

---

# Adaptive Learn-then-Test: Statistically Valid and Efficient Hyperparameter Selection

---

Matteo Zecchin<sup>1</sup> Sangwoo Park<sup>1</sup> Osvaldo Simeone<sup>1</sup>

## Abstract

We introduce adaptive learn-then-test (aLTT), an efficient hyperparameter selection procedure that provides finite-sample statistical guarantees on the population risk of AI models. Unlike the existing learn-then-test (LTT) technique, which relies on conventional p-value-based multiple hypothesis testing (MHT), aLTT implements sequential data-dependent MHT with early termination by leveraging e-processes. As a result, aLTT can reduce the number of testing rounds, making it particularly well-suited for scenarios in which testing is costly or presents safety risks. Apart from maintaining statistical validity, in applications such as online policy selection for offline reinforcement learning and prompt engineering, aLTT is shown to achieve the same performance as LTT while requiring only a fraction of the testing rounds.

## 1. Introduction

### 1.1. Context and Motivation

The safe and reliable deployment of AI applications, or apps for short, hinges on the possibility of certifying their performance (Seshia et al., 2022; Tegmark & Omohundro, 2023). Depending on the problem, this may require controlling the missed detection probability in medical imaging (Lu et al., 2022; Mehrtash et al., 2020), ensuring safety measures for control policies (Lindemann et al., 2023; Zecchin et al., 2024), or verifying the correctness of the answers given by a large language model (LLM) (Quach et al., 2023).

In practice, before deployment, AI apps can be often *calibrated* by selecting *hyperparameters* based on data set aside for this purpose or based on rounds of real-world testing.

---

<sup>1</sup>Centre for Intelligent Information Processing Systems, Department of Engineering, King’s College London, London, United Kingdom. Correspondence to: Sangwoo Park, Matteo Zecchin <sangwoo.park@kcl.ac.uk, matteo.l.zecchin@kcl.ac.uk>.

In the current scaling-centric era of large, Internet-scale, data sets (Xiao, 2024), hyperparameter optimization is typically modeled as a bandit problem in which the goal is the minimization of the training loss over a set of *discrete* configurations. The configurations may include prompt designs for the fine-tuning of large language models (LLMs) (Zhou et al., 2023; Quach et al., 2023; Schneider et al., 2024), architectural choices related to model scale (Nasir et al., 2024) or to the timing of early decisions (Schuster et al., 2022), or policy settings for offline reinforcement learning (Paine et al., 2020; Fujimoto & Gu, 2021).

The learn-then-test (LTT) method introduced by (Angelopoulos et al., 2021) has recently emerged as a formal framework to address hyperparameter selection over a discrete space. LTT treats the calibration task of *hyperparameter selection* as a *multiple hypothesis testing* (MHT) problem. Accordingly, it associates each hyperparameter in a candidate set to the null hypothesis that the hyperparameter does not meet a reliability requirement on the population risk. Hypotheses are tested using p-values, and the probability of mistakenly detecting a hyperparameter as reliable is guaranteed via *family-wise error rate* (FWER)-controlling statistical procedures (Keselman & Rogan, 1977). This way, LTT ensures finite-sample, high-probability guarantees on the population risk of the selected hyperparameters.

To prevent p-hacking (Head et al., 2015), LTT’s guarantees apply only to *non-adaptive* MHT procedures. However, related work on hyperparameter *optimization* has shown the significant benefits that can be accrued by *adaptive* exploration strategies that test hyperparameters sequentially in a data-driven manner (Swersky et al., 2014; Rakotoarison et al., 2024). This work aims at improving the data efficiency of LTT by leveraging recent advances in sequential MHT based on e-processes (Vovk & Wang, 2021; Waudby-Smith & Ramdas, 2024; Xu et al., 2021).

### 1.2. Related Work

Finite-sample statistical guarantees on inferential outputs can be obtained via *conformal prediction* methods (Shafer & Vovk, 2008; Angelopoulos et al., 2023) and, more generally, via conformal risk control and risk-controlling set-valued predictions (Angelopoulos et al., 2022; Bates et al., 2021).




Figure 1. An example application of aLTT to reliable prompt optimization (Zhou et al., 2023; Quach et al., 2023; Schneider et al., 2024). A set  $\Lambda$  of candidate prompts for a movie recommender is generated using an LLM and/or prior experience. Prompts serve as an example of a discrete set of hyperparameters to be optimized using aLTT. The goal is to select a subset  $\hat{\Lambda}^{\text{rel}} \subseteq \Lambda$  of prompts that guarantee a sufficiently high recommendation accuracy. To this end, aLTT applies a sequence of data-dependent testing rounds with adaptive termination. Specifically, at each testing round  $t$ , aLTT estimates the performance of a subset of hyperparameters  $T^t \subseteq \Lambda$  through held-out data or real-world testing. The subset to be tested is selected based on prior testing outcomes, and the process stops as soon as a sufficiently large reliable subset  $\hat{\Lambda}^{\text{rel}}$  is identified. An additional post-calibration selection step can be applied to choose a single hyperparameter  $\hat{\lambda}$  from the selected subset  $\hat{\Lambda}^{\text{rel}}$  based on users preferences.

These methods calibrate classification or regression models by setting a threshold hyperparameter based on held-out data to control the size of a prediction set. Calibrated predictors can be leveraged in predict-then-optimize control tasks to offer reliability guarantees (Vovk & Bendtsen, 2018; Lindemann et al., 2023; Zecchin et al., 2024).

Beyond prediction sets, the problem of calibrating an AI app via the selection of hyperparameters from a candidate pool has been addressed through the LTT framework (Angelopoulos et al., 2021). For example, as illustrated in Figure 1, for the problem of prompt engineering (Zhou et al., 2023), the initial set of candidate hyperparameters encompasses instruction prompt templates generated by an LLM and/or via prior experience. Leveraging p-value-based MHT via FWER-controlling procedures, LTT selects a subset of candidates that come with high-probability population risk guarantees.

FWER guarantees are often too conservative, potentially resulting in empty calibration sets. The *false discovery rate* (FDR) is an alternative and less strict criterion that is often preferred for MHT in fields such as genetics (van den Oord & Sullivan, 2003), neuroimaging (Genovese et al., 2002), online advertising (Berman & Van den Bulte, 2022), and finance (Harvey & Liu, 2020).

*E-values* have gained popularity in MHT due to their advantages over p-values (Vovk & Wang, 2021; Shafer & Vovk, 2019; Ramdas et al., 2020). Similarly to p-values, e-values measure the statistical plausibility of the null hypothesis. Specifically, an e-value can be thought of as a special type of p-value that has additional robustness properties. The key property of interest in this paper is that, unlike p-values, e-values can be readily combined to obtain *e-processes*, making it possible to devise *sequential* testing strategies with *anytime safety* (Wang & Ramdas, 2022; Ramdas et al., 2023). E-processes have been applied to problems such as sequential change detection (Shin et al., 2022), multiple

bandit testing (Xu et al., 2021), two-sample testing (Shekhar & Ramdas, 2023), and mean estimation of bounded random variables (Waudby-Smith & Ramdas, 2024).

*Hyperparameter optimization* is a vast field focused on the optimization of the hyperparameters of training algorithms, such as learning rate, weight decay, and dropout rate (Swersky et al., 2014; Pedregosa, 2016; Maclaurin et al., 2015; Lindauer et al., 2022). Hyperparameter optimization typically operates in continuous domains, and it can serve as a preliminary step for the identification of candidate hyperparameters. While hyperparameter optimization does not provide statistical guarantees on the population risk, the goal of hyperparameter selection methods such as LTT is to formally test a subset of candidate hyperparameters for statistical validity.

### 1.3. Main Contributions

The main contributions of this paper are as follows.

- We introduce adaptive LTT (aLTT), a data-efficient hyperparameter selection method that provides finite-sample guarantees on the population risk of AI apps. The main technical underpinning of aLTT is *e-process-based MHT*, which supports statistical validity while enabling data-dependent sequential testing (Xu et al., 2021). Unlike LTT, as illustrated in Figure 1, aLTT adaptively tests subsets of hyperparameters that are chosen based on the evidence accumulated in the previous rounds, allowing also for the early termination of the calibration process. aLTT guarantees rigorous control over FWER and FDR, while significantly reducing the number of testing rounds.
- We study two practical scenarios requiring hyperparameter selection, namely online policy selection for offline reinforcement learning (Fujimoto & Gu, 2021) and automated prompt engineering (Zhou et al., 2023). In both cases, aLTT is shown to deliver reliable and effective hyperparameters using only a small fraction of the testing rounds required

by LTT.

## 2. Problem Definition

### 2.1. Setting

Let  $\mathcal{M}_\lambda$  be an AI app whose operation is determined by a vector of hyperparameters  $\lambda$ . The performance of a hyperparameter vector  $\lambda$  when tested at input data  $Z$  is measured by a risk function  $R(\lambda, Z) \in [0, 1]$ . Accordingly, the *population risk* with respect to an unknown data distribution  $P_Z$  is defined as

$$R(\lambda) = \mathbb{E}_{P_Z}[R(\lambda, Z)]. \quad (1)$$

As illustrated in Figure 1, for a given discrete subset  $\Lambda = \{\lambda_1, \dots, \lambda_N\}$  of hyperparameters and a user-specified reliability level  $\alpha \in [0, 1]$ , we aim at determining the subset of hyperparameters in set  $\Lambda$  that conforms with the required reliability level  $\alpha$ , i.e.,

$$\Lambda^{\text{rel}} = \{\lambda \in \Lambda : R(\lambda) \leq \alpha\}. \quad (2)$$

The complementary set, comprising unreliable hyperparameters, is accordingly defined as

$$\Lambda^{\text{unrel}} = \Lambda \setminus \Lambda^{\text{rel}} = \{\lambda \in \Lambda : R(\lambda) > \alpha\}. \quad (3)$$

Since identifying the entire set  $\Lambda^{\text{rel}}$  in (2) is impossible owing to the lack of knowledge about the data distribution  $P_Z$ , the goal is producing a subset of hyperparameters  $\hat{\Lambda}^{\text{rel}} \subseteq \Lambda$  that contains as many reliable hyperparameters from subset  $\Lambda^{\text{rel}}$  as possible while controlling the number of unreliable hyperparameters from subset  $\Lambda^{\text{unrel}}$  mistakenly included in subset  $\hat{\Lambda}^{\text{rel}}$ .

### 2.2. Performance Criteria

**Definition 2.1** ( $(\alpha, \delta)$ -FWER-controlling set). For a given reliability level  $\alpha \in [0, 1]$  and an error level  $\delta \in [0, 1]$ , a hyperparameter subset  $\hat{\Lambda}^{\text{rel}} \subseteq \Lambda$  is  $(\alpha, \delta)$ -FWER-controlling set if it satisfies the requirement

$$\text{FWER}(\hat{\Lambda}^{\text{rel}}) := \Pr \left[ |\Lambda^{\text{unrel}} \cap \hat{\Lambda}^{\text{rel}}| \geq 1 \right] \leq \delta. \quad (4)$$

where the probability is evaluated with respect to the distribution of the subset  $\hat{\Lambda}^{\text{rel}}$ .

