IN-CONTEXT REINFORCEMENT LEARNING THROUGH BAYESIAN FUSION OF CONTEXT AND VALUE PRIOR

Anonymous authors
Paper under double-blind review

ABSTRACT

In-context reinforcement learning (ICRL) promises fast adaptation to unseen environments without parameter updates, but current methods either cannot improve beyond the training distribution or require near-optimal data, limiting practical adoption. We introduce SPICE, a Bayesian ICRL method that learns a prior over Q-values via deep ensemble and updates this prior at test-time using in-context information through Bayesian updates. To recover from poor priors resulting from training on sub-optimal data, our online inference follows an Upper-Confidence Bound rule that favours exploration and adaptation. In bandit settings, we prove this principled exploration reaches regret-optimal behaviour even when pretrained only on suboptimal trajectories. We validate these findings empirically across bandit and control benchmarks. SPICE achieves near-optimal decisions on unseen tasks, substantially reduces regret compared to prior ICRL and meta-RL approaches while rapidly adapting to unseen tasks and remaining robust under distribution shift.

1 INTRODUCTION

Following the success of transformers with in-context learning abilities Vaswani et al. (2017), In-Context Reinforcement Learning (ICRL) emerged as a promising paradigm Chen et al. (2021); Zheng et al. (2022). ICRL aims to adapt a policy to new tasks using only a context of logged interactions and no parameter updates. This approach is particularly attractive for practical deployment in domains where training classic online RL is either risky or expensive, where abundant historical logs are available, or where fast gradient-free adaptation is required. Examples include robotics, autonomous driving or buildings energy management systems. ICRL improves upon classic offline RL by amortising knowledge across tasks, as a single model is pre-trained on trajectories from many environments and then used at test time with only a small history of interactions from the test task. The model must make good decisions in new environments using this in-context dataset as the only source of information Moeini et al. (2025).

Existing ICRL approaches suffer from three main limitations. First, behaviour-policy bias from supervised training objectives: methods trained with Maximum Likelihood Estimation (MLE) on actions inherit from the same distribution as the behaviour policy. When the behaviour policy is suboptimal, the learned model performs poorly. Many ICRL methods fail to improve beyond the pretraining data distribution and essentially perform imitation learning Dong et al.; Lee et al. (2023). Second, existing methods lack uncertainty quantification and inference-time control. Successful online adaptation requires epistemic uncertainty over action values to enable temporally coherent exploration. Most ICRL methods expose logits but not actionable posteriors over Q-values, which are needed for principled exploration like Upper Confidence Bound (UCB) or Thompson Sampling (TS) Lakshminarayanan et al. (2017); Osband et al. (2016; 2018); Auer (2002); Russo et al. (2018). Third, current algorithms have unrealistic data requirements that make them unusable in most real-world deployments. Algorithm Distillation (AD) Laskin et al. (2022) requires learning traces from trained RL algorithms, while Decision Pretrained Transformers (DPT) Lee et al. (2023) needs optimal policy to label actions. Recent work has attempted to loosen these requirements, like Decision Importance Transformers (DIT) Dong et al. and In-Context Exploration with Ensembles (ICEE) Dai et al. (2024). However, these methods lack explicit measure of uncertainty and test-time controller for exploration and efficient adaptation.

To address these limitations, we introduce SPICE (**S**haping **P**olicies **I**n-**C**ontext with **E**nsemble prior), a Bayesian ICRL algorithm that maintains a prior over Q-values using a deep ensemble and updates this prior with state-weighted evidence from the context dataset. The resulting per-action posteriors can be used greedily in offline settings or with a posterior-UCB rule for online exploration, enabling test-time adaptation to unseen tasks without parameter updates. In bandit settings, we show that the SPICE inference controller achieves the same optimal logarithmic regret rate as UCB without assuming the model is trained on optimal data. We test our algorithm in bandit and dark room environments to compare against prior work, demonstrating that our algorithm achieves near-optimal decision making on unseen tasks while substantially reducing regret compared to prior ICRL and meta-RL approaches. This work paves the way for real-world deployment of ICRL methods, which should feature good uncertainty quantification and test-time adaptation to new tasks without relying on unrealistic optimal control trajectories for training.

2 Related Work

Meta-RL. Classical meta-reinforcement learning aims to learn to adapt across tasks with limited experience. Representative methods include RL^2 Duan et al. (2016), gradient-based meta-learning such as MAML Finn et al. (2017); and probabilistic context–variable methods such as PEARL Rakelly et al. (2019). These approaches typically require online interaction and task-aligned adaptation loops during deployment.

Sequence modelling for decision-making. Treating control as sequence modelling has proven effective with seminal works such as Decision Transformer (DT) Chen et al. (2021) and Trajectory Transformer models Janner et al. (2021). Scaling variants extend DT to many games and longer horizons Lee et al. (2022); Correia & Alexandre (2023), while Online Decision Transformer (ODT) blends offline pretraining with online fine-tuning via parameter updates Zheng et al. (2022). These works paved the way for in context decision making.

