the desired α coverage frequency without making an assumptions on the data-generating distribution (Proposition 4.1 in Gibbs & Candes (2021)). As T approaches ∞ , $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \text{err}_t$ approaches α . This guarantees ACI gives the $1 - \alpha$ long-term empirical coverage frequency regardless of the underlying data generation process.

B.1 INTERMITTENT ADAPTIVE CONFORMAL PREDICTION

To achieve Intermittent Adaptive Conformal Inference (IACI), we update α_t at each timestep t with Equation 27:

$$\alpha_{t+1} = \alpha_t + \frac{\gamma}{p_t}(\alpha - \text{err}_t) \text{obs}_t \quad (27)$$

Recall $\text{err}_t = \mathbb{1}_{y_t \notin C_t(x_t)}$ and obs_t represents whether y_t was observed at timestep t , and $p_t = \mathbb{P}(\text{obs}_t | x_t) \in (0, 1]$. The calibrated prediction interval for y_t becomes:

$$C_t(x_t) = [\hat{q}_{\alpha_{t_0}}(x_t) - \hat{Q}_{S_t^{\text{obs}}}(1 - \alpha_t), \hat{q}_{\alpha_{h_i}}(X_t) + \hat{Q}_{S_t^{\text{obs}}}(1 - \alpha_t)], \quad (28)$$

where $\hat{Q}_{S_t^{\text{obs}}}(1 - \alpha_t) := (1 - \alpha_t)(1 + \frac{1}{|\mathcal{D}_{\text{calib}}|})$ -th adaptive empirical quantile of S_t^{obs} , the set of nonconformity scores that have been observed, weighted by their probability of observation. Since we will only observe feedback with probability p_t at each timestep, our set of nonconformity values will not be the full set of scores at every timestep. Instead, we have access to some subset of nonconformity scores S_t^{obs} , where p_i is the probability of element $s(x_i, y_i; \hat{f})$ being in the subset S_t^{obs} for $1 \leq i \leq t$. We define $\hat{Q}_{S_t^{\text{obs}}}(c)$:

$$\hat{Q}_{S_t^{\text{obs}}}(c) := \inf \left\{ m : \left(\frac{1}{|D_{t-k:t}|} \sum_{(x_i, y_i) \in D_{t-k:t}} \frac{1}{p_i} \cdot \text{obs}_i \cdot \mathbb{1}_{\{s(x_i, y_i) \leq m\}} \right) \geq c \right\} \quad (29)$$

At best, we can say that in expectation, the inclusion criteria for any element in the summation term is equivalent for the summation in $\hat{Q}_t^{\text{obs}}(x)$ and the summation in $\hat{Q}_t(x)$: $\mathbb{E}_{\text{obs}_i \sim p_i} [\frac{1}{p_i} \cdot \text{obs}_i \cdot \mathbb{1}_{\{s(X_i, Y_i) \leq m\}}] = \mathbb{1}_{\{s(X_i, Y_i) \leq m\}}$.

Proposition 2. IACI Coverage Guarantee. With this intermittent feedback update, with probability one, where $M_t = \min \{p_1, \dots, p_t\} \in (0, 1]$, we have that for all $T \in \mathbb{N}$:

$$\left| \frac{1}{T} \sum_{t=1}^T \text{err}_t - \alpha \right| \leq \frac{\max\{\alpha_1, 1 - \alpha_1\} + \frac{\gamma}{M_{T+1}}}{T\gamma} \quad (30)$$

At $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \text{err}_t$ approaches α , given $M_\infty = \min \{p_1, \dots, p_\infty\} \in (0, 1]$. This guarantees IACI gives the $1 - \alpha$ long-term empirical coverage frequency regardless of the underlying data generation process. The proof follows from Gibbs & Candes (2021) as is described in the next subsection (Section B.2).

B.2 PROOF FOR INTERMITTENT ADAPTIVE CONFORMAL INFERENCE

Assumptions. We will assume throughout that with probability one, $\alpha_1 \in [0, 1]$, $\alpha \in (0, 1)$, $p_t \in (0, 1] \forall 1 \leq t \leq \infty$, and $\hat{Q}_t^{obs}(x)$ is non-decreasing with $\hat{Q}_{S_t^{obs}}(c) = -\infty$ for all $c < 0$, and $\hat{Q}_{S_t^{obs}}(c) = \infty$ for all $c > 1$.

Lemma 2. *With probability one, we have that $\forall t \in [1, T]$,*

$$\alpha_t \in \left[-\frac{\gamma}{M_T}, 1 + \frac{\gamma}{M_T}\right] \quad (31)$$

where $M_T = \min_{1 \leq r \leq T} p_r$.