The FWER guarantee (4) imposes that the probability that the calibration set  $\hat{\Lambda}^{\text{rel}}$  contains an unreliable hyperparameter is bounded by  $\delta$ .

**Definition 2.2** ( $(\alpha, \delta)$ -FDR-controlling set). For a given reliability level  $\alpha \in [0, 1]$  and an error level  $\delta \in [0, 1]$ , a hyperparameter subset  $\hat{\Lambda}^{\text{rel}} \subseteq \Lambda$  is  $(\alpha, \delta)$ -FDR-controlling set if it satisfies the inequality

$$\text{FDR}(\hat{\Lambda}^{\text{rel}}) := \mathbb{E} \left[ \frac{|\Lambda^{\text{unrel}} \cap \hat{\Lambda}^{\text{rel}}|}{|\hat{\Lambda}^{\text{rel}}|} \middle| |\hat{\Lambda}^{\text{rel}}| \geq 1 \right] \leq \delta, \quad (5)$$

with the average evaluated with respect to the distribution of the subset  $\hat{\Lambda}^{\text{rel}}$ .

Accordingly, a testing procedure that outputs  $(\alpha, \delta)$ -FDR-controlling sets guarantees that the expected fraction of unreliable hyperparameters in the predicted set  $\hat{\Lambda}^{\text{rel}}$  is bounded by  $\delta$ . Thus, ensuring the  $(\alpha, \delta)$ -FWER condition automatically also guarantees the  $(\alpha, \delta)$ -FDR requirement.

Since any FWER or FDR level can be trivially satisfied by a procedure that returns the empty set  $\hat{\Lambda}^{\text{rel}} = \emptyset$ , it is important to gauge the informativeness of the testing procedure via the *true positive rate* (TPR), which corresponds to the expected fraction of reliable models in the predicted set  $\hat{\Lambda}^{\text{rel}}$ , i.e.,

$$\text{TPR}(\hat{\Lambda}^{\text{rel}}) = \mathbb{E} \left[ \frac{|\Lambda^{\text{rel}} \cap \hat{\Lambda}^{\text{rel}}|}{|\hat{\Lambda}^{\text{rel}}|} \right]. \quad (6)$$

### 2.3. Sequential and Adaptive Hyperparameter Selection

To produce the estimated subset of reliable hyperparameters,  $\hat{\Lambda}^{\text{rel}}$ , we adopt a general sequential testing procedure that, at each round  $t \geq 1$ , operates as follows.

① *Hyperparameter subset selection*: A subset of hyperparameters  $\mathcal{I}^t \subseteq \Lambda$  is selected for testing.

② *Testing*: Empirical risk estimates  $\mathcal{R}^t = \{R(\lambda_i, Z_i^t)\}_{\lambda_i \in \mathcal{I}^t}$  are obtained, one for each candidate hyperparameter  $\lambda_i$  in the selected subset  $\mathcal{I}^t$ , using held-out data or real-world testing. The random variable  $Z_i^t \sim P_Z$  describes the data used to test hyperparameter  $\lambda_i$  at round  $t$ . The random variables  $\{Z_i^t\}_{i \in \mathcal{I}^t}$  can be arbitrarily dependent, and thus one may reuse the same data to test all hyperparameters  $\lambda \in \mathcal{I}^t$ .

③ *Evidence update*: Evidence accumulated up to time  $t$ , including both the observed risks and the subset of queried models, is updated as  $\mathcal{D}^t = \mathcal{D}^{t-1} \cup \{(\mathcal{I}^t, \mathcal{R}^t)\}$ .

The testing procedure outlined above is fully specified by the tuple  $\Pi = (\{\mathcal{Q}^t\}_{t \geq 1}, \mathcal{A}, T)$ , encompassing a family of acquisition policies  $\{\mathcal{Q}^t\}_{t \geq 1}$ , a decision rule  $\mathcal{A}$ , and a calibration horizon  $T$ , which are defined as follows.

- *Acquisition policy*: At each round  $t$ , the acquisition policy  $\mathcal{Q}^t$  determines the hyperparameters  $\mathcal{I}^t$  to be tested at step ①. If the policy  $\mathcal{Q}^t$  uses the evidence  $\mathcal{D}^{t-1}$  to select the hyperparameter set  $\mathcal{I}^t$ , it is said to be *adaptive*; otherwise, it is *non-adaptive*. Both *adaptive* and *non-adaptive* acquisition policies can incorporate prior knowledge, which we denote as  $\mathcal{D}^0$ .

- *Decision rule*: The decision rule  $\mathcal{A}$  uses the evidence  $\mathcal{D}^T$  available at the end of the last calibration round  $T$  to

produce the estimated set  $\hat{\Lambda}^{\text{rel}}$  of reliable hyperparameters.

- *Number of calibration rounds:* The number of calibration rounds  $T$ , is said to be *adaptive*, if the stopping condition  $T = t$  is determined by the evidence  $\mathcal{D}^{t-1}$ . Otherwise, when it is predetermined based solely on prior knowledge  $\mathcal{D}^0$ , the calibration horizon  $T$  is said to be *non-adaptive*.

### 3. (Non-Adaptive) Learn-then-Test

In this section, we review LTT, a *non-adaptive* hyperparameter selection procedure devised to meet the  $(\alpha, \delta)$ -FWER guarantee (Angelopoulos et al., 2021). LTT associates to each hyperparameter  $\lambda_i \in \Lambda$  the null hypothesis

$$\mathcal{H}_i : R(\lambda_i) > \alpha \quad (7)$$

that the population risk  $R(\lambda_i)$  in (1) violates the target reliability level  $\alpha$ . For each null hypothesis  $\mathcal{H}_i$  a p-value  $P_i$  is a non-negative random variable that satisfies the inequality

$$\Pr[P_i \leq x | \mathcal{H}_i] \leq x \quad (8)$$

for every  $x \in [0, 1]$ . By the definition (8) a p-value  $P_i$  provides evidence for the validity of hypothesis  $\mathcal{H}_i$ . This is in the sense that a small value of  $P_i$  is unlikely to occur if  $\mathcal{H}_i$  is true.

Given a pre-specified calibration horizon  $T$  and a non-adaptive acquisition policy  $\{Q^t\}_{t \geq 1}$ , LTT queries at each round  $t \geq 1$  the subset of hyperparameters  $Q^t(\mathcal{D}^0) = \mathcal{I}^t$ , obtaining the corresponding risk estimates  $\mathcal{R}^t = \{R(\lambda_i, Z_i^t)\}_{\lambda_i \in \mathcal{I}^t}$ .

LTT uses the accumulated evidence at the end of the testing process,  $\mathcal{D}^T$ , to compute a valid p-value for the null hypothesis (7) using, for instance, the Hoeffding-Bentkus concentration inequality introduced in (Bates et al., 2021). Based on the collection of p-values  $\mathcal{P} = \{P_i\}_{i=1}^N$ , LTT selects a subset of hyperparameters  $\hat{\Lambda}^{\text{LTT}}$  using a FWER-controlling algorithm  $\mathcal{A}^{\text{FWER}}(\mathcal{P})$ . A variant of LTT that is FDR-controlling can be readily obtained by using an FDR-controlling selection rule  $\mathcal{A}^{\text{FDR}}(\mathcal{P})$ . Examples of FWER and FDR controlling procedures are provided in the Supplementary Material.

### 4. Adaptive Learn-Then-Test

In this section, we introduce *adaptive LTT* (aLTT), a hyperparameter selection scheme that supports *adaptive* acquisition policies and an *adaptive* number of calibration rounds. The algorithmic description of aLTT is given in Algorithm 1.

#### 4.1. Hypothesis Testing via E-Processes

The proposed aLTT scheme applies MHT based on e-values and e-processes (Shafer, 2021; Ramdas et al., 2023). For

each null hypothesis  $\mathcal{H}_i$  in (7), an *e-value*  $E_i$  is a non-negative random variable with an expectation no larger than 1 when  $\mathcal{H}_i$  is true, i.e.,

$$\mathbb{E}[E_i | \mathcal{H}_i] \leq 1. \quad (9)$$

By Markov's inequality, an e-value  $E_i$  can be turned into a p-value  $P_i$  as  $P_i = 1/E_i$ , since the inequality (8) is satisfied as

$$\Pr\left[\frac{1}{E_i} \leq x \mid \mathcal{H}_i\right] \leq \mathbb{E}[E_i | \mathcal{H}_i] x \leq x, \quad \forall x \in [0, 1]. \quad (10)$$

For each null hypothesis  $\mathcal{H}_i$ , given an observation  $Z$  and a fixed  $\mu \in (0, 1/(1 - \alpha))$ , a valid e-value is given by (Waudby-Smith & Ramdas, 2024)

$$E_i = (1 + \mu(\alpha - R(\lambda_i, Z))). \quad (11)$$

The e-value (11) has the interpretation of *wealth growth* in a *betting* setting. Accordingly, one can think of parameter  $\mu$  in (11) as the amount of the current wealth that the gambler bets on the hypothesis  $\mathcal{H}_i$  being false, i.e., on the validity of the assumption  $R(\lambda_i) \leq \alpha$  that the hyperparameter  $\lambda_i$  is reliable. In fact, if  $\mu > 0$ , when  $R(\lambda_i, Z) \leq \alpha$ , the gambler's wealth in (11) increases; while, when  $R(\lambda_i, Z) > \alpha$ , the quantity (9) the gambler's wealth (11) decreases.

An e-process for hypothesis  $\mathcal{H}_i$  is a sequence of random variables  $\{E_i^t\}_{t \geq 1}$  such that, for any stopping time  $T$ , which may depend on all previously collected evidence, the random variable  $E_i^T$  is a valid e-value. Using the e-value (11), considering the general iterative testing framework in Section 2.3, an e-process for the null hypothesis  $\mathcal{H}_i$  in (7) can be obtained as the product

$$E_i^t = \prod_{\tau \leq t: \lambda_i \in \mathcal{I}^\tau} (1 + \mu_i^\tau(\alpha - R(\lambda_i, Z_i^\tau))), \quad (12)$$

where the betting strategy  $\mu_i^t \in (0, 1/(1 - \alpha))$  can be optimized as a function of the past risk estimates  $\{R(\lambda_i, Z_i^\tau)\}_{\tau < t}$  and  $E_i^0 = 1$ . Based on the discussion above, the e-process (12) represents the wealth accumulated up to time  $t$  by a gambler making sequential bets  $\{\mu_i^\tau\}_{\tau < t}$  (Shafer & Vovk, 2019; Waudby-Smith & Ramdas, 2024). As such, the gambler's wealth up to time  $t$ ,  $E_i^t$ , provides evidence against the null hypothesis that the hyperparameter  $\lambda_i$  is unreliable.