In-context RL via supervised pretraining. Two influential ICRL methods are Algorithm Distillation (AD) Laskin et al. (2022), which distills the learning dynamics of a base RL algorithm into a Transformer that improves in-context without gradients, and Decision-Pretrained Transformer (DPT) Lee et al. (2023), which is trained to map a query state and in-context experience to optimal actions and is theoretically connected to posterior sampling. Both rely on labels generated by strong/optimal policies (or full learning traces) and therefore inherit behaviour-policy biases from the data Moeini et al. (2025). DIT (Dong et al.) improves over behaviour cloning by reweighting a supervised policy with in-context advantage estimates, but it remains a purely supervised objective: it exposes no calibrated uncertainty, produces no per-action posterior, and lacks any inference-time controller or regret guarantees. ICEE (Dai et al., 2024) induces exploration–exploitation behaviour inside a Transformer at test time, yet it does so heuristically, without explicit Bayesian updates, calibrated posteriors, or theoretical analysis. By contrast, SPICE is the first ICRL method to (i) learn an explicit value prior with uncertainty from suboptimal data, (ii) perform Bayesian context fusion at test time to obtain per-action posteriors, and (iii) act with posterior-UCB, yielding principled exploration and a provable $O(\log K)$ regret bound with only a constant warm-start term.

3 Bayesian In-Context Decision Making

In this section, we introduce the key components of our method. In Sec. 3.1, we present the model architecture and key design choices for its training. In Sec. 3.2, we introduce our test-time controller, which is further analysed theoretically in Sec. 4 and empirically in Sec. 5 and Sec. 6. The complete method is summarised in Sec. 3.3.

3.1 Definition of the Sequence Model

Consider a setting in which we draw tasks $T \sim \mathcal{T}$ with state space $\mathcal{S}$, action space $\mathcal{A}$, horizon H, rewards $r_t \in \mathbb{R}$ and discount $\gamma \in [0, 1]$. At test time, given a task T, the agent receives a multi-episode in-context dataset $\mathcal{C} = \{(s_t, a_t, r_t, s_{t+1})\}$ collected from T and must select an action a for a query state s without any parameter updates. The dataset $\mathcal{C}$ may be provided offline or collected through online rollouts during test time. Our approach focuses on discrete action spaces $\mathcal{A}$, though it extends naturally to continuous actions. [1]

The trunk of our model is composed of a transformer architecture. Following prior work (Lee et al., 2023), a causal GPT-2 transformer is used to encode sequences of transitions. Each transition is embedded using a single linear layer $\mathbf{h}_t = \mathrm{Linear}\big([\, s_t,\ a_t,\ s'_t,\ r_t\,]\big) \in \mathbb{R}^D$, where D is the hidden size dimension. A sequence is processed as

$$\big(\underbrace{[\, s_{\text{qry}},\ 0,\ 0,\ 0]}_{\text{query token}},\ \underbrace{[\, s_1, a_1, s'_1, r_1\,], \ldots, [\, s_H, a_H, s'_H, r_H\,]}_{\text{context transitions}}\big)$$

The query token is placed first and filled with dummy slots for (a, s', r); the transformer then outputs a hidden vector at each position. Two decoder heads are used:

- **Policy head** $\pi_\theta(a \mid \cdot) = \mathrm{softmax}(\mathrm{Linear}(\mathbf{h}))$.
- **Value ensemble head** $Q_{\phi_k}(a \mid \cdot)$, $k = 1, \ldots, K$, further described in the following.

Q-value ensemble Deep ensembles can be used as Bayesian neural networks, as they capture the different modes of the Bayesian posterior (Fort et al., 2019; Wilson & Izmailov, 2020). The epistemic uncertainty can be estimated from the model standard deviation of the models, or model disagreement Lakshminarayanan et al. (2017). To account for model uncertainty, we model the Q-value using a deep ensemble of K separate MLP f_k. To ensure diversity in the ensemble, a randomised prior (Osband et al., 2018) p_k is added to each Q-value network in the form of a frozen and randomly initialised MLP $Q_{\phi_k} = f_k + \alpha p_k$,where $\alpha > 0$ scales the prior. In addition, to avoid model collapse, the loss is regularised using an anchor loss (Pearce et al., 2018), where $\phi^{(0)}$ represents the initial weights.

$$\mathcal{L}_{\text{anchor}} = \sum_{k=1}^{K} \sum_{j} \big\| \phi_{k,j} - \phi_{k,j}^{(0)} \big\|_2^2. \tag{1}$$

The query hidden vector is concatenated with a given action a encoded as a one-hot vector and passed to each Q-value head:

$$Q_{\phi_k}(a \mid s, \mathcal{C}) \;=\; f_k\big([\mathbf{h}_{\text{qry}};\ \mathrm{onehot}(a)]\big) \;+\; \alpha\, p_k\big([\mathbf{h}_{\text{qry}};\ \mathrm{onehot}(a)]\big), \tag{2}$$

The ensemble mean and standard deviation are used as calibrated value prior for ICRL

$$\bar{Q}(a) \;=\; \tfrac{1}{K} \sum_{k=1}^{K} Q_{\phi_k}(a), \qquad \sigma_Q(a) \;=\; \sqrt{\tfrac{1}{K-1} \sum_{k=1}^{K} \big(Q_{\phi_k}(a) - \bar{Q}(a)\big)^2}. \tag{3}$$

Learning a policy with Advantage and Uncertainty Reweightings The policy head π_θ is trained to optimise a weighted cross-entropy loss:

$$\mathcal{L}_\pi \;=\; \mathbb{E}_b \left[\frac{1}{H} \sum_{t=1}^{H} \omega_b \,\Big(-\log \pi_\theta(a_b^\star \,\big|\, \mathbf{h}_{b,t}) \Big) \right], \qquad \omega_b = \omega_{\text{IS}} \cdot \omega_{\text{adv}} \cdot \omega_{\text{epi}}. \tag{4}$$

given a training example x_b (i.e. a context with horizon H) and $a_b^\star$ its action label. The multiplicative weight ω_b is the product of three weight factors described below.