Proof. Observe that

$$\sup_{1 \leq t \leq T} |\alpha_{t+1} - \alpha_t| = \sup_{1 \leq t \leq T} \left| \frac{\gamma}{p_t} (\alpha - \text{err}_t) \right| \quad (32)$$

$$\leq \sup_{1 \leq t \leq T} \frac{\gamma}{M_T} |(\alpha - \text{err}_t)| = \frac{\gamma}{M_T} \sup_{1 \leq t \leq T} |(\alpha - \text{err}_t)| \quad (33)$$

$$< \frac{\gamma}{M_T} \quad (34)$$

The rest of the proof follows from Lemma 4.1 in Gibbs & Candes (2021). We will write it out explicitly here.

Lower Bound. Assume towards contradiction that with positive probability the set $\{\alpha_t\}_{t \in [1, T]}$ is such that $\inf_{t \in [1, T]} \{\alpha_t\} < -\frac{\gamma}{M_T}$, which means there exists some element in the set $\{\alpha_t\}_{t \in [1, T]}$ such that the element is less than $-\frac{\gamma}{M_T}$. Let α_{t+1} be the first element in set $\{\alpha_t\}_{t \in [1, T]}$ such that $\alpha_{t+1} < -\frac{\gamma}{M_T}$.

We know by assumption that $\alpha_1 \in [0, 1]$, so α_{t+1} cannot be the first value. So $\alpha_t > \alpha_{t+1}$, else the latter would not be the first element such that $\alpha_{t+1} < -\frac{\gamma}{M_T}$. Since we know that $\alpha_t - \alpha_{t+1} < \frac{\gamma}{M_T}$ (from the observation above), then

$$\alpha_t - \alpha_{t+1} < \frac{\gamma}{M_T} \quad (35)$$

$$\Leftrightarrow \alpha_t < \frac{\gamma}{M_T} + \alpha_{t+1} \quad (36)$$

$$\Leftrightarrow \alpha_t < \frac{\gamma}{M_T} + \alpha_{t+1} < \frac{\gamma}{M_T} + \left(-\frac{\gamma}{M_T}\right) \quad (37)$$

$$\Leftrightarrow \alpha_t < 0. \quad (38)$$

However, if $\alpha_t < 0$, then $\hat{Q}_{S_t^{obs}}(1 - \alpha_t) = \infty \Rightarrow \text{err}_t = 0$. This is because the quantile will be the trivially large infinite quantile, meaning there will definitely be no undercoverage at t . Since $\frac{\gamma}{p_t} \alpha$ is positive by definition,

$$\Rightarrow \alpha_{t+1} = \alpha_t + \frac{\gamma}{p_t} (\alpha - \text{err}_t) = \alpha_t + \frac{\gamma}{p_t} (\alpha - 0) \geq \alpha_t \quad (39)$$

We have reached a contradiction with α_{t+1} being the first infimum value reached.

Upper Bound. The upper bound argument is symmetric but we will write it out explicitly. Assume towards contradiction that with positive probability the set $\{\alpha_t\}_{t \in [1, T]}$ is such that $\sup_{t \in [1, T]} \{\alpha_t\} > 1 + \frac{\gamma}{M_T}$, which means there exists some element in the set $\{\alpha_t\}_{t \in [1, T]}$ such that the element is greater than $1 + \frac{\gamma}{M_T}$. Let α_{t+1} be the first element in set $\{\alpha_t\}_{t \in [1, T]}$ such that $\alpha_{t+1} > 1 + \frac{\gamma}{M_T}$.

We know by assumption that $\alpha_1 \in [0, 1]$, so α_{t+1} cannot be the first value. So there exists some $\alpha_t < \alpha_{t+1}$, where $\alpha_{t+1} > 1 + \frac{\gamma}{M_T}$. Since we know that $\alpha_{t+1} - \alpha_t < \frac{\gamma}{M_T}$ (from the observation

above), then

$$\alpha_{t+1} - \alpha_t < \frac{\gamma}{M_T} \quad (40)$$

$$\Leftrightarrow -\alpha_t < \frac{\gamma}{M_T} - \alpha_{t+1} \quad (41)$$

$$\Leftrightarrow \alpha_t > \alpha_{t+1} - \frac{\gamma}{M_T} > (1 + \frac{\gamma}{M_T}) - \frac{\gamma}{M_T} \quad (42)$$

$$\Leftrightarrow \alpha_t > 1. \quad (43)$$

$$(44)$$

However, if $\alpha_t > 1$, then $\hat{Q}_{S_t^{obs}}(1 - \alpha_t) = -\infty \Rightarrow \text{err}_t = 1$. This is because the quantile will be the trivially small negative infinite quantile, meaning there will definitely be miscoverage at t . Then,

$$\Rightarrow \alpha_{t+1} = \alpha_t + \frac{\gamma}{p_t}(\alpha - \text{err}_t) = \alpha_t + \frac{\gamma}{p_t}(\alpha - 1) \leq \alpha_t \quad (45)$$

since $\frac{\gamma}{p_t}(\alpha - 1)$ is negative by definition of α . We have reached a contradiction with α_{t+1} being the first supremum value reached. \square