In a similar way, an *anytime-valid p-value* for hypothesis  $\mathcal{H}_i$  is a sequence of random variables  $\{P_i^t\}_{t \geq 1}$  such that, for any stopping time  $T$ , the random variable  $P_i^T$  is a valid p-value. Given an e-process  $\{E_i^t\}_{t \geq 1}$  for hypothesis  $\mathcal{H}_i$ , the sequence

$$P_i^t = \frac{1}{\max_{\tau \leq t} E_i^\tau} \quad (13)$$

is an anytime-valid p-value for the hypothesis  $\mathcal{H}_i$  (Ramdas et al., 2023).



**Algorithm 1** Adaptive Learn-Then-Test (aLTT)

---

**Require:** Candidate hyperparameters  $\Lambda$ , prior knowledge  $\mathcal{D}^0$ , reliability level  $\alpha$ , error tolerance level  $\delta$ , acquisition policy  $\{\mathcal{Q}^t\}_{t \geq 1}$ , FWER/FDR-controlling selection rule  $\mathcal{A}^{\text{FWER}}(\cdot)/\mathcal{A}^{\text{FDR}}(\cdot)$ , betting strategy  $\{\mu_i^t\}_{t \geq 1, i=1, \dots, N}$ , maximum number of iterations  $t_{\max}$  and minimal hyperparameter set cardinality  $d$

**Ensure:**  $(\alpha, \delta)$ -FWER/FDR-controlling hyperparameter set  $\hat{\Lambda}^{\text{aLTT}, T}$

$t \leftarrow 1$

**while**  $t \leq t_{\max} \wedge |\hat{\Lambda}^{\text{aLTT}, t}| \leq d$  **do**

Select hyperparameters  $\mathcal{I}^t = \mathcal{Q}^t(\mathcal{E}^{t-1})$  and receive risk estimates  $\mathcal{R}^t = \{R(\lambda_i, Z_i^t)\}_{\lambda_i \in \mathcal{I}^t}$

Update evidence  $\mathcal{D}^t = \mathcal{D}^{t-1} \cup \{(\mathcal{I}^t, \mathcal{R}^t)\}$  and e-processes  $\mathcal{E}^t$  as in (14)

**if** FWER-control **then**

Compute p-values  $\mathcal{P}^t$  as in (13)

$\hat{\Lambda}^{\text{aLTT}, t} \leftarrow \mathcal{A}^{\text{FWER}}(\mathcal{P}^t)$

**else if** FDR-control **then**

$\hat{\Lambda}^{\text{aLTT}, t} \leftarrow \mathcal{A}^{\text{FDR}}(\mathcal{E}^t)$

**end if**

$t \leftarrow t + 1$

**end while**

**Return:**  $\hat{\Lambda}^{\text{aLTT}, T}$

---

## 4.2. Adaptive Acquisition Policy

aLTT applies an *adaptive* acquisition policy  $\{\mathcal{Q}^t\}_{t \geq 1}$  and an *adaptive* calibration horizon  $T$ . Specifically, at each calibration round  $t \geq 1$ , aLTT's acquisition policy  $\mathcal{Q}^t$  uses the e-processes  $\mathcal{E}^{t-1} = \{E_i^{t-1}\}_{i=1}^N$  in (12) to choose which subset of hyperparameters,  $\mathcal{I}^t$ , to test next. Examples of acquisition functions  $\mathcal{I}^t = \mathcal{Q}^t(\mathcal{E}^{t-1})$  will be provided in the next section.

For the selected hyperparameters in set  $\mathcal{I}^t$ , aLTT obtains the risk estimates  $\mathcal{R}^t = \{R(\lambda_i, Z_i^t)\}_{\lambda_i \in \mathcal{I}^t}$  and updates the associated e-processes using the recursive formula (12), i.e.,

$$E_i^t = \begin{cases} (1 + \mu_i^t(\alpha - R(\lambda_i, Z_i^t)))E_i^{t-1}, & \text{if } \lambda_i \in \mathcal{I}^t \\ E_i^{t-1}, & \text{otherwise.} \end{cases} \quad (14)$$

With this information, a prediction set  $\hat{\Lambda}^{\text{aLTT}, t}$  is evaluated by employing either an FWER-controlling method  $\mathcal{A}^{\text{FWER}}(\mathcal{P}^t)$  based on the p-values  $\mathcal{P}^t = \{P_i^t\}_{i=1}^N$  in (13); or an FDR-controlling procedure  $\mathcal{A}^{\text{FDR}}(\mathcal{E}^t)$ , such as the e-Benjamini-Hochberg (eBH) method, reviewed in the Supplementary Material (Wang & Ramdas, 2022), using directly the e-values (14).

aLTT terminates the calibration procedure whenever there are at least  $d$  hyperparameters in set  $\hat{\Lambda}^{\text{aLTT}, t}$ , i.e.,  $|\hat{\Lambda}^{\text{aLTT}, t}| \geq d$ , or a maximum number of iterations  $t_{\max}$  have been reached. This allows aLTT to stop the data acquisition phase early when a sufficiently large number of reliable hyperparameters have been identified. .

## 4.3. Hyperparameter Subset Selection

At the end of the calibration process, aLTT uses the current e-processes  $\mathcal{E}^T$  to generate the final prediction set  $\hat{\Lambda}^{\text{aLTT}, T}$ . By the anytime validity properties explained in Section 4.1, if an FWER-controlling method  $\mathcal{A}^{\text{FWER}}(\mathcal{P}^T)$  is used, the resulting set  $\hat{\Lambda}^{\text{aLTT}, T}$  is  $(\alpha, \delta)$ -FWER-controlling; while if an FDR-controlling method is used, the resulting set  $\hat{\Lambda}^{\text{aLTT}, T} = \mathcal{A}^{\text{FDR}}(\mathcal{E}^T)$  is  $(\alpha, \delta)$ -FDR-controlling.

**Theorem 4.1.** *Given a hyperparameter set  $\Lambda$ , a reliability level  $\alpha \in [0, 1]$ , and an error level  $\delta \in [0, 1]$ , aLTT with an FWER or FDR-controlling selection rule returns a final prediction set  $\hat{\Lambda}^{\text{aLTT}, T}$  that satisfies  $(\alpha, \delta)$ -FWER control or  $(\alpha, \delta)$ -FDR control, respectively.*

## 5. Applications

### 5.1. Online Policy Selection for Offline Reinforcement Learning

Offline reinforcement learning enables the training of control policies based on a fixed data set collected by using a possibly unknown behavior policy, without any online interaction with the environment (Levine et al., 2020). However, the estimate of the performance of the trained policies obtained from offline data can differ substantially from the actual performance in the real world. This makes it practically essential to validate the policies' performance via online interaction with the environment (Paine et al., 2020; Liu et al., 2023).

To reduce the cost and potential harm of online validation procedures, the number of online interactions of the pre-trained candidate policies with the real world must be kept to a minimum (Garcia & Fernández, 2015). To this end, in




Figure 2. True positive rate of LTT and aLTT with  $\epsilon$ -greedy acquisition policy for  $\epsilon \in \{0.25, 0.5, 0.75, 0.95\}$  and non-adaptive acquisition. On the left panel, the prediction sets satisfy FWER control while on the right FDR control. In both cases, the tolerance level is  $\delta = 0.1$ .

this subsection, we investigate the potential benefits of the proposed aLTT framework as a means to select a subset of candidate policies that enjoy performance guarantees with respect to the real-world environment.

#### 5.1.1. PROBLEM DEFINITION

We assume a standard Markov decision process (MDP)  $\mathcal{E} = \{\mathcal{S}, \mathcal{A}, P_{s'|s,a}, P_{r|s,a}\}$  specified by a state space  $\mathcal{S}$ ; an action space  $\mathcal{A}$ ; a transition kernel  $P_{s'|s,a}$ , defining the conditional distribution of the next state  $s' \in \mathcal{S}$  given the current state  $s \in \mathcal{S}$  and action  $a \in \mathcal{A}$ ; and a conditional reward distribution  $P_{r|s,a}$  given state  $s \in \mathcal{S}$  and action  $a \in \mathcal{A}$ . We assume that the reward  $r$  is bounded and normalized in the  $[0, 1]$  interval.

We are given a set of pre-trained control policies  $\Pi = \{\pi_1, \dots, \pi_N\}$ , mapping an observed state  $s \in \mathcal{S}$  to the random action  $a \sim \pi_i(s) \in \mathcal{A}$ . Each policy  $\pi_i$  is identified by a hyperparameter  $\lambda_i$ . The goal is to select a subset  $\hat{\Lambda}^{\text{rel}} \subset \Lambda$  of hyperparameters that yield reliable policies by using a limited amount of interactions with the real world.

Reliability is measured via the cumulative reward obtained by a policy  $\lambda$  on the MDP  $\mathcal{E}$ , which is defined as

$$R(\lambda, Z) = \frac{1}{K} \sum_{k=1}^K r^k, \quad (15)$$

where  $K$  is the length of the episode, and the per-episode random variable  $Z$  encompasses the initial state  $s^1 \sim P_{s^1}$  along with the sequence  $a^1, r^1, s^2, a^2, r^2, \dots, s^K, a^K, r^K$  with actions  $a^k \sim \pi_i(s^k)$ , rewards  $r^k \sim P_{r^k|s^k, a^k}$  and MDP transitions  $s^{k+1} \sim P_{s^{k+1}|s^k, a^k}$ . For a user-specified reliability level  $\alpha \in (0, 1)$ , the subset of reliable policies  $\Lambda^{\text{rel}} \subseteq \Lambda$  includes all policies in  $\Lambda$  with average cumulative reward larger than  $\alpha$ ,  $\Lambda^{\text{rel}} = \{\lambda \in \Lambda : R(\lambda) = \mathbb{E}_{P_Z}[R(\lambda, Z)] >$

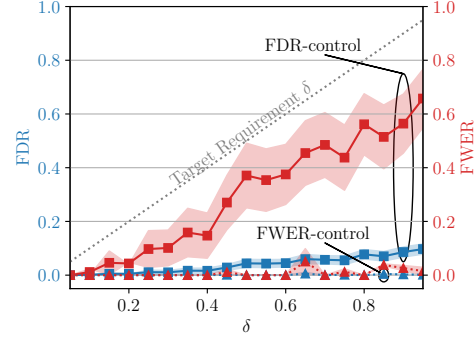


Figure 3. Comparison of the FWER and FDR levels obtained by aLTT under FDR-control (solid lines) and FWER-control (dashed lines) for different maximum tolerated error (FWER or FDR) levels  $\delta$ .

$\alpha\}$ , while the complementary set  $\Lambda^{\text{unrel}} = \Lambda \setminus \Lambda^{\text{rel}}$  includes the policies that do not satisfy the given reliability requirement.