(i) Propensity correction. Offline datasets reflect the action selection of the behaviour policy $\pi_b(\cdot \mid s)$, which induces a mismatch between the supervised training target and the uniform reference action distribution. To remove this behaviour-policy bias and recover the target likelihood under a uniform distribution $\pi_u(\cdot \mid s)$, labeled samples can be re-weighted with an importance ratio (Dai et al., 2024):

$$\omega_{\text{IS}} = \mathrm{clip}\left(\frac{\pi_u(a_b^\star \mid s)}{\pi_b(a_b^\star \mid s)}\ ,\ 0,\ c_{\text{iw}} \right), \qquad \pi_u(a \mid s) = \tfrac{1}{|\mathcal{A}|}. \tag{5}$$

Intuitively, overrepresented actions under π_b are downweighted, and rare but informative actions are upweighted.

[1]For continuous action settings, one can replace the categorical policy head with a parametric density (e.g., Gaussian), concatenate raw action vectors instead of one-hot encodings in the value ensemble, and perform posterior updates using kernel-weighted statistics in action space.

(ii) Advantage weighting. Inspired by (Wang et al., 2018; Dai et al., 2024; Peng et al., 2019), we upweight transitions whose estimated advantage is positive so that the trunk allocates more capacity to reward-relevant behaviours, thereby improving learning from suboptimal data. The advantage is estimated using the Q-value ensemble:

$$\omega_{\text{adv}} = \text{clip}\Big(\exp\Big(\tfrac{A(s,a_b^\star)}{\tau_{\text{adv}}}\Big),\ \varepsilon,\ c_{\text{adv}}\Big), \quad A(s,a) := \Big(\tfrac{1}{K}\sum_k Q_{\phi_k}(s,a)\Big) - \tfrac{1}{|\mathcal{A}|}\sum_{a'}\Big(\tfrac{1}{K}\sum_k Q_{\phi_k}(s,a')\Big). \tag{6}$$

(iii) Epistemic weighting. Building on ensemble-based uncertainty estimation and randomised priors (Lakshminarayanan et al., 2017; Osband et al., 2018; Pearce et al., 2018; Wilson & Izmailov, 2020), we emphasise samples with higher ensemble standard deviation, concentrating computation on regions of epistemic uncertainty so that the model learns most from poorly covered areas and provides a more informative posterior for exploration (σ_Q is from Eq. 3):

$$\omega_{\text{epi}} = \text{clip}(1 + \lambda_\sigma\ \sigma_Q(s,a_b^\star),\ \varepsilon,\ c_{\text{epi}})\,, \tag{7}$$

Learning a Q-value ensemble with TD(n) and Bayesian Shrinkage The Q-value ensemble is trained using only the logged context tuples, combining TD(n) regression Sutton et al. (1998); Dayan (1992) with Bayesian shrinkage Murphy (2007; 2012). For each context window of length H with transitions $\{(s_t,a_t,r_t,s_{t+1})\}_{t=1}^{H}$, the ensemble mean is $\bar{Q}(s,a) = \frac{1}{K}\sum_{k=1}^{K} Q_{\phi_k}(s,a)$. A n-step bootstrapped targets per time step t can be constructed as:

$$y_t^{(n)} \;=\; \sum_{i=0}^{n-1}\gamma^i r_{t+i} \;+\; \gamma^n\,\mathbf{1}[\,t{+}n \le H\,]\ \max_{a'}\bar{Q}(s_{t+n},a'), \tag{8}$$

where the next n observed rewards are summed along the logged trajectory. A bootstrap term is added only if the context still contains a state s_{t+n}[2].

To learn the ensemble, the loss function is composed of two terms $\mathcal{L}_Q = \mathcal{L}_{\text{TD}} + \mathcal{L}_{\text{shrink}}$.

$\mathcal{L}_{\textbf{TD}}$ - TD(n) regression on taken actions. For each (s_t,a_t) the ensemble mean is regressed to the TD(n) target:

$$\mathcal{L}_{\text{TD}} \;=\; \mathbb{E}\Big[\big(\bar{Q}(s_t,a_t) - y_t^{(n)}\big)^2\Big]. \tag{9}$$

$\mathcal{L}_{\textbf{shrink}}$ - Bayesian shrinkage to per-action posterior means. To improve statistical stability, per-action predictions are shrunk toward conjugate posterior means computed from the same TD(n) targets. For each action a we form counts and empirical TD(n) averages over the context:

$$c_a \;=\; \sum_{t=1}^{H}\mathbb{1}[a_t = a], \qquad \bar{y}_a \;=\; \frac{\sum_{t:\,a_t=a} y_t^{(n)}}{\max(1,c_a)}.$$

With prior mean μ_0, prior variance v_0, and likelihood variance σ^2, the per-action posterior mean is

$$m_a^{\text{post}} \;=\; \underbrace{\frac{\sigma^2}{\sigma^2 + c_a v_0}}_{w_a}\,\mu_0 \;+\; \big(1-w_a\big)\,\bar{y}_a, \qquad w_a = \frac{\sigma^2}{\sigma^2 + c_a v_0}. \tag{10}$$

The following loss shrinks the per-action time-average of the ensemble toward m_a^{post} for actions observed in the context is:

$$\mathcal{L}_{\text{shrink}} \;=\; \frac{1}{\sum_a \mathbb{1}[c_a > 0]}\sum_{a:\,c_a>0}\Big(\underbrace{\tfrac{1}{H}\sum_{t=1}^{H}\bar{Q}(s_t,a)}_{\text{per-action average over context states}} - m_a^{\text{post}}\Big)^2. \tag{11}$$

[2]Bandits arise as the special case $n{=}1$ (and $\gamma{=}0$).