Lemma 2 will enable us to prove Proposition 2. The proof is derivative of the constant feedback ACI proof in (Gibbs & Candes, 2021), and the key idea is to bound the expectation of α_{t+1} .

Proof. Examine the expectation of α_{t+1} , conditional on $D_{T+1} := \{\text{err}_t, p_t\}_{t \leq T+1}$:

$$\mathbb{E}_{\text{obs}_1 \sim p_1, \dots, \text{obs}_t \sim p_{T+1}}[\alpha_{T+1} | D_{T+1}] = \alpha_1 + \gamma \sum_{t=1}^{T+1} (\alpha - \text{err}_t) \quad (46)$$

We will abbreviate the left hand side by $\mathbb{E}[\alpha_{T+1} | D_{T+1}]$. Because the expected value cannot exceed the range of the value of α_{T+1} , we infer that $\mathbb{E}[\alpha_{T+1} | D_{T+1}] \in [-\frac{\gamma}{M_{T+1}}, 1 + \frac{\gamma}{M_{T+1}}]$.

First, we observe from Lemma 2 that

$$\mathbb{E}[\alpha_{T+1} | D_{T+1}] = \alpha_1 + \sum_{t=1}^T \gamma(\alpha - \text{err}_t) \in [-\frac{\gamma}{M_{T+1}}, 1 + \frac{\gamma}{M_{T+1}}] \quad (47)$$

$$\Rightarrow -\sum_{t=1}^T \gamma(\alpha - \text{err}_t) = \alpha_1 - \mathbb{E}[\alpha_{T+1} | D_{T+1}] \quad (48)$$

$$\Rightarrow \left| \frac{1}{T} \sum_{t=1}^T \text{err}_t - \alpha \right| = \frac{|\alpha_1 - \mathbb{E}[\alpha_{T+1} | D_{T+1}]|}{T\gamma}. \quad (49)$$

To bound the right hand side above, consider in turn the two cases of $\alpha_1 - \mathbb{E}[\alpha_{T+1} | D_{T+1}] \geq 0$ and $\alpha_1 - \mathbb{E}[\alpha_{T+1} | D_{T+1}] < 0$.

Starting with Case 1,

$$\alpha_1 - \mathbb{E}[\alpha_{T+1} | D_{T+1}] \geq 0 \Rightarrow \alpha_1 \geq \mathbb{E}[\alpha_{T+1} | D_{T+1}] \Rightarrow \alpha_{T+1} \in [-\frac{\gamma}{M_{T+1}}, \alpha_1]. \quad (50)$$

This case corresponds to the following equivalence: $|\alpha_1 - \mathbb{E}[\alpha_{T+1} | D_{T+1}]| = \alpha_1 - \mathbb{E}[\alpha_{T+1} | D_{T+1}]$.

Negating $\alpha_{T+1} \geq -\frac{\gamma}{M_{T+1}}$, we get $-\alpha_{T+1} \leq \frac{\gamma}{M_{T+1}}$. Thus,

$$|\alpha_1 - \mathbb{E}[\alpha_{T+1} | D_{T+1}]| = \alpha_1 - \mathbb{E}[\alpha_{T+1} | D_{T+1}] \leq \alpha_1 + \frac{\gamma}{M_{T+1}} \quad (51)$$

$$\Rightarrow \frac{|\alpha_1 - \mathbb{E}[\alpha_{T+1} | D_{T+1}]|}{T\gamma} \leq \frac{\alpha_1 + \frac{\gamma}{M_{T+1}}}{T\gamma}. \quad (52)$$

In Case 2,

$$\alpha_1 - \mathbb{E}[\alpha_{T+1}|D_{T+1}] < 0 \Rightarrow \alpha_1 < \mathbb{E}[\alpha_{T+1}|D_{T+1}] \Rightarrow \mathbb{E}[\alpha_{T+1}|D_{T+1}] \in (\alpha_1, 1 + \frac{\gamma}{M_{T+1}}]. \quad (53)$$