Control policies are tested sequentially by following the procedure described in Section 2.3, such that at each calibration round  $t$  a policy  $\lambda_i^t \in \Lambda$  is tested using online interactions with the MDP  $\mathcal{E}$  to obtain the episodic reward value  $R_i^t = R(\lambda_i^t, Z_i^t)$  in (15). For an error threshold  $\delta \in [0, 1]$ , the goal of reliable online policy selection is to return a prediction set  $\hat{\Lambda}^{\text{rel}} \subseteq \Lambda$  that is either  $(\alpha, \delta)$ -FDR controlling or  $(\alpha, \delta)$ -FWER controlling with a TPR that is as large as possible.

#### 5.1.2. RESULTS

In our experiments, we consider the Half Cheetah control problem from the OpenAI Gym MuJoCo tasks (Todorov et al., 2012) and use control policies obtained via the offline reinforcement learning algorithm TD3+BC (Fujimoto & Gu, 2021). The TD3+BC algorithm leverages an offline data set  $\mathcal{D}$  to optimize policies by maximizing the standard deterministic policy gradient objective (Silver et al., 2014), combined with a behavioral cloning regularization term, whose strength is controlled by a hyperparameter  $\lambda$  in [Eq. 5](Silver et al., 2014). We produce  $N = 20$  different control policies by setting the hyperparameter  $\lambda$  in the TD3+BC training objective on an evenly spaced grid in the interval  $[0.25, 5]$ . Unless stated otherwise, we consider a target reliability  $\alpha = 0.57$  and a target FDR requirement  $\delta = 0.1$ .

We evaluate aLTT with an  $\epsilon$ -greedy acquisition policy  $Q^t$  that, at every calibration round  $t$ , with probability  $1 - \epsilon$ , selects the hyperparameter  $\lambda_i^t$  not included in  $\hat{\Lambda}^{\text{aLTT}, t}$  that is associated with the largest e-process value; otherwise, it picks uniformly at random a hyperparameter not in  $\hat{\Lambda}^{\text{aLTT}, t}$ . For reference, we also consider aLTT with a non-adaptive acquisition policy that, at each round  $t$ , picks uniformly at random the hyperparameter to be tested regardless of the

prediction outcome  $\hat{\Lambda}^{\text{aLTT}, t}$  and the e-process values. As a benchmark, we implement LTT with a random uniform acquisition policy and p-values obtained from the e-processes as in (13). Recall that LTT produces a decision at the end of the calibration process, here at round  $T = 5000$ .

Finally, the value of the parameter  $\mu_i^t$  in aLTT is set by following the *approximate growth rate adaptive to the particular alternative* (aGRAPA) betting strategy in (Waudby-Smith & Ramdas, 2024) with other adaptive and non-adaptive betting strategies evaluated in the Supplementary Material.

In Figure 2, we compare the TPR of LTT and aLTT as a function of the calibration round  $t$ . We target FWER control on the left and FDR control on the right. By construction, LTT returns uninformative hyperparameter sets up until the termination of the testing procedure. The performance of LTT is the same as aLTT with a non-adaptive acquisition policy, i.e. with  $\epsilon = 1$ , at  $t = T$ . aLTT with an  $\epsilon$ -greedy acquisition function can benefit from the accumulated evidence to adaptively determine the models to test next. As  $\epsilon$  decreases, and thus the acquisition policy becomes increasingly driven by evidence, the TPR increases from 0.32 to 0.85 in the case of FWER control and from 0.4 to 0.85 in the case of FDR control. Finally, we note that the TPR of the schemes under FDR control is larger than that obtained under FWER control, reflecting the stricter reliability requirement of FWER control.

In Figure 3, we report the FWER and FDR of aLTT with  $\epsilon = 0.25$  as a function of the tolerated error level  $\delta$ . As the error level  $\delta$  increases, the empirical FWER and FDR increase accordingly, remaining below the maximum target level  $\delta$ . However, since FWER control is a stricter requirement, the aLTT prediction set obtained under FWER control delivers lower FDR and FWER levels as compared to aLTT with FDR control.

## 5.2. Reliable Automated Prompt Engineering

Prompt engineering focuses on designing and refining input instructions for LLMs (Reynolds & McDonell, 2021; Shin et al., 2020). Recent studies have demonstrated the effectiveness of methods that automate this search using LLMs as prompt generators (Zhou et al., 2023; Zhang et al., 2023). However, while LLMs are capable of producing high-quality instruction templates, at a level comparable to human annotators, supervision and testing remain essential to filter out poorly performing prompts.

In this section, we propose applying aLTT to automated prompt engineering (Zhou et al., 2023) to generate instructions with statistical performance guarantees. As illustrated in Figure 1, given a set  $\Lambda$  of candidate instructions generated by an LLM, we use aLTT to sequentially test the instruc-

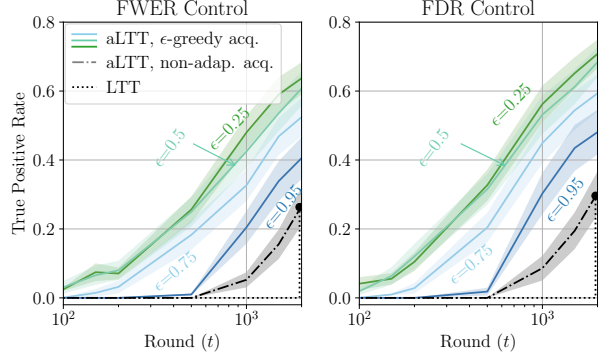


Figure 4. True positive rate of LTT and aLTT with  $\epsilon$ -greedy acquisition policy for  $\epsilon \in \{0.25, 0.5, 0.75, 0.95\}$  and non-adaptive acquisition. On the left panel, the prediction sets satisfy FWER control, while on the right panel they meet FDR requirements. In both cases, the tolerance level is  $\delta = 0.1$ . Results are averaged over the tasks in (Honovich et al., 2022) that yield non-empty reliable prompts set.

tions and identify a sufficiently large number of prompts that meet user-defined reliability requirements.

### 5.2.1. PROBLEM DEFINITION

Consider a target LLM  $f(\cdot)$  that, when fed the concatenation  $[\lambda, X]$  of an instruction  $\lambda$  and data  $X$ , produces an output text  $f([\lambda, X])$ . For example, in the task of movie recommendation illustrated in Figure 1, each prompt  $\lambda$  corresponds to an instruction to generate movie titles similar to those provided in the input list  $X$ ; while the output  $f([\lambda, X])$  represents a list of recommended movies. Following (Zhou et al., 2023; Zhang et al., 2023), the candidate set  $\Lambda$  of instructions, corresponding to the hyperparameters to be tested, is generated using a separate LLM. In particular, we use the Llama3 8B Instruct LLM (Dubey et al., 2024) as the target model  $f(\cdot)$ , while the initial instruction set  $\Lambda$  is generated through the Llama3.3 70B Instruct model (meta, 2025) using the forward generation mode detailed in (Zhou et al., 2023).

For each instruction  $\lambda$ , the ability of the prompted model  $f([\lambda, X])$  to generate high-quality outputs for a test datum  $Z = (X, Y)$  is measured by a task-dependent loss function  $R(\lambda, Z) = \ell(f([\lambda, X]), Y) \in [0, 1]$ . Focusing on tasks from the instruction induction data set (Honovich et al., 2022), we specifically adopt the 0-1 loss tailored to the given task (Zhou et al., 2023). The goal is to determine a subset  $\Lambda^{\text{rel}} \subseteq \Lambda$  that complies with the condition (2) with a target execution error  $\alpha = 0.2$ , and FWER and FDR requirements given by the tolerance parameter  $\delta = 0.1$ .

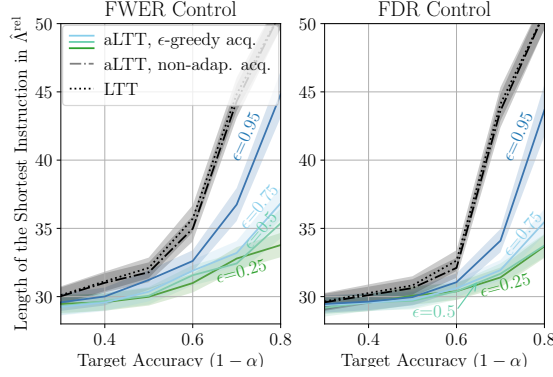


Figure 5. Length of the shortest instruction in the predicted set of reliable hyperparameter  $\hat{\Lambda}^{\text{rel}}$  returned by LTT and aLTT. Instructions are tested under different accuracy requirements and with a fixed testing budget  $T = 2000$ .

### 5.2.2. RESULTS

For the outlined prompt engineering problem, we evaluate aLTT using an  $\epsilon$ -greedy acquisition strategy with  $\epsilon \in \{0.25, 0.5, 0.75, 0.95\}$ , as well as a non-adaptive acquisition strategy whereby the instruction to be tested is selected randomly and independently from prior testing rounds. All schemes employ an online Newton step betting strategy (Waudby-Smith & Ramdas, 2024) for the e-value (12).

In Figure 4, we report the TPR of LTT and aLTT as a function of the number of rounds  $t$ . Results are averaged over tasks that yield a final non-empty set  $\hat{\Lambda}^{\text{rel}}$ . Plots for individual tasks are included in Appendix B.3. Under both FDR and FWER control, aLTT reduces the number of testing rounds, or equivalently the number of LLM executions, required to achieve a given TPR. In particular, LTT produces a non-empty set only at the end of the testing process, attaining a final TPR less than half of the TPR achieved by aLTT with  $\epsilon$ -greedy acquisition strategy and  $\epsilon = 0.25$ . Moreover, aLTT with  $\epsilon = 0.25$  identifies 50% of the reliable instructions within the first 1000 testing rounds, while the non-adaptive testing procedure identifies less than 10% of the reliable instructions with the same number of rounds.

The improved efficiency of aLTT translates into superior performance in downstream tasks. To illustrate this point, we follow the described hyperparameter selection process with a post-selection phase that identifies a single hyperparameter  $\hat{\lambda}$  from the estimated set  $\hat{\Lambda}^{\text{rel}}$  of reliable instructions (see Figure 1). Specifically, we adopt the shortest instruction  $\lambda \in \hat{\Lambda}^{\text{rel}}$ . This post-selection criterion is motivated by the computational benefits of processing shorter prompts.

In Figure 5, we report the length of the selected instruction  $\hat{\lambda}$  as a function of the target execution accuracy  $1 - \alpha$ . Re-

sults are averaged over tasks with a final non-empty set  $\hat{\Lambda}^{\text{rel}}$ , while plots for individual tasks are included in Appendix B.3. Stricter accuracy requirements are seen to reduce the number of discovered reliable instructions, leading to an increase in the minimal instruction length within  $\hat{\Lambda}^{\text{rel}}$ . However, across all reliability levels  $\alpha$ , aLTT consistently delivers the shortest instructions, outperforming alternative schemes. This advantage is to be attributed to the data-adaptive acquisition strategy implemented by aLTT, which enables the discovery of a larger number of reliable instructions.