Training Objective The full training loss $\mathcal{L} = \mathcal{L}_\pi + \lambda_Q \mathcal{L}_Q + \lambda_{\text{anchor}} \mathcal{L}_{\text{anchor}}$ is optimised using AdamW (Loshchilov & Hutter, 2019). We checkpoint the transformer and heads jointly, and optionally detach policy weights ω_b during $\mathcal{L}_\pi$ computation to prevent Q-network gradient interference.

See Appendix A for the full algorithm and implementation details , and Appendix C.3 for an ablation isolating the effect of the weighting terms in Eq. (4)–(7).

3.2 Test-Time Bayesian Fusion of Context and Value Prior

This section presents the key component of our algorithm: a test-time controller that combines information from the ensemble prior and context, following a UCB principle for action selection.

At the query state s we form an action-wise posterior by combining the ensemble prior $(\bar{Q}, \sigma_Q)$ with state-weighted statistics extracted from the context. Let $w_t(s) \in [0, 1]$ denote a kernel weight that measures how similar context state s_t is to the query s. Instances of such kernels are uniform, cosine or RBF kernels Cleveland & Devlin (1988); Watson (1964). For each action, the state-weighted counts and targets are

$$c_a(s) = \sum_t w_t(s)\, \mathbb{1}[a_t = a], \qquad \tilde{y}_a(s) = \frac{\sum_t w_t(s)\, \mathbb{1}[a_t = a]\; y_t}{\max\big(1, c_a(s)\big)}. \tag{12}$$

The target y_t can be chosen as immediate reward or an n-step bootstrapped return :

$$y_t^{(n)} = \sum_{i=0}^{n-1} \gamma^i r_{t+i} + \gamma^n \max_{a'} \bar{Q}(s_{t+n}, a'). \tag{13}$$

Given Eq. 3, a choice of kernel, and the weighted evidence $\big(c_a(s), \tilde{y}_a(s)\big)$, SPICE composes a conjugate-style posterior per action by precision additivity Murphy (2007; 2012):

$$\text{prior:} \quad \mu_a^{\text{pri}} = \bar{Q}(a), \quad v_a^{\text{pri}} = \max\{\sigma_Q(a)^2,\ v_{\min}\}, \qquad \text{likelihood variance: } \sigma^2, \tag{14}$$

$$\text{posterior:} \quad v_a^{\text{post}} = \left(\frac{1}{v_a^{\text{pri}}} + \frac{c_a(s)}{\sigma^2}\right)^{-1}, \qquad m_a^{\text{post}} = v_a^{\text{post}}\left(\frac{\mu_a^{\text{pri}}}{v_a^{\text{pri}}} + \frac{c_a(s)\ \tilde{y}_a(s)}{\sigma^2}\right). \tag{15}$$

Based on this posterior distribution, we propose the following action selection:

- **Online**, the policy follows a posterior-UCB rule with exploration parameter $\beta_{\text{ucb}} > 0$ Auer (2002), allowing exploration and adaptation to the task:

$$a^\star \;=\; \arg\max_a \left(m_a^{\text{post}} + \beta_{\text{ucb}} \sqrt{v_a^{\text{post}}}\right). \tag{16}$$

- **Offline**, the policy act greedily: $a^\star = \arg\max_a m_a^{\text{post}}$.

Hyperparameter choices are listed in Appendix A.5 and the pseudocode for Bayesian fusion appears in Algorithm 1 (Appendix A.1).

3.3 SPICE: Overall Algorithm

We now introduce the full algorithm in Algo. 1. SPICE combines context and ensemble priors in a Bayesian update, providing calibrated posteriors for UCB-based exploration. Unlike prior works (Lee et al., 2023; Dai et al., 2024), this design enables coherent adaptation from suboptimal data (see Fig. 1 for a pipeline overview).

4 Regret bound of the SPICE algorithm

A key component of SPICE is the use of a posterior-UCB rule at inference time that leverages both ensemble prior and in-context data. Importantly, we show in this section that the resulting online controller achieves optimal logarithmic regret despite being pretrained on sub-optimal data. Any prior miscalibration from pretraining manifests only as a constant warm-start term without affecting the asymptotic convergence rate. We establish this result formally in a bandit setting and provide in the next sections empirical validation across bandit problems and extend to other MDPs.