This case corresponds to the following equivalence: $|\alpha_1 - \mathbb{E}[\alpha_{T+1}|D_{T+1}]| = \mathbb{E}[\alpha_{T+1}|D_{T+1}] - \alpha_1 = -1(\alpha_1 - \mathbb{E}[\alpha_{T+1}|D_{T+1}])$. Plugging in $\mathbb{E}[\alpha_{T+1}|D_{T+1}] < 1 + \frac{\gamma}{M_{T+1}}$,

$$|\alpha_1 - \mathbb{E}[\alpha_{T+1}|D_{T+1}]| = \mathbb{E}[\alpha_{T+1}|D_{T+1}] - \alpha_1 \leq 1 + \frac{\gamma}{M_{T+1}} - \alpha_1 \quad (54)$$

$$\frac{|\alpha_1 - \mathbb{E}[\alpha_{T+1}|D_{T+1}]|}{T\gamma} \leq \frac{(1 - \alpha_1) + \frac{\gamma}{M_{T+1}}}{T\gamma}. \quad (55)$$

Lastly, we merge the two cases by taking the maximum over $\{\alpha_1, 1 - \alpha_1\}$ to come up with an upper bound that covers both cases.

$$\left| \frac{1}{T} \sum_{t=1}^T \text{err}_t - \alpha \right| = \frac{|\alpha_1 - \mathbb{E}[\alpha_{T+1}|D_{T+1}]|}{T\gamma} \leq \frac{\max\{\alpha_1, 1 - \alpha_1\} + \frac{\gamma}{M_{T+1}}}{T\gamma}. \quad (56)$$

Taking the limit as $T \rightarrow \infty$, if M_∞ is bounded, we get $\lim_{T \rightarrow \infty} \left| \frac{1}{T} \sum_{t=1}^T \text{err}_t - \alpha \right| = 0$, as claimed. \square

C EXTENDED EXPERIMENTS: INTERMITTENT QUANTILE TRACKING

We experiment with these four predictors because of their use in prior conformal literature (Angelopoulos et al., 2024a) and in order to ensure that under different base model conditions, IQT maintains coverage close to the desired level. (1) Autoregressive (**AR**) model with 3 lags, (2) **Theta** model with $\theta = 2$ (Assimakopoulos & Nikolopoulos, 2000), (3) **Prophet** model Taylor & Letham (2018), and (4) **Transformer** model (Vaswani et al., 2017). Consistent with prior works, for all base models except for transformer, we retrain the base model after each timestep; for the transformer, we retrain every 100 timesteps. We set lookback window $k = 100$ timesteps for the Google and Amazon stock price data, and $k = 300$ for the Elec2 dataset.

Amazon stock price data under partial ($p_t = 0.5$) feedback. Figure 7 shows the prediction interval sizes for 1 seed, and coverage averaged over 5 seeds for Amazon stock price data under partial ($p_t = 0.5$) feedback. We see that the interval size for **IQT-pd** is larger than **IQT-pi** for high learning rate $lr = 1$, but the size of intervals is comparable for smaller learning rates. Table 2 shows the performance metrics averaged over 5 seeds.




Figure 7: Amazon stock price data under partial ($p_t = 0.5$) feedback. We show the prediction interval sizes for 1 seed, and coverage averaged over 5 seeds for Amazon stock price data under partial observations.

Amazon stock price data under partial ($p_t = 0.9$) feedback. Figure 8 shows the prediction interval sizes for 1 seed, and coverage averaged over 5 seeds for Amazon stock price data under frequent ($p_t = 0.9$) feedback. We see that the interval size for **IQT-pd** and **IQT-pi** are very similar across learning rates. Table 3 shows the performance metrics averaged over 5 seeds.

Table 1: **IQT on Amazon Stocks (*infrequent observations*): Trends Across Models.** We test four base models, set $p_t = 0.1, \forall t$ to simulate seeing the true price only 10% of the time, and report the mean across 5 seeds.

Metric	lr	AR Base Model		Prophet Base Model		Theta Base Model		Transformer Base Model	
		IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd
Coverage	1.0	0.903	0.937	0.906	0.935	0.899	0.921	0.902	0.923
	0.1	0.876	0.903	0.876	0.906	0.815	0.899	0.808	0.902
	0.01	0.690	0.876	0.690	0.876	0.430	0.815	0.360	0.808
Longest err seq	1.0	13.2	13.2	11.6	30.6	27	41.2	28.8	30
	0.1	6.8	20.8	10	11.6	51.8	27	89	28.8
	0.01	12.6	6.8	13.8	10	145.2	51.8	212.8	89
Avg set size	1.0	47.452	425.435	46.056	414.977	105.040	711.829	131.025	982.469
	0.1	16.872	47.452	16.495	46.056	82.463	105.040	118.040	131.025
	0.01	9.428	16.872	9.470	16.495	35.458	82.463	46.929	118.040

Table 2: **IQT on Amazon Stocks (*partial observations*): Trends Across Models.** We test four base models, set $p_t = 0.5, \forall t$ to simulate seeing the true price only 50% of the time, and report the mean across 5 seeds.