## 6. Conclusion

We introduced aLTT, a novel framework for hyperparameter selection that implements data-dependent sequential testing via early termination. Unlike the existing LTT, which builds on p-value multiple hypothesis testing (MHT), aLTT is based on sequential MHT via e-processes (Xu et al., 2021). In practical scenarios, including the problem of prompt engineering, this results in more efficient and flexible calibration procedures that maintain statistical validity and the same discovery power as LTT by using only a fraction of testing rounds.

Potential extensions of the aLTT framework include the study of scenarios characterized by distribution shift, data reuse (Wang et al., 2025) and simulation-aided calibration.

## Acknowledgments

This work was supported by the European Union’s Horizon Europe project CENTRIC (101096379). The work of Osvaldo Simeone was also supported by the Open Fellowships of the EPSRC (EP/W024101/1) and by the EPSRC project (EP/X011852/1). Thanks also to the Advanced Research and Invention Agency (ARIA) for supporting broader foundational work in this space.

## Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

## References

- Angelopoulos, A. N., Bates, S., Candès, E. J., Jordan, M. I., and Lei, L. Learn then test: Calibrating predictive algorithms to achieve risk control. *arXiv preprint arXiv:2110.01052*, 2021.
- Angelopoulos, A. N., Bates, S., Fisch, A., Lei, L., and Schuster, T. Conformal risk control. *arXiv preprint arXiv:2208.02814*, 2022.



- Angelopoulos, A. N., Bates, S., et al. Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4):494–591, 2023.
- Baruah, S. K., Cohen, N. K., Plaxton, C. G., and Varvel, D. A. Proportionate progress: A notion of fairness in resource allocation. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pp. 345–354, 1993.
- Bates, S., Angelopoulos, A., Lei, L., Malik, J., and Jordan, M. Distribution-free, risk-controlling prediction sets. *Journal of the ACM (JACM)*, 68(6):1–34, 2021.
- Bauer, P. Multiple testing in clinical trials. *Statistics in medicine*, 10(6):871–890, 1991.
- Benjamini, Y. and Hochberg, Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300, 1995.
- Benjamini, Y. and Yekutieli, D. The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*, pp. 1165–1188, 2001.
- Berman, R. and Van den Bulte, C. False discovery in A/B testing. *Management Science*, 68(9):6762–6782, 2022.
- Bonferroni, C. Teoria statistica delle classi e calcolo delle probabilita. *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze*, 8:3–62, 1936.
- Dabney, W., Rowland, M., Bellemare, M., and Munos, R. Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- de Sant Ana, P. M. and Marchenko, N. Radio access scheduling using CMA-ES for optimized QoS in wireless networks. In *2020 IEEE Globecom Workshops (GC Wkshps)*, pp. 1–6. IEEE, 2020.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Elbamby, M. S., Perfecto, C., Liu, C.-F., Park, J., Samarakoon, S., Chen, X., and Bennis, M. Wireless edge computing with latency and reliability guarantees. *Proceedings of the IEEE*, 107(8):1717–1737, 2019.
- Farzaneh, A., Park, S., and Simeone, O. Quantile learn-then-test: Quantile-based risk control for hyperparameter optimization. *IEEE Signal Processing Letters*, 2024.
- Ferdosian, N., Othman, M., Ali, B. M., and Lun, K. Y. Greedy-knapsack algorithm for optimal downlink resource allocation in LTE networks. *Wireless Networks*, 22:1427–1440, 2016.
- Fujimoto, S. and Gu, S. S. A minimalist approach to offline reinforcement learning. *Advances in Neural Information Processing Systems*, 34:20132–20145, 2021.
- Garcia, J. and Fernández, F. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- Genovese, C. R., Lazar, N. A., and Nichols, T. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage*, 15(4):870–878, 2002.
- Harvey, C. R. and Liu, Y. False (and missed) discoveries in financial economics. *The Journal of Finance*, 75(5): 2503–2553, 2020.
- Head, M. L., Holman, L., Lanfear, R., Kahn, A. T., and Jennions, M. D. The extent and consequences of p-hacking in science. *PLoS biology*, 13(3):e1002106, 2015.
- Honovich, O., Shaham, U., Bowman, S. R., and Levy, O. Instruction induction: From few examples to natural language task descriptions. *arXiv preprint arXiv:2205.10782*, 2022.
- Keselman, H. and Rogan, J. C. The Tukey multiple comparison test: 1953–1976. *Psychological Bulletin*, 84(5): 1050, 1977.
- Laros III, J. H., Pedretti, K., Kelly, S. M., Shu, W., Ferreira, K., Van Dyke, J., Vaughan, C., Laros III, J. H., Pedretti, K., Kelly, S. M., et al. Energy delay product. *Energy-Efficient High Performance Computing: Measurement and Tuning*, pp. 51–55, 2013.
- Levine, S., Kumar, A., Tucker, G., and Fu, J. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Li, C., Song, S., Zhang, J., and Letaief, K. B. Maximizing energy efficiency in wireless networks with a minimum average throughput requirement. In *2012 IEEE wireless communications and networking conference (WCNC)*, pp. 1130–1134. IEEE, 2012.
- Lindauer, M., Eggensperger, K., Feurer, M., Biedenkapp, A., Deng, D., Benjamins, C., Ruhkopf, T., Sass, R., and Hutter, F. SMAC3: A versatile Bayesian optimization package for hyperparameter optimization. *Journal of Machine Learning Research*, 23(54):1–9, 2022.
- Lindemann, L., Cleaveland, M., Shim, G., and Pappas, G. J. Safe planning in dynamic environments using conformal prediction. *IEEE Robotics and Automation Letters*, 2023.

- Liu, V., Nagarajan, P., Patterson, A., and White, M. When is offline policy selection sample efficient for reinforcement learning? *arXiv preprint arXiv:2312.02355*, 2023.
- Lu, C., Lemay, A., Chang, K., Höbel, K., and Kalpathy-Cramer, J. Fair conformal predictors for applications in medical imaging. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 12008–12016, 2022.
- Maclaurin, D., Duvenaud, D., and Adams, R. Gradient-based hyperparameter optimization through reversible learning. In *International Conference on Machine Learning*, pp. 2113–2122. PMLR, 2015.
- Mahapatra, R., Nijssure, Y., Kaddoum, G., Hassan, N. U., and Yuen, C. Energy efficiency tradeoff mechanism towards wireless green communication: A survey. *IEEE Communications Surveys & Tutorials*, 18(1):686–705, 2015.
- Mehrtash, A., Wells, W. M., Tempany, C. M., Abolmaesumi, P., and Kapur, T. Confidence calibration and predictive uncertainty estimation for deep medical image segmentation. *IEEE Transactions on Medical Imaging*, 39(12):3868–3878, 2020.
- meta. Llama. <https://LLaMA.meta.com/LLaMA3>, 2025. [Online; accessed 28-Jan-2025].
- Nasir, M. U., Earle, S., Togelius, J., James, S., and Cleghorn, C. Llmatic: neural architecture search via large language models and quality diversity optimization. In *proceedings of the Genetic and Evolutionary Computation Conference*, pp. 1110–1118, 2024.
- Nokia. Wireless suite, 2020. URL <https://github.com/nokia/wireless-suite>.
- Paine, T. L., Paduraru, C., Michi, A., Gulcehre, C., Zolna, K., Novikov, A., Wang, Z., and de Freitas, N. Hyperparameter selection for offline reinforcement learning. *arXiv preprint arXiv:2007.09055*, 2020.
- Pedregosa, F. Hyperparameter optimization with approximate gradient. In *International Conference on Machine Learning*, pp. 737–746. PMLR, 2016.
- Quach, V., Fisch, A., Schuster, T., Yala, A., Sohn, J. H., Jaakkola, T. S., and Barzilay, R. Conformal language modeling. *arXiv preprint arXiv:2306.10193*, 2023.
- Rakotoarison, H., Adriaensen, S., Mallik, N., Garibov, S., Bergman, E., and Hutter, F. In-context freeze-thaw Bayesian optimization for hyperparameter optimization. *arXiv preprint arXiv:2404.16795*, 2024.
- Ramdas, A. and Wang, R. Hypothesis testing with e-values. *arXiv preprint arXiv:2410.23614*, 2024.
- Ramdas, A., Ruf, J., Larsson, M., and Koolen, W. Admissible anytime-valid sequential inference must rely on non-negative martingales. *arXiv preprint arXiv:2009.03167*, 2020.
- Ramdas, A., Grünwald, P., Vovk, V., and Shafer, G. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601, 2023.
- Reynolds, L. and McDonnell, K. Prompt programming for large language models: Beyond the few-shot paradigm. In *Extended abstracts of the 2021 CHI conference on human factors in computing systems*, pp. 1–7, 2021.
- Schneider, L., Wistuba, M., Klein, A., Golebiowski, J., Zappella, G., and Merra, F. A. Hyperband-based bayesian optimization for black-box prompt selection. *arXiv preprint arXiv:2412.07820*, 2024.
- Schuster, T., Fisch, A., Gupta, J., Dehghani, M., Bahri, D., Tran, V., Tay, Y., and Metzler, D. Confident adaptive language modeling. *Advances in Neural Information Processing Systems*, 35:17456–17472, 2022.
- Seshia, S. A., Sadigh, D., and Sastry, S. S. Toward verified artificial intelligence. *Communications of the ACM*, 65(7):46–55, 2022.
- Shafer, G. Testing by betting: A strategy for statistical and scientific communication. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 184(2):407–431, 2021.
- Shafer, G. and Vovk, V. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3), 2008.
- Shafer, G. and Vovk, V. *Game-theoretic foundations for probability and finance*, volume 455. John Wiley & Sons, 2019.
- Shekhar, S. and Ramdas, A. Nonparametric two-sample testing by betting. *IEEE Transactions on Information Theory*, 2023.
- Shin, J., Ramdas, A., and Rinaldo, A. E-detectors: a non-parametric framework for sequential change detection. *arXiv preprint arXiv:2203.03532*, 2022.
- Shin, T., Razeghi, Y., Logan IV, R. L., Wallace, E., and Singh, S. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 2020.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. Deterministic policy gradient algorithms. In *International Conference on Machine Learning*, pp. 387–395. Pmlr, 2014.