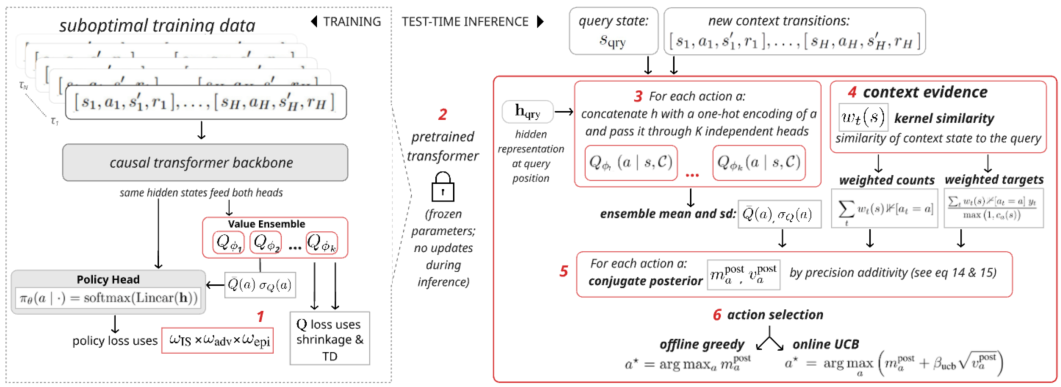

Figure 1: **Training and Test-Time Inference.** *Red boxes highlight novel contributions.* SPICE pretrains on suboptimal data using a transformer with (1) a novel value ensemble that learns calibrated uncertainty while improving policy training via advantage/epistemic reweighting. Gradient-free adaptation on new tasks: (2) frozen transformer processes query/context; (3) ensemble extracts value priors; (4) kernel weighting computes context evidence; (5) Bayesian fusion yields per-action posteriors; (6) action selection. SPICE learns explicit uncertainty from suboptimal data, enabling both improved training and principled gradient-free exploration.

4.1 Definitions and Assumptions

Consider a stochastic A-armed bandit setting with unknown means $\{\mu_a\}_{a=1}^{A} \subset \mathbb{R}$. At each round $t \in 1, ..., K$ the algorithm chooses a_t and receives a reward $r_t = \mu_{a_t} + \varepsilon_t$, where $(\varepsilon_t)_{t \geq 1}$ are independent mean zero $\sigma-$sub-Gaussian noise variables. The best-arm mean is defined as $\mu_\star = \max_{a \in [A]} \mu_a$ and the gap of arm a as $\Delta_a = \mu_\star - \mu_a$.

Without loss of generality, we scale rewards so that means satisfy $\mu_a \in [0.1]$ for all $a \in [A]$. Hence $0 \leq \mu_\star - \mu_a \leq 1$ and the per-round regret is at most 1. Assuming that each reward distribution is σ^2-sub-Gaussian, a current assumption in bandit analysis(Whitehouse et al., 2023; Han et al., 2024), one can derive the following tail bound for any arm a and round $t \geq 1$ with $n_{a,t}$ pulls and empirical mean $\widehat{\mu}_{a,t}$ for all $\varepsilon > 0$

$$\Pr\left(\left|\widehat{\mu}_{a,t} - \mu_a\right| > \varepsilon\right) \leq 2\exp\left(-\frac{n_{a,t}\varepsilon^2}{2\sigma^2}\right) \tag{17}$$

By setting $\varepsilon = \sigma\sqrt{\frac{2\log t}{n_{a,t}}}$, one can show that with probability at least $1 - O(\frac{1}{t^2})$

$$\left|\widehat{\mu}_{a,t} - \mu_a\right| \leq \sigma\sqrt{\frac{2\log t}{n_{a,t}}}, \tag{18}$$

i.e the deviation of the empirical mean from the true mean is bounded by $\sigma\sqrt{2\log t / n_{a,t}}$ with high probability (Hoeffding's inequality; see (Hoeffding, 1963; Boucheron & Thomas, 2012)).

Definition 1 (SPICE posterior)**.** *Let the ensemble prior for arm a be Gaussian with mean μ_a^{pri} and variance $v_a^{pri} > 0$, estimated from the value ensemble at the query (see Section 3.2). The prior pseudo-count is defined as*

$$N_a^{pri} := \frac{\sigma^2}{v_a^{pri}}, \quad \Longrightarrow \quad m_{a,t}^{post} = \frac{N_a^{pri}\mu_a^{pri} + n_{a,t}\widehat{\mu}_{a,t}}{N_a^{pri} + n_{a,t}}, \quad v_{a,t}^{post} = \frac{\sigma^2}{N_a^{pri} + n_{a,t}} \tag{19}$$

where $n_{a,t}$ and $\widehat{\mu}_{a,t}$ are the number of pulls and the empirical mean of arm a up to round t (these updates follow Normal-Normal conjugacy; see Murphy, 2007; 2012.) .

Definition 2 (SPICE inference)**.** *SPICE acts using a posterior-UCB rule at inference time*

$$a_t \in \arg\max_{a \in \mathcal{A}} \left\{ m_{a,t-1}^{post} + \beta_t\sqrt{v_{a,t-1}^{post}} \right\}, \quad \beta_t = \sqrt{2\log t} \tag{20}$$

The schedule $\beta_t = \sqrt{2\log t}$ mirrors the classical UCB1 analysis (Auer et al., 2002).

4.2 Inference-time Regret Bound

We now derive a regret bound for SPICE inference-time controller. The proof is given in Sec. B.

Theorem 1 (Regret-optimality with warm start)**.** *Under the assumption of σ^2-sub-Gaussian reward distributions, the SPICE inference controller satisfies*

$$\mathbb{E}\Big[\sum_{t=1}^{K}(\mu_\star - \mu_{a_t})\Big] \le \sum_{a\neq\star}\Big(\frac{32\sigma^2\log K}{\Delta_a} + 4N_a^{pri}\left|\mu_a^{pri} - \mu_a\right|\Big) + O(1). \tag{21}$$

Thus the cumulative regret of SPICE has an optimal logarithm rate in K and any sub-optimal pre-training results only in a constant warm-start term $\sum_{a\neq\star} 4N_a^{\text{pri}}\left|\mu_a^{\text{pri}} - \mu_a\right|$ that does not scale with K. The leading $O(\log K)$ term matches the classical UCB1 proof (Auer et al., 2002). The additive warm-start term depends on the prior pseudo-count $N_a^{\text{pri}} = \sigma^2/v_a^{\text{pri}}$, which behaves as prior data in a Bayesian sense (Gelman et al., 1995).