Metric	lr	AR		Prophet		Theta		Transformer	
		IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd
Marginal coverage	1.0	0.902	0.905	0.901	0.904	0.906	0.904	0.900	0.904
	0.1	0.891	0.896	0.883	0.893	0.890	0.895	0.879	0.885
	0.01	0.846	0.869	0.737	0.812	0.843	0.870	0.705	0.796
Longest err seq	1.0	5.6	5.8	7.4	8	5.6	6.6	7.8	8
	0.1	4.6	4	12.8	8.2	7.6	5	23.4	11.4
	0.01	7.2	7	89.4	50.6	10	10	191.2	97.6
Avg set size	1.0	43.816	78.447	76.013	139.875	44.449	77.741	100.127	183.529
	0.1	17.535	19.918	66.940	55.233	17.581	19.854	92.513	72.867
	0.01	13.684	14.952	70.972	81.854	13.678	15.085	99.109	118.457

Table 3: **IQT on Amazon Stocks (*frequent observations*): Trends Across Models.** We test four base models, set $p_t = 0.9, \forall t$ to simulate seeing the true price 90% of the time, and report the mean across 5 seeds.

Metric	lr	AR		Prophet		Theta		Transformer	
		IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd
Marginal coverage	1.0	0.902	0.903	0.901	0.903	0.902	0.902	0.903	0.902
	0.1	0.894	0.894	0.886	0.886	0.8932	0.894	0.885	0.887
	0.01	0.870	0.873	0.797	0.805	0.871	0.873	0.786	0.798
Longest err seq	1.0	2.8	2.6	3	3.6	3.2	3.6	3.2	3
	0.1	3.6	3.8	7	6.6	5.4	5	10.6	9
	0.01	7	7	57.8	44.2	10	10	97.8	98.4
Avg set size	1.0	43.221	46.624	71.646	78.889	42.979	46.270	93.068	103.777
	0.1	17.663	17.982	54.235	52.265	17.687	18.121	69.700	66.835
	0.01	14.627	14.832	79.736	81.078	14.732	14.954	116.272	117.762




Figure 8: Amazon stock price data under frequent ($p_t = 0.9$) feedback. We show the prediction interval sizes for 1 seed, and coverage averaged over 5 seeds for Amazon stock price data under partial observations.

Table 4: **IQT on Google Stocks (infrequent observations): Trends Across Models.** We test four base models, set $p_t = 0.1, \forall t$ to simulate seeing the true price only 10% of the time, and report the mean across 5 seeds.

Metric	lr	AR		Prophet		Theta		Transformer	
		IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd
Marginal coverage	1.0	0.925	0.949	0.928	0.929	0.929	0.929	0.907	0.927
	0.1	0.910	0.925	0.897	0.928	0.906	0.929	0.793	0.907
	0.01	0.812	0.910	0.639	0.897	0.692	0.906	0.234	0.793
Longest err seq	1.0	11.4	27.8	21.4	42.4	17.2	29.2	23.2	29
	0.1	6	11.4	35.6	21.4	53.4	17.2	109.2	23.2
	0.01	9.2	6	69.6	35.6	103	53.4	873.4	109.2
Avg set size	1.0	87.834	911.476	134.392	1017.708	93.068	815.183	218.270	1445.050
	0.1	24.305	87.834	104.224	134.392	49.364	93.068	269.950	218.270
	0.01	15.070	24.305	54.653	104.224	28.910	49.364	109.453	269.950

Table 5: **IQT on Google Stocks (partial observations): Trends Across Models.** We test four base models, set $p_t = 0.5, \forall t$ to simulate seeing the true price only 50% of the time, and report the mean across 5 seeds.

Metric	lr	AR		Prophet		Theta		Transformer	
		IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd
Marginal coverage	1.0	0.909	0.909	0.902	0.908	0.906	0.903	0.902	0.899
	0.1	0.905	0.905	0.896	0.892	0.899	0.901	0.901	0.894
	0.01	0.891	0.901	0.858	0.887	0.864	0.890	0.661	0.809
Longest err seq	1.0	5.2	6.4	8	7.6	5.4	7.8	8.2	8.4
	0.1	4	4	13.8	11	13.8	7.4	16	13.2
	0.01	6	6	49.4	32.2	80.8	55.6	239.8	104.2
Avg set size	1.0	59.236	101.424	91.522	168.522	64.736	116.381	131.415	241.408
	0.1	22.021	25.670	71.972	63.493	36.04	34.893	156.157	106.766
	0.01	18.820	19.830	91.394	96.245	42.575	43.364	246.053	276.761

C.1 GOOGLE STOCK PRICE DATASET RESULTS

For the Google stock price dataset, Tables 4, 5, and 6 show the results for infrequent, partial, and frequent feedback. The differences between **IQT-pd** and **IQT-pi** for the Google dataset are very similar to the Amazon stock price dataset.