- Stanczak, S., Wiczanowski, M., and Boche, H. *Fundamentals of resource allocation in wireless networks: theory and algorithms*, volume 3. Springer Science & Business Media, 2009.
- Swersky, K., Snoek, J., and Adams, R. P. Freeze-thaw Bayesian optimization. *arXiv preprint arXiv:1406.3896*, 2014.
- Tegmark, M. and Omohundro, S. Provably safe systems: the only path to controllable AGI. *arXiv preprint arXiv:2309.01933*, 2023.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.
- van den Oord, E. J. and Sullivan, P. F. False discoveries and models for gene discovery. *Trends in Genetics*, 19(10): 537–542, 2003.
- Vovk, V. and Bendtsen, C. Conformal predictive decision making. In *Conformal and Probabilistic Prediction and Applications*, pp. 52–62. PMLR, 2018.
- Vovk, V. and Wang, R. E-values: Calibration, combination and applications. *The Annals of Statistics*, 49(3):1736–1754, 2021.
- Wang, H., Dandapanthula, S., and Ramdas, A. Anytime-valid fdr control with the stopped e-bh procedure. *arXiv preprint arXiv:2502.08539*, 2025.
- Wang, R. and Ramdas, A. False discovery rate control with e-values. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(3):822–852, 2022.
- Waudby-Smith, I. and Ramdas, A. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(1):1–27, 2024.
- Xiao, L. Rethinking conventional wisdom in machine learning: From generalization to scaling. *arXiv preprint arXiv:2409.15156*, 2024.
- Xu, Z., Wang, R., and Ramdas, A. A unified framework for bandit multiple testing. *Advances in Neural Information Processing Systems*, 34:16833–16845, 2021.
- Zecchin, M., Park, S., and Simeone, O. Forking uncertainties: Reliable prediction and model predictive control with sequence models via conformal risk control. *IEEE Journal on Selected Areas in Information Theory*, 2024.
- Zhang, Z., Zhang, A., Li, M., and Smola, A. Automatic chain of thought prompting in large language models. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=5NTt8GFjUHkr>.
- Zhou, Y., Muresanu, A. I., Han, Z., Paster, K., Pitis, S., Chan, H., and Ba, J. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=92gvk82DE->.

## A. FWER and FDR-Controlling Procedures

In this section we review popular FWER and FDR-controlling strategies applied to the problem of hyperparameter selection.

### A.1. FWER Control

#### A.1.1. BONFERRONI CORRECTION

For a set of p-values  $\mathcal{P} = \{P_i\}_{i=1}^N$ , a simple  $(\alpha, \delta)$ -FWER-controlling procedure is given by the Bonferroni correction  $\mathcal{A}_{\text{Bon}}^{\text{FWER}}(\mathcal{P})$ , which returns the hyperparameter set

$$\hat{\Lambda}_{\text{Bon}}^{\text{rel}} = \left\{ \lambda_i : P_i \leq \frac{\delta}{N} \right\}. \quad (16)$$

The hyperparameter set  $\hat{\Lambda}_{\text{Bon}}^{\text{rel}}$  is guaranteed to be  $\delta$ -FWER-controlling (Bonferroni, 1936).

#### A.1.2. FIXED-SEQUENCE TESTING

If the designer has access to prior knowledge  $\mathcal{D}^0$  about which hyperparameters are likely to be more reliable, it is possible to increase the power of the statistical test, and thus the cardinality of the hyperparameter  $\hat{\Lambda}^{\text{rel}}$ , by using a *fixed-sequence testing* procedure  $\mathcal{A}_{\text{FS}}^{\text{FWER}}(\mathcal{P}, \mathcal{D}^0)$ . In fixed sequence hypothesis testing, the hypotheses  $\{\mathcal{H}_i\}_{i=1}^N$  are ordered based on the prior knowledge  $\mathcal{D}^0$  from the most likely to be reliable to the least likely to be reliable. Denote the  $k$ -th hypothesis in the corresponding ordered sequence as  $\mathcal{H}_{(k)}$ , and the associated p-value as  $P_{(k)}$ . With fixed-sequence testing, the hypotheses are sequentially tested at a reliability level  $\alpha$  until one is accepted. Accordingly, the resulting hyperparameter set contains all hypotheses up until the first acceptance, i.e.,

$$\hat{\Lambda}_{\text{FS}}^{\text{rel}} = \left\{ \lambda_{(j)} : P_{(i)} \leq \delta, \forall i \leq j \right\}. \quad (17)$$

For any ordering of the p-values, the hyperparameter set  $\hat{\Lambda}_{\text{FS}}^{\text{rel}}$  is guaranteed to be  $\delta$ -FWER-controlling (Bauer, 1991).

### A.2. FDR Control

#### A.2.1. BENJAMINI-HOCHBERG PROCEDURE

Given a set of independent p-values  $\mathcal{P}$ , denote the  $k$ -th smallest value in the set as  $P_{(k)}$  and the associated hyperparameter as  $\lambda_{(k)}$ . For an error level  $\delta$ , the Benjamini-Hochberg procedure  $\mathcal{A}_{\text{BH}}^{\text{FDR}}(\mathcal{P})$  returns the prediction set

$$\hat{\Lambda}_{\text{BH}}^{\text{rel}} = \left\{ \lambda_{(i)} : P_{(i)} \leq \frac{i\delta}{N} \right\}. \quad (18)$$

By (18), BH applies a larger threshold to hyperparameters  $\lambda \in \Lambda$  that have larger p-values and are thus less likely to be reliable. If the p-values in set  $\mathcal{P}$  are independent, the Benjamini-Hochberg (BH) procedure  $\mathcal{A}_{\text{BH}}^{\text{FDR}}(\mathcal{P})$  guarantees FDR control at a level  $\delta$  (Benjamini & Hochberg, 1995).

#### A.2.2. BENJAMINI-YEKUTIELI PROCEDURE

In case of arbitrarily dependent p-values in set  $\mathcal{P}$  the error level  $\delta$  has to be adjusted by a multiplicative factor  $(\sum_{n=1}^N 1/n)^{-1}$ . The resulting FDR-controlling procedure, also known as the Benjamini-Yekutieli (BY) procedure, yields the set

$$\hat{\Lambda}_{\text{BY}}^{\text{rel}} = \left\{ \lambda_{(i)} : P_{(i)} \leq \frac{i\delta}{\sum_{n=1}^N N/n} \right\}. \quad (19)$$

The BY procedure  $\mathcal{A}_{\text{BY}}^{\text{FDR}}(\mathcal{P})$  is  $\delta$ -FDR-controlling (Benjamini & Yekutieli, 2001).



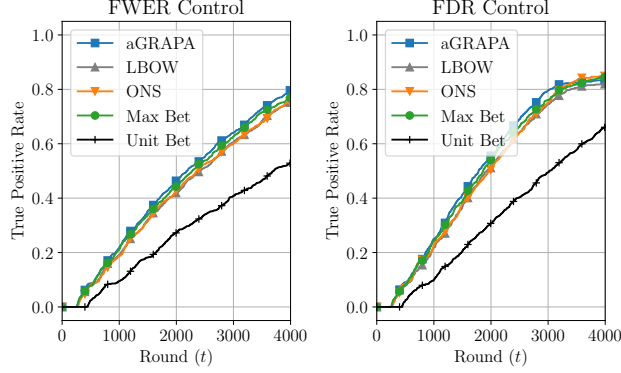


Figure 6. True positive rate of aLTT with  $\epsilon$ -greedy acquisition policy with  $\epsilon = 0.25$  under different betting strategies. On the left panel, the prediction sets satisfy FWER control while on the right FDR control. In both cases, the tolerance level is  $\delta = 0.1$ .

### A.2.3. E-BENJAMINI-HOCHBERG PROCEDURE

Alternatively, given the set of e-values  $\mathcal{E}^t = \{E_i^t\}_{i=1}^N$ , denote by the  $k$ -th largest e-value in  $\mathcal{E}$  as  $E_{(k)}^t$  and the associated hyperparameter as  $\lambda_{(k)}$ . For an error level  $\delta$ , the e-Benjamini-Hochberg procedure  $\mathcal{A}_{\text{eBH}}^{\text{FDR}}(\mathcal{E})$  outputs the hyperparameter subset

$$\hat{\Lambda}_{\text{eBH}}^{\text{rel}} = \left\{ \lambda_{(i)} : E_{(i)}^t \geq \frac{N}{i\delta} \right\}. \quad (20)$$

Following the same basic principles underlying BH in (18), eBH applies a smaller threshold to hyperparameters  $\lambda \in \Lambda$  that have smaller e-values and are thus less likely to be reliable. The eBH procedure returns an  $\delta$ -FDR-controlling set even for arbitrarily dependent e-values  $E_{(k)}^t$  without the need to adjust the error level  $\delta$  as in the case of the p-value-based BH procedure (Xu et al., 2021).

## B. Additional Experiments

### B.1. Online Policy Selection for Offline Reinforcement Learning

#### B.1.1. EFFECT OF THE BETTING STRATEGY

In Section 5.1, we studied the TPR of aLTT under the *approximate growth rate adaptive to the particular alternative* (aGRAPA) betting strategy. In the following, we compare its performance against non-adaptive betting strategies, specifically the *maximum bet* (MaxBet) strategy, which sets the bet  $\mu_i^t$  to the maximum allowed value of  $1/\alpha$ , and the *unitary bet* (UnitBet) strategy, where  $\mu_i^t = 1$ . Additionally, we evaluate two alternative adaptive online betting strategies based on approximate wealth maximization (Waudby-Smith & Ramdas, 2024): *lower-bound on the wealth* (LBOW) and *online Newton step* (ONS).

Under the same online policy selection for offline reinforcement learning as set up in Section 5.1, in Figure 6 we compare the TPR of aLTT with an  $\epsilon$ -greedy acquisition policy. The left panel targets FWER control, while the right focuses on FDR control. All adaptive betting strategies exhibit similar performance, with aGRAPA showing a slight advantage. Interestingly, in this scenario, the non-adaptive MaxBet strategy performs surprisingly well, while UnitBet achieves the lowest TPR.

#### B.1.2. QUANTILE RISK CONTROL

In this experiment, we consider the same set-up of Section 5.1 with the caveat that, instead of identifying policies with a large enough average reward, we focus on policies with a small enough quantile risk — a risk definition that explicitly captures safety (Dabney et al., 2018; Farzaneh et al., 2024). The quantile risk of a hyperparameter  $\lambda$  at a level  $q$  is defined as

$$R_q(\lambda) = \inf_r \{r : \Pr[R(\lambda, Z) \leq r] > 1 - q\}. \quad (21)$$




Figure 7. True positive rate of LTT and aLTT with  $\epsilon$ -greedy acquisition policy for  $\epsilon \in \{0.25, 0.5, 0.75, 0.95\}$  and non-adaptive acquisition. On the left panel, the prediction sets satisfy FWER control while on the right FDR control. In both cases, a policy is considered reliable if the quantile risk at level  $q = 0.1$  is lower than  $\alpha = 0.57$ , and the tolerance level is  $\delta = 0.1$ .




Figure 8. True positive rate of aLTT as a function of the number of testing rounds, using a top- $K$   $\epsilon$ -greedy acquisition policy with  $\epsilon = 0.25$  and  $K \in \{1, 2, 3, 4\}$ .