This theorem yields the following corollaries highlighting the impact of the prior quality on the regret bound.

Corollary 1 (Bound of well-calibrated priors)**.** *If the ensemble prior is perfectly calibrated, then $\mu_a^{pri} = \mu_a$ for all arms a and the warm-start term vanishes. SPICE then reduces to classical UCB*

$$\mathbb{E}[R_K] \le \sum_{a\neq\star}\frac{32\sigma^2\log K}{\Delta_a^2} + O(1). \tag{22}$$

Corollary 2 (Bound on weak priors)**.** *If the ensemble prior has infinite variance, $v_a^{pri} \to \infty$ and therefore $N_A^{pri} \to 0$. The warm-start term vanishes and SPICE reduces to classical UCB Eq. 22.*

The regret bound shows that SPICE inherits the optimal $O(\log K)$ rate of UCB while adding a constant warm-start cost from pretraining. The posterior mean in Eq. 19 is a convex combination of the empirical and prior means and the variance is shrinking at least as fast as $O(1/n_{a,t})$. Early decisions are influenced by the prior, but as $n_{a,t}$ grows, the bias term vanishes and learning relies entirely on observed rewards. A miscalibrated confident prior increases the warm-start constant but does not affect asymptotics, a well-calibrated prior eliminates the warm-start entirely and an uninformative prior ($v_a^{\text{pri}} \to \infty$) reduces SPICE to classical UCB. In practice, this means that SPICE can exploit structure from suboptimal pretraining when it is useful, while remaining safe in the long run, as its regret matches UCB regardless of the prior quality.

5 Learning in Bandits

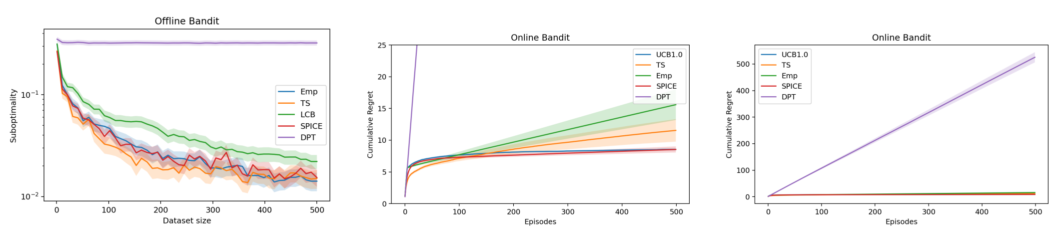

(a) Offline: suboptimality vs. context size h. Lower is better.

(b) Online: cumulative regret (zoomed). Lower is better.

(c) Online: cumulative regret (full scale).

Figure 2: **Bandit performance evaluation.** **(a)** Offline selection quality. **(b)** Online cumulative regret (zoomed view). **(c)** Online cumulative regret (full scale). Shaded regions are $\pm$ SEM over N=200 test environments.

We test our algorithm using the DPT evaluation protocol (Lee et al., 2023). Each task is a stochastic A-armed bandit with Gaussian rewards. Unless noted, A=5 and horizon H=500. Further details are given in Appendix C and Appensix A.5.

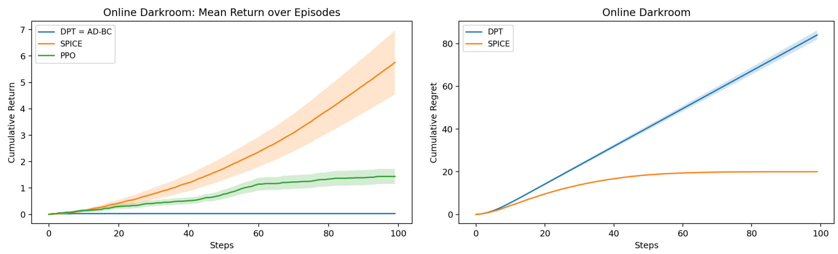

(a) Online: cumulative return over H=100 steps (higher is better).

(b) Online: cumulative regret (lower is better).

Figure 4: **Darkroom (MDP) results.** Models are pretrained on uniformly collected, weak-last labeled trajectories and evaluated online on N=100 held-out tasks for H=100 steps. Shaded regions denote $\pm$ SEM across tasks.

Results are presented in Fig. 2 and Fig. 3. **Offline**, SPICE and TS achieve the lowest suboptimality across h, while LCB is competitive early but remains above TS/SPICE. DPT is flat and far from optimal in this weak-data regime. **Online**, SPICE attains the lowest cumulative regret among learned methods and tracks the classical UCB closely (Fig. 2b and Fig. 2c). Under increasing reward noise, SPICE, TS, UCB, and Emp degrade smoothly with small absolute changes, whereas DPT's final regret remains two orders of magnitude larger, indicating failure to adapt from weak logs (Fig. 3).

SPICE achieves logarithmic online regret from suboptimal pretraining. Its posterior-UCB controller inherits $O(\log H)$ regret, with any prior miscalibration contributing only a constant warm-start term; the empirical curves match this prediction. Even with non-optimal pretraining, Bayesian fusion quickly overrides prior bias as evidence accrues, while DPT remains tied to its supervised labels.