C.2 ELEC2 DATASET RESULTS

For the Elec2 dataset, Tables 7, 8, and 9 show the results for infrequent, partial, and frequent feedback.

D FURTHER EXAMINATION: SIMULATED EXPERIMENTS

Implementation Details. Both **EnsembleDagger** and **ConformalDagger** are implemented as 7-layer multilayer perceptions (with hidden sizes [64,128,472,512,256,64,42]). We use ReLU

Table 6: **IQT on Google Stocks (*frequent observations*): Trends Across Models.** We test four base models, set $p_t = 0.9, \forall t$ to simulate seeing the true price 90% of the time, and report the mean across 5 seeds.

Metric	lr	AR		Prophet		Theta		Transformer	
		IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd
Marginal coverage	1.0	0.903	0.902	0.902	0.903	0.902	0.902	0.902	0.901
	0.1	0.897	0.898	0.889	0.889	0.898	0.898	0.892	0.892
	0.01	0.895	0.897	0.880	0.881	0.889	0.891	0.791	0.808
Longest err seq	1.0	2.8	3	3.2	2.8	3.6	2.8	3	3.6
	0.1	4	4	8	7.8	6.8	6.6	10	9.4
	0.01	6	6	32.4	31.4	58	57.6	108.4	105
Avg set size	1.0	53.161	57.376	85.427	93.799	61.876	68.177	121.994	134.782
	0.1	21.068	21.564	59.732	57.680	31.262	31.240	106.240	99.712
	0.01	19.096	19.396	94.682	94.370	43.206	42.881	278.926	279.124

Table 7: **IQT on Elec2 Dataset (*infrequent observations*): Trends Across Models.** We test four base models, set $p_t = 0.1, \forall t$ to simulate seeing the true price only 10% of the time, and report the mean across 5 seeds.

Metric	lr	AR		Prophet		Theta		Transformer	
		IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd
Marginal coverage	1.0	0.918	0.908	0.900	0.934	0.909	0.949	0.903	0.904
	0.1	0.897	0.918	0.901	0.900	0.911	0.909	0.908	0.903
	0.01	0.809	0.897	0.665	0.901	0.857	0.911	0.663	0.908
Longest err seq	1.0	10.4	24.8	19.6	28.6	15.2	24.4	20.2	25.4
	0.1	4	10.4	19.2	19.6	9.2	15.2	10.8	20.2
	0.01	5.2	4	29.8	19.2	13	9.2	29.4	10.8
Avg set size	1.0	0.120	1.867	0.846	6.374	0.252	2.489	0.905	6.565
	0.1	0.091	0.120	0.528	0.846	0.085	0.252	0.556	0.905
	0.01	0.061	0.091	0.307	0.528	0.059	0.085	0.349	0.556

Table 8: **IQT on Elec2 Dataset (*partial observations*): Trends Across Models.** We test four base models, set $p_t = 0.5, \forall t$ to simulate seeing the true price only 50% of the time, and report the mean across 5 seeds.

Metric	lr	AR		Prophet		Theta		Transformer	
		IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd
Marginal coverage	1.0	0.909	0.905	0.898	0.908	0.907	0.909	0.902	0.899
	0.1	0.895	0.896	0.901	0.896	0.901	0.899	0.903	0.897
	0.01	0.892	0.894	0.865	0.887	0.899	0.901	0.886	0.903
Longest err seq	1.0	4.6	5.2	8.6	7.4	5.2	6.4	6.8	7.6
	0.1	3.8	5	7.2	7.4	5.2	4.8	7.8	7.6
	0.01	3.2	2.8	27	14.2	4.2	4.6	11	9.8
Avg set size	1.0	0.187	0.325	0.721	1.304	0.234	0.443	0.738	1.287
	0.1	0.093	0.101	0.475	0.4721	0.074	0.087	0.487	0.478
	0.01	0.087	0.088	0.481	0.500	0.068	0.070	0.497	0.517

Table 9: **IQT on Elec2 Dataset (*frequent observations*): Trends Across Models.** We test four base models, set $p_t = 0.9, \forall t$ to simulate seeing the true price 90% of the time, and report the mean across 5 seeds.