Controlling the quantile risk ensures guarantees on the average fraction of times the policy risk exceeds the level  $R_q(\lambda)$ . For the RL setting, we deal with rewards, so the inequalities in the definition are reversed.

We aim to identify policies with a quantile risk at the level  $q = 0.1$  that is smaller than  $\alpha = 0.57$ . In Figure 7, we report the true positive rate of aLTT using the  $\epsilon$ -greedy acquisition policy and LTT. aLTT significantly outperforms LTT, identifying up to 75% of the policies that satisfy the given reliability requirement within the maximum testing round,  $T$ . In contrast, LTT is unable to identify any policy due to its non-adaptive testing strategy.

### B.1.3. TOP- $K$ $\epsilon$ -GREEDY ACQUISITION STRATEGY

In Figure 2, we instantiate aLTT using an  $\epsilon$ -greedy acquisition strategy. With probability  $1 - \epsilon$ , the strategy selects the hypothesis associated with the largest e-process among those that have not yet been rejected. With probability  $\epsilon$ , it selects a hypothesis at random from the remaining unrejected set. A natural generalization of this strategy is the top- $K$   $\epsilon$ -greedy acquisition strategy. In this variant, with probability  $1 - \epsilon$ , aLTT tests the set of the  $K$  hypotheses associated with the largest e-processes that have not yet been rejected; with probability  $\epsilon$ , it tests  $K$  hypotheses selected uniformly at random from the entire set of unrejected hypotheses.

In Figure 8, we report the FDR of aLTT using this top- $K$   $\epsilon$ -greedy strategy with  $\epsilon = 0.25$  and varying values of  $K \in \{1, 2, 3, 4\}$ . As  $K$  increases, the TPR performance of aLTT improves. This is expected, as the effective number of evaluations per testing round increases proportionally with  $K$  allowing aLTT to discover reliable models in less testing rounds.




Figure 9. Size of the prediction set returned by LTT and aLTT with a top- $K$   $\epsilon$ -greedy acquisition policy for varying reliability levels  $\alpha$ . As  $\alpha$  increases, the reliability requirement becomes stricter, requiring a larger expected reward  $R(\lambda) = \mathbb{E}_{P_Z}[R(\lambda, Z)] > \alpha$ . For all values of  $\alpha$ , aLTT is able to discover a larger number of reliable policies.

#### B.1.4. RELIABILITY LEVEL

In the following, we compare the efficiency of aLTT and LTT under different values of the reliability level  $\alpha$ . In Figure 9, we vary  $\alpha \in [0.565, 0.58]$  and report the expected size of the prediction set returned by LTT and aLTT after  $T = 5000$  testing rounds.

In the context of the online policy selection problem, a higher reliability level  $\alpha$  imposes a stricter requirement on the expected reward of the policies, namely  $R(\lambda) = \mathbb{E}_{P_Z}[R(\lambda, Z)] > \alpha$ . As a result, the size of the prediction sets decreases as  $\alpha$  increases. However, for all values of  $\alpha$ , aLTT is capable of identifying a larger number of reliable policies compared to LTT. For example, under FWER control, when  $\alpha = 0.575$ , LTT fails to find any reliable policies, whereas aLTT discovers up to four reliable policies.

### B.2. Reliable Hyperparameter Selection for Wireless Resource Allocation

In wireless communication systems, resource allocation is an essential functionality that regulates access to the spectrum for users and services (Stanczak et al., 2009). The performance of resource allocation policies is evaluated by using key performance indicators (KPIs) such as throughput, delay, and energy efficiency. Despite the randomness inherent in the network conditions, some services require strict reliability guarantees in terms of KPIs. For instance, gaming applications must meet latency constraints (Elbamy et al., 2019), streaming connections are subject to throughput requirements (Li et al., 2012), and battery-powered transmitters have strict energy-efficiency constraints (Mahapatra et al., 2015).

#### B.2.1. PROBLEM DEFINITION

As illustrated in Figure 10, we consider a downlink resource allocation problem in which, at every transmission frame  $k \geq 1$ , a base station serves users by communicating bits from the corresponding queues. This is done by assigning physical resource blocks (PRBs) to users based on a descriptor  $s_k$  of the network conditions that includes the state of users' transmission buffers, users' priorities, and channel conditions. Following the simulation software Nokia Wireless Suite (Nokia, 2020), we consider three different types of PRBs assignment policies: a proportional fair scheme (Baruah et al., 1993), a knapsack allocation policy (Ferdosian et al., 2016), and a learning-based scheme based on a Covariance Matrix Adaptation Evolution Strategy (CMA-ES) (de Sant Ana & Marchenko, 2020). Overall, a resource allocation scheme is specified by the hyperparameters  $\lambda = \{\lambda_1, \lambda_2\}$ , with  $\lambda_1 \in \{\text{proportional fair, knapsack, CMA-ES}\}$  identifying the scheduling policy, and  $\lambda_2$  determining the transmit power level. We let the hyperparameter  $\lambda_2$  vary from 0 dB to 24 dB in 1 dB increments. This way, the set  $\Lambda$  contains 75 candidate hyperparameters.

An instance of the resource allocation problem is described by the random quantity  $Z$ , which encompasses channel conditions and packet generation. Furthermore, the performance of the resource allocation hyperparameter  $\lambda$  is measured by the following KPIs: (i) *average transmission delay* of a packet  $T_{\text{tx}}(\lambda, Z)$ ; (ii) *overall energy efficiency*, which is defined as the ratio between the number of transmitted bits and the overall transmitted energy within the episode,  $\rho(\lambda, Z) = \sum_{k=1}^K B^k / \sum_{k=1}^K E^k$  where  $B^k$  and  $E^k$  are the number of transmitted bits and transmit energy at slot  $k$ ; (iii) *average queue length occupancy*  $Q(\lambda, Z)$ , which is normalized by the maximum queue size  $Q_{\text{max}}$ ; (iv) *average energy-delay product*,

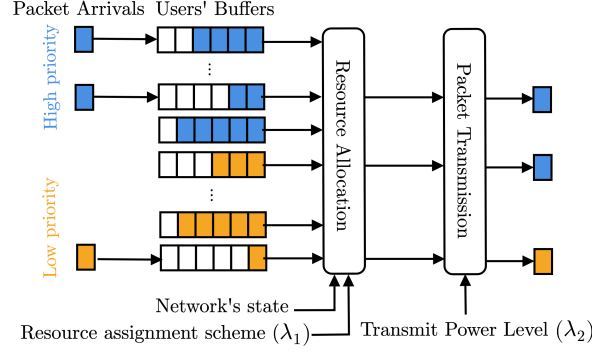


Figure 10. System diagram of the considered resource allocation problem. At each transmission frame  $k$ , new packets are randomly generated and stored in the corresponding users' buffers. A resource allocation scheme, determined by hyperparameter  $\lambda_1$ , assigns resources to the users, deciding whose users to schedule at round  $k$ . Packets from the scheduled users' queues are then transmitted over a wireless channel using a transmit power dictated by the hyperparameter  $\lambda_2$ .




Figure 11. Average energy-delay product of the resource allocation policy returned by LTT and aLTT with  $\epsilon$ -greedy acquisition policy with  $\epsilon \in \{0.25, 0.5, 0.95\}$  under different queue occupancy efficiency requirements. The final hyperparameter is obtained by selecting the hyperparameter in the estimated set of reliable policies  $\hat{\Lambda}^{\text{rel}}$  that is associated with the lowest empirical energy-delay product. On the left panel, the prediction sets satisfy FWER control, while on the right FDR control, with the tolerance level  $\delta = 0.1$ .

$\text{EDP}(\lambda, Z) = T_{\text{tx}}(\lambda, Z) \sum_{k=1}^K E^k / K$ , quantifying the overall performance in terms of energy and delay (Laros III et al., 2013).

### B.2.2. RESULTS

We study a setting with differentiated service requirements. In particular, we enforce a constraint on the average queue occupancy and energy efficiency for high-priority users, while making a best-effort attempt at minimizing the energy-delay product for low-priority users. Writing as  $Q^{\text{HI}}(\lambda) = \mathbb{E}_{P_Z}[Q^{\text{HI}}(\lambda, Z)]$  the average queue occupancy of high-priority users, and as  $\rho^{\text{HI}} = \mathbb{E}_{P_Z}[\rho^{\text{HI}}(\lambda, Z)]$  the average energy efficiency of high-priority users, and as  $(\alpha_1, \alpha_2)$  the corresponding requirements. A hyperparameter is reliable if it meets both the queue occupancy and energy efficiency requirements for high-priority users, i.e.,  $\Lambda^{\text{rel}} = \{\lambda \in \Lambda : Q^{\text{HI}}(\lambda) \leq \alpha_1 \cap \rho^{\text{HI}}(\lambda) \geq \alpha_2\}$ . An e-process for the hypothesis  $\mathcal{H}_i : Q^{\text{HI}}(\lambda_i) > \alpha_1 \cap \rho^{\text{HI}}(\lambda_i) < \alpha_2$  can be obtained as the minimum between an e-process for hypothesis  $\mathcal{H}_i^1 : Q^{\text{HI}}(\lambda_i) > \alpha_1$ , and an e-process for the hypothesis  $\mathcal{H}_i^2 : \rho^{\text{HI}}(\lambda_i) < \alpha_2$ . Note that any e-merging function can be applied to ensure validity over the joint hypothesis (see, e.g., (Vovk & Wang, 2021; Ramdas & Wang, 2024)), we choose minimum for simplicity.

In Figure 11, we compare the performance of LTT and aLTT for a calibration horizon  $T = 4000$  by choosing in the estimated set of reliable hyperparameters  $\hat{\Lambda}^{\text{rel}}$  the hyperparameter  $\hat{\lambda}$  that minimizes the empirical energy-delay product of low-priority users based on all the data  $\mathcal{D}^T$  collected across the testing steps. We vary the queue occupancy reliability level  $\alpha_1$  while fixing the energy-efficiency requirement  $\alpha_2 = 0.01$  Mbit/Joule. As the queue occupancy target  $\alpha_1$  increases, the requirement becomes less stringent, and the energy-delay product of the low-priority users under the returned policy decreases. In both cases—FWER control on the left and FDR control on the right—the performance of aLTT is significantly



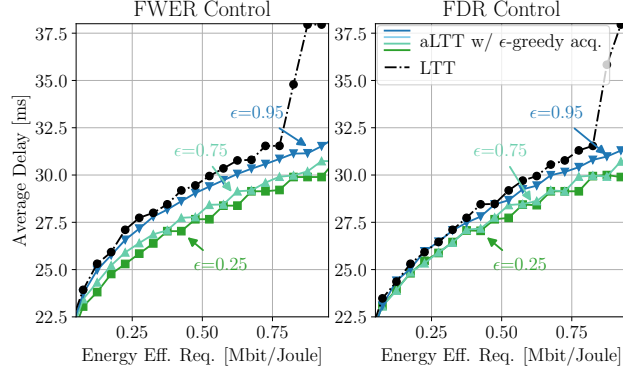


Figure 12. Average delay of the resource allocation policy returned by LTT and aLTT with  $\epsilon$ -greedy acquisition policy with  $\epsilon \in \{0.25, 0.5, 0.95\}$  under different energy efficiency requirements. The final policy is obtained by selecting the policy in the estimated set of reliable policies  $\hat{\Lambda}^{\text{rel}}$  that is associated to the lowest empirical delay. On the left panel, the prediction sets satisfy FDR control, while on the right FWER control, with the tolerance level  $\delta = 0.1$ .

superior to that of LTT. What is more, in this example, the looser guarantees provided by FDR do not entail any reliability loss, since the FWER of the hyperparameter  $\hat{\lambda}$  returned by FDR-controlling procedures was found to be significantly below  $\delta = 0.1$ .