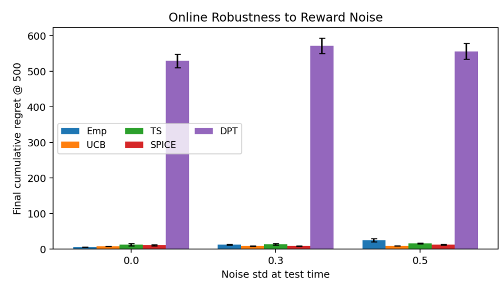

Figure 3: **Robustness to reward noise.** Final regret at H=500 for different noise levels ($\sigma \in \{0.0, 0.3, 0.5\}$). Bars are $\pm$SEM over N=200 test environments.

6 Learning in Markov Decision Processes

The Darkroom is a 10×10 gridworld with A=5 discrete actions and a sparse reward of 1 only at the goal cell. We pretrain on 100,000 environments using trajectories from a uniform behaviour policy and the "weak-last" label (the last action in the context), which provides explicitly suboptimal supervision. Testing uses N=100 held-out goals, horizon H=100, and identical evaluation for all methods. Further details are given in Appendix C and Appensix A.5.

Under weak supervision, DPT = AD-BC, as DPT is trained by cross-entropy to predict a single action label from the [query; context] sequence. With the "weak-last" dataset this label is simply the last action taken by a uniform behaviour policy. Algorithm Distillation (AD) with a behaviour-cloning teacher (AD-BC) optimises the same loss on the same targets, so both reduce to contextual behaviour cloning on suboptimal labels. Lacking reward-aware targets or calibrated uncertainty, the resulting policy remains bound to the behaviour and fails to adapt online, hence the flat returns and near-linear regret.

In this environment, SPICE adapts quickly and achieves high return with a regret curve that flattens after a short warm-up (Figs. 4a–4b). DPT, identical to AD-BC in this regime, exhibits near-linear regret and essentially zero return. We include PPO as a single-task RL reference for sample-efficiency; it improves but remains far below SPICE.

7 Discussion

SPICE addresses limitations of current ICRL methods using minimal changes to the sequence-modelling recipe: a lightweight value ensemble is attached to a shared transformer and learns the value prior at the query state; the transformer trunk is learned using a weighted loss to shape better representations feeding into the value ensemble; at inference, the value ensemble prior is fused with state-weighted statistics extracted from the provided context of the test task, resulting in per-action posteriors that can be used greedily offline or with a posterior-UCB rule for principled exploration online. SPICE is designed to learn a good-enough structural prior from the suboptimal data to leverage knowledge such as reward sparsity and consistent action effects across different environments. The value ensemble provides calibrated uncertainty that behaves as if the prior contributed a small number of virtual samples: it influences the posterior in the first few steps but is quickly outweighted as more data from the test environment is collected. This equips SPICE with two advantages: a strong warm start from weak data and principles posterior-UCB exploration, enabling rapid adaptation to new tasks and low regret in practice.

Theoretically, we show that SPICE achieves optimal $O(\log K)$ regret in stochastic bandits, with any pretraining miscalibration contributing only to a constant warm-start term. We validate this empirically, demonstrating that SPICE achieves logarithmic regret when trained on suboptimal data, while sequence-only ICRL baselines achieve lower return and linear regret (Fig. 4). Similarly, SPICE performs nearly optimal in offline selection on held-out tasks in weak data regimes, a setting where classic ICRL perform extremely poorly (Fig. 2a).

Limitations and Future Work. SPICE uses kernel-weighted counts to extrapolated state proximity at inference. The kernel choice can be critical in highly non-stationary or partially observable settings, where poorly chosen kernels can either over-fit or over-smooth context evidence. Additionally, the regret analysis assumes sub-Gaussian reward noise and focuses on the bandit regime; extending the guarantees to MDPs with long horizons is an exciting research direction. Another limitation is that SPICE assumes that the ensemble produces reasonably calibrated priors. If the prior is systematically misspecified, the posterior fusion may inherit its bias. This can slow early adaptation despite the regret guarantees.

8 Conclusion

We introduce SPICE, a Bayesian in-context reinforcement learning method that i) learns a value ensemble prior from suboptimal data via TD(n) regression and Bayesian shrinkage, ii) performs Bayesian context fusion at test time to obtain per-action posteriors and iii) acts with a posterior-UCB controller, performing principled exploration. The design is simple: attach lightweight value heads to a Transformer trunk and keep adaptation entirely gradient-free. SPICE addresses two persistent challenges in ICRL: behaviour-policy bias during pretraining and the lack of calibrated value uncertainty at inference. Theoretically, we show that the SPICE controller has optimal logarithmic regret and any pretraining miscalibration contributes only to a constant warm-start term. Empirical results show that SPICE achieves near-optimal offline decisions and online regret under distribution shift on bandits and control tasks.

9 Reproducibility Statement

We provide the details needed to reproduce all results. Algorithmic steps and test-time inference are given in Appendix A.1; model architecture, losses, and all hyperparameters are listed in Appendix A; data generators, evaluation protocols, and an ablation study are specified in Appendix C, with metrics, horizons, and noise levels matched to the DPT protocol Lee et al. (2023). Figures report means $\pm$ s.e.m. over the stated number of tasks and seeds, and we fix random seeds for every run. We use only standard benchmarks and public baselines; no external or proprietary data are required. We will release code, configuration files, and checkpoints upon publication to facilitate exact replication.

REFERENCES

Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of machine learning research*, 3(Nov):397–422, 2002.

Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.