Metric	lr	AR		Prophet		Theta		Transformer	
		IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd	IQT-pi	IQT-pd
Marginal coverage	1.0	0.902	0.901	0.901	0.901	0.902	0.902	0.902	0.902
	0.1	0.900	0.902	0.897	0.898	0.897	0.899	0.897	0.898
	0.01	0.903	0.903	0.881	0.883	0.893	0.894	0.899	0.902
Longest err seq	1.0	2.6	3	3.8	3.4	2.6	2.8	3.4	3.2
	0.1	2.6	2.8	4.6	4.8	4.2	3.8	5.8	5.2
	0.01	2.4	2.8	15.6	15	4.6	4.8	10	10
Avg set size	1.0	0.179	0.192	0.677	0.734	0.219	0.243	0.670	0.724
	0.1	0.096	0.0977	0.446	0.437	0.070	0.071	0.444	0.438
	0.01	0.091	0.092	0.492	0.492	0.066	0.067	0.513	0.518

activations between each layer. We train using a learning rate of 0.001 and a batch size of 32. We train the initial policy for 200 iterations, and fine-tune between each interactive deployment episode for 100 iterations. The **EnsembleDagger** safety classifier is implemented as a 4-layer perceptron ([64,128,64,42]) with ReLU activation between each layer and we apply a sigmoid function on the output to classify the output.

A closer look at expert shift under $p_t = 0.2$ intermittent feedback . When the expert shift occurs at episode 5 (Figure 9), the human provides feedback occasionally to the learners which deviates from the predicted actions. **ConformalDagger** increases its calibrated uncertainty based on these human inputs, and once uncertainty exceeds the threshold, the probability for asking for help converges to 1, causing the robot to ask for more human feedback such that it can learn the new goal.

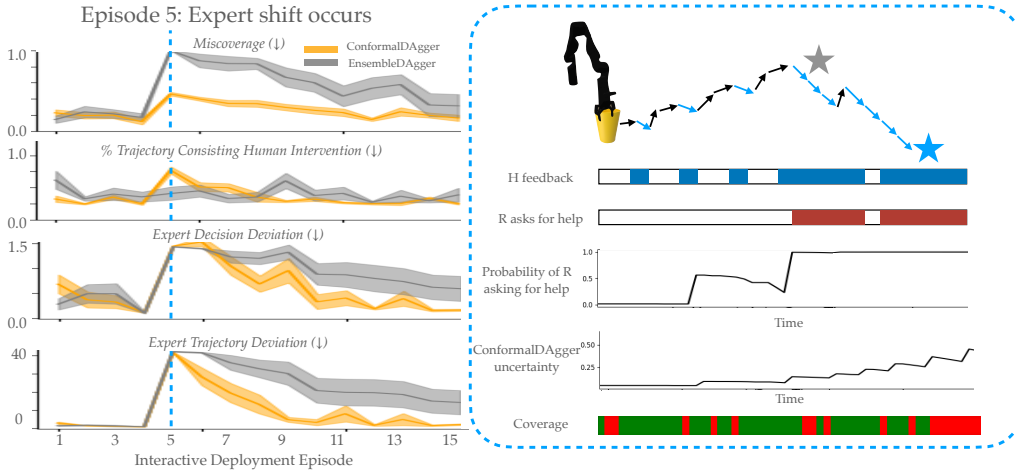


Figure 9: At episode 5, the **ConformalDagger** learner uncertainty enables the robot to ask for more human feedback to gather information about the shifted goal.

Simulated results under partial ($p_t = 0.5$) intermittent feedback . Under $p_t = 0.5$ intermittent feedback (Figure 10), **ConformalDagger** is able to detect the shift and drift more quickly than **EnsembleDagger**. We find that as the **ConformalDagger** algorithm under the stationary expert becomes more noisy has increased miscoverage. This is due to the expert labels decreasing the conformal parameters, q^{hi} , q^{lo} as the expert gives feedback, causing the intervals to become too small. Under the short time horizon, the miscoverage rate is higher than the desired level.

Simulated results under frequent ($p_t = 0.9$) intermittent feedback . Under $p_t = 0.9$ intermittent feedback (Figure 11), both algorithms are able to quickly adapt to expert shift and drift. This is because both algorithms during deployment are receiving human labels extremely frequently. Similar to partial feedback, as the **ConformalDagger** algorithm under the stationary expert becomes more noisy has increased miscoverage. **ConformalDagger** decreases the value of the conformal parameters, q^{hi} , q^{lo} as the expert gives feedback, causing the intervals to become too small, giving miscoverage higher than the desired level.