### B.2.3. MINIMIZING AVERAGE DELAY UNDER AN ENERGY EFFICIENCY REQUIREMENT

Now we consider the wireless resource allocation problem under a reliability requirement on the average energy efficiency. Specifically, we consider the population risk  $R(\lambda) = \mathbb{E}_{P_Z}[\rho(\lambda, Z)]$ , so that the set of reliable hyperparameters is given by  $\Lambda^{\text{rel}} = \{\lambda \in \Lambda : R(\lambda) \geq \alpha\}$  for the given target energy efficiency  $\alpha$ . After  $T = 1000$  testing steps, with each step consisting of a resource allocation episode of  $K = 2500$  frames, the hyperparameter selection scheme selects from the estimated subset of reliable hyperparameters  $\hat{\Lambda}^{\text{rel}}$ , the hyperparameter  $\hat{\lambda}$  that minimizes the empirical communication delay  $\hat{T}_{\text{tx}}(\lambda, \mathcal{D}^T) = 1/|\mathcal{D}^T| \sum_{Z \in \mathcal{D}^T} \hat{T}_{\text{tx}}(\lambda, Z)$  based on all the data  $\mathcal{D}^T$  collected across the testing steps.

In Figure 12 we report the average delay  $T_{\text{tx}}(\hat{\lambda}) = \mathbb{E}_{P_Z}[T_{\text{tx}}(\hat{\lambda}, Z)]$ , estimated on hold-out data, of the selected policy  $\hat{\lambda}$  as a function of the reliability level  $\alpha$  for LTT and aLTT with an  $\epsilon$ -greedy acquisition function, when setting the target error level  $\delta = 0.1$ . On the left panel, we consider FWER control, while the right presents the performance with FDR control.

As the reliability requirement  $\alpha$  increases, becoming more stringent, the average delay of the selected hyperparameter  $\hat{\lambda}$  increases. However, thanks to adaptive testing, aLTT returns a larger prediction set  $\hat{\Lambda}^{\text{rel}}$  providing hyperparameters  $\hat{\lambda}$  with lower delays as compared to LTT. Furthermore, due to the less stringent FDR requirements, FDR-controlling testing procedures generally return policies with average delays smaller than FWER-controlling.

### B.3. Reliable Prompt Selection for Natural Language Processing

In Figures 4 and 5, we report the performance averaged over ten tasks from the instruction induction dataset (Honovich et al., 2022). Complementing the results in the main text, here we provide a breakdown of the methods' performance across the ten tasks. In Figure 13, we show the TPR of the schemes considered in Section 5.2.2 under FDR and FWER control with tolerance  $\delta = 0.1$ . Across all tasks, aLTT consistently outperforms LTT, which returns empty sets for many tasks.

In Figure 14, we plot the length of the shortest instruction in the predicted set  $\hat{\Lambda}^{\text{rel}}$  returned by the different hyperparameter selection schemes as a function of the execution accuracy  $1 - \alpha$ . If a prediction scheme returns an empty predicted hyperparameter set, we set the instruction length to the length of the longest instruction in the candidate set  $\Lambda$ . For all tasks, aLTT can identify shorter instructions compared to LTT under the same testing budget.

#### B.3.1. EXAMPLE PROMPTS

A list of the tested prompts together with their accuracy levels can be found at: <https://github.com/kclip/aLTT>.

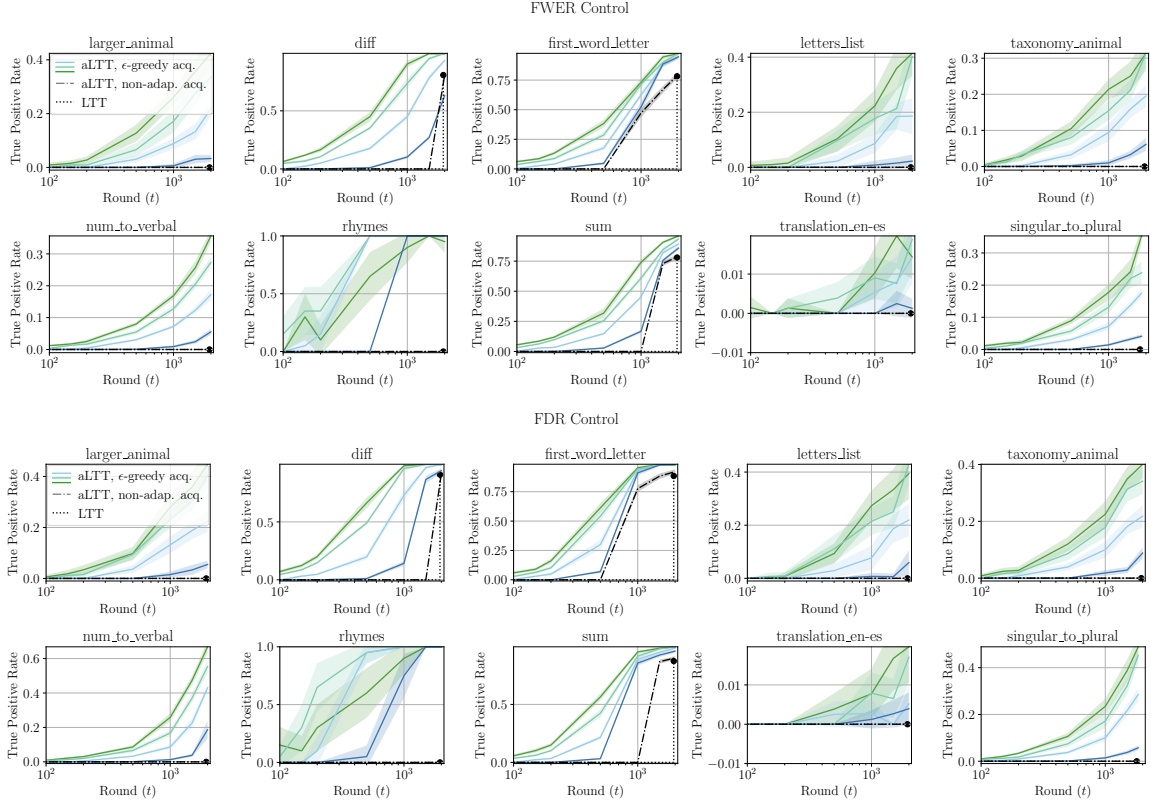


Figure 13. Breakdown of the average performance of methods presented in Figure 4 across the 10 tasks in the instruction induction dataset (Honovich et al., 2022) for which a non-empty hyperparameter sets  $\hat{\Lambda}^{\text{rel}}$  is returned. Across all tasks, aLTT delivers the highest true positive rate.

git

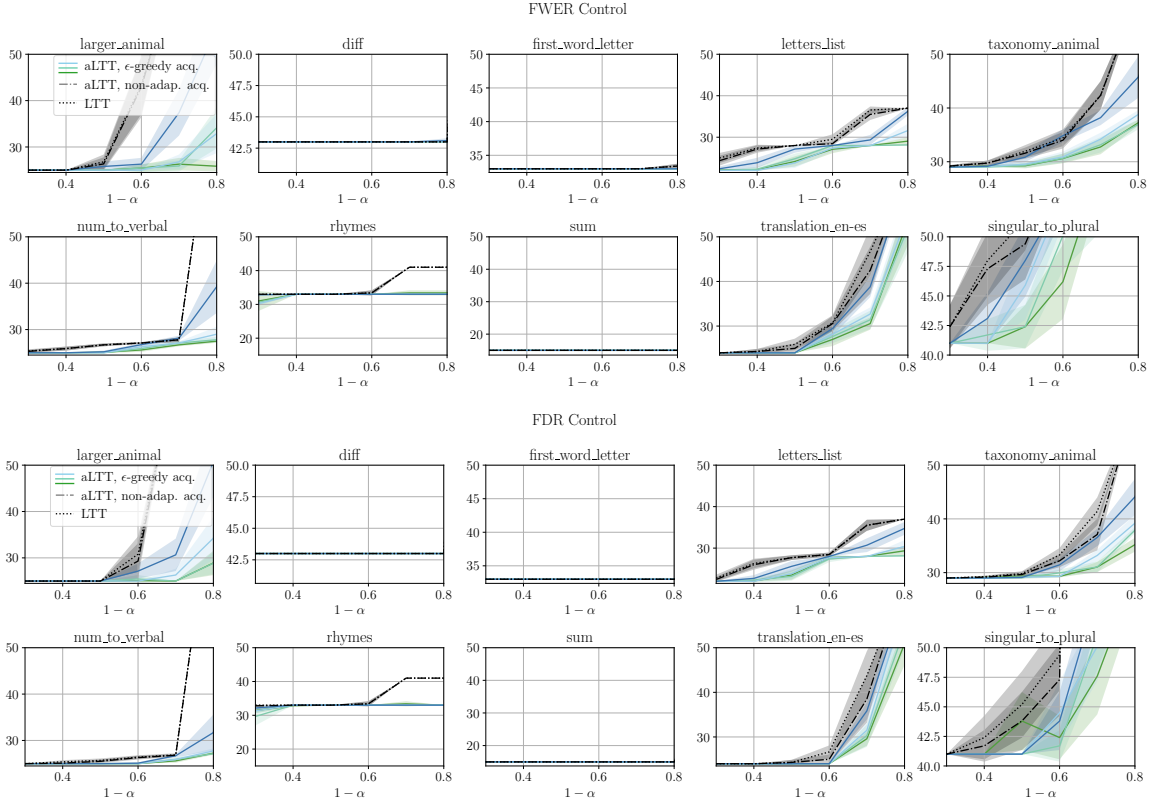


Figure 14. Length of the shortest instruction in  $\hat{\Lambda}^{\text{rel}}$  as a function of target execution accuracy  $1 - \alpha$  for each task in the instruction induction dataset (Honovich et al., 2022) for which the methods presented in Figure 5 return non-empty hyperparameter sets  $\hat{\Lambda}^{\text{rel}}$ .