E IMPLEMENTATION AND TASK DETAILS: REAL EXPERIMENTS

E.1 LEARNING POLICY TRAINING DETAILS

We record robot and human actions at 15hz. The initial policy π_0^r is trained for 60K iterations with a batch size of 100. We use a weight decay of 1.0e-6, learning rate of 0.0001; for learning-rate scheduling, we used a cosine schedule with linear warmup (Nichol & Dhariwal, 2021). The number of training diffusion iterations is 100, and number of inference diffusion iterations is 16. The policy is trained on an NVIDIA RTX A6000 GPU.

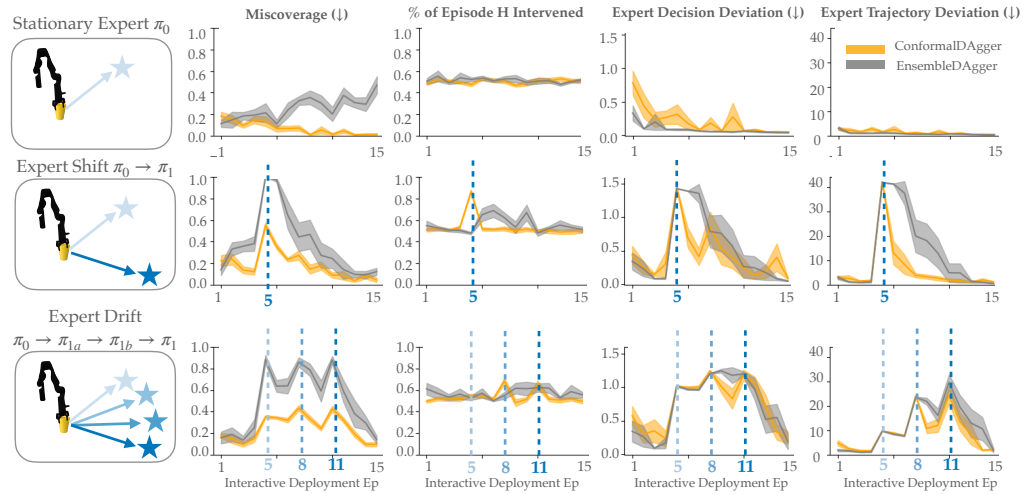


Figure 10: Under $p_t = 0.5$ intermittent feedback, **ConformalDagger** learns the shift and drift more quickly than **EnsembleDagger**.

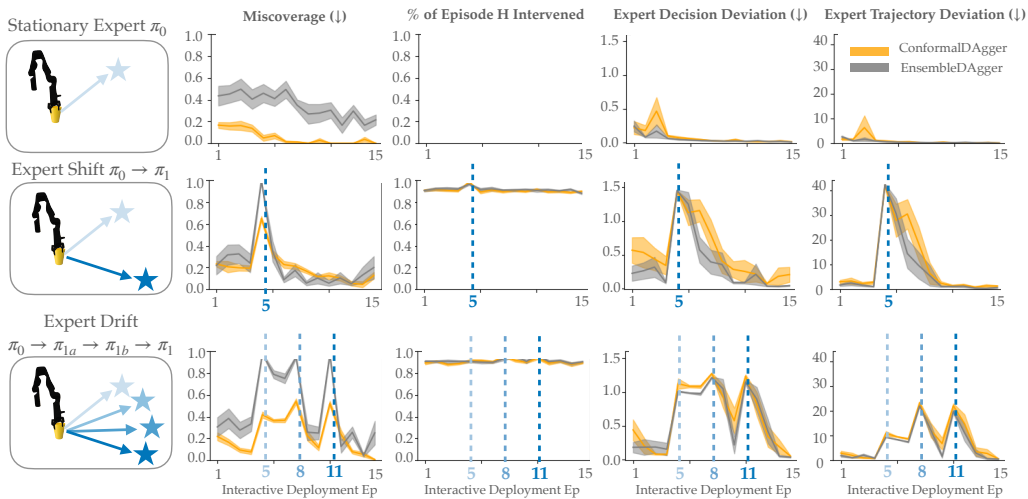


Figure 11: Under $p_t = 0.9$ intermittent feedback, both **ConformalDagger** and **EnsembleDagger** learn the shift and drift quickly.

A look at the real-world feedback request interface. Figure 12 shows the real-world interface for requesting help from the user during the interactive deployment episodes. The user provides teleoperated actions via a Meta Quest 3, and the robot’s uncertainty is displayed on a computer screen next to the robot. When the robot needs help, the robot pauses its execution and presents an alert notification on the screen.

Sponging Task Specifications. We present the task setup details (Figure 13) so that readers can also reproduce this task.



Figure 12: When the robot needs help, the robot pauses its execution and presents an alert notification on the screen.




Figure 13: Sponge task specifications.