Stéphane Boucheron and Maud Thomas. Concentration inequalities for order statistics. 2012.

David Brandfonbrener, Alberto Bietti, Jacob Buckman, Romain Laroche, and Joan Bruna. When does return-conditioned supervised learning work for offline reinforcement learning? *Advances in Neural Information Processing Systems*, 35:1542–1553, 2022.

Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.

William S Cleveland and Susan J Devlin. Locally weighted regression: an approach to regression analysis by local fitting. *Journal of the American statistical association*, 83(403):596–610, 1988.

André Correia and Luís A Alexandre. Hierarchical decision transformer. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1661–1666. IEEE, 2023.

Zhenwen Dai, Federico Tomasi, and Sina Ghiassian. In-context exploration-exploitation for reinforcement learning. *arXiv preprint arXiv:2403.06826*, 2024.

Peter Dayan. The convergence of td (λ) for general λ. *Machine learning*, 8(3):341–362, 1992.

Juncheng Dong, Moyang Guo, Ethan X Fang, Zhuoran Yang, and Vahid Tarokh. In-context reinforcement learning from suboptimal historical data.

Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. Rl^2: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.

Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*, 2019.

Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 1995.

Qiyang Han, Koulik Khamaru, and Cun-Hui Zhang. Ucb algorithms for multi-armed bandits: Precise regret and adaptive inference, 2024. URL `https://arxiv.org/abs/2412.06126`.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.

Michael Janner, Qiyang Li, and Sergey Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34:1273–1286, 2021.

Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International conference on machine learning*, pp. 652–661. PMLR, 2016.

Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169*, 2021.

Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.

Michael Laskin, Luyu Wang, Junhyuk Oh, Emilio Parisotto, Stephen Spencer, Richie Steigerwald, DJ Strouse, Steven Hansen, Angelos Filos, Ethan Brooks, et al. In-context reinforcement learning with algorithm distillation. *arXiv preprint arXiv:2210.14215*, 2022.

Jonathan Lee, Annie Xie, Aldo Pacchiano, Yash Chandak, Chelsea Finn, Ofir Nachum, and Emma Brunskill. Supervised pretraining can learn in-context reinforcement learning. *Advances in Neural Information Processing Systems*, 36:43057–43083, 2023.

Kuang-Huei Lee, Ofir Nachum, Mengjiao Sherry Yang, Lisa Lee, Daniel Freeman, Sergio Guadarrama, Ian Fischer, Winnie Xu, Eric Jang, Henryk Michalewski, et al. Multi-game decision transformers. *Advances in neural information processing systems*, 35:27921–27936, 2022.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019. URL https://arxiv.org/abs/1711.05101.

Amir Moeini, Jiuqi Wang, Jacob Beck, Ethan Blaser, Shimon Whiteson, Rohan Chandra, and Shangtong Zhang. A survey of in-context reinforcement learning. *arXiv preprint arXiv:2502.07978*, 2025.

Kevin P Murphy. Conjugate bayesian analysis of the gaussian distribution. *def*, 1(2σ2):16, 2007.

Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.

Ashvin Nair, Abhishek Gupta, Murtaza Dalal, and Sergey Levine. Awac: Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.

Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped dqn. *Advances in neural information processing systems*, 29, 2016.

Ian Osband, John Aslanides, and Albin Cassirer. Randomized prior functions for deep reinforcement learning. *Advances in neural information processing systems*, 31, 2018.

Tim Pearce, Mohamed Zaki, Alexandra Brintrup, N Anastassacos, and A Neely. Uncertainty in neural networks: Bayesian ensembling. *stat*, 1050:12, 2018.

Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.

Alexandre Piche, Rafael Pardinas, David Vazquez, Igor Mordatch, and Chris Pal. Implicit offline reinforcement learning via supervised learning. *arXiv preprint arXiv:2210.12272*, 2022.

Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *International conference on machine learning*, pp. 5331–5340. PMLR, 2019.

Daniel J Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, Zheng Wen, et al. A tutorial on thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.

Adith Swaminathan and Thorsten Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *International conference on machine learning*, pp. 814–823. PMLR, 2015a.

Adith Swaminathan and Thorsten Joachims. The self-normalized estimator for counterfactual learning. *advances in neural information processing systems*, 28, 2015b.

Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International conference on machine learning*, pp. 2139–2148. PMLR, 2016.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Qing Wang, Jiechao Xiong, Lei Han, Han Liu, Tong Zhang, et al. Exponentially weighted imitation learning for batched historical data. *Advances in Neural Information Processing Systems*, 31, 2018.

Geoffrey S Watson. Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 359–372, 1964.

Justin Whitehouse, Aaditya Ramdas, and Steven Wu. On the sublinear regret of GP-UCB. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=G3aubF5Wnw.

Andrew G Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33:4697–4708, 2020.

Qinqing Zheng, Amy Zhang, and Aditya Grover. Online decision transformer. In *international conference on machine learning*, pp. 27042–27059. PMLR, 2022.

APPENDIX

A IMPLEMENTATION AND EXPERIMENTAL DETAILS

A.1 SPICE ALGORITHM AND ARCHITECTURE

Algorithm 1 SPICE: Training and Test-Time Bayesian Fusion

0: **Inputs:** ensemble size K, prior scale α, horizon H, discount γ, TD(n) length n, kernel (ϕ, τ), noise variance σ^2, prior-variance floor $v_{\min}$
0: **Model:** GPT-2 trunk; policy head π_θ; value heads $Q_{\phi_k} = f_k + \alpha\, p_k$ with frozen priors p_k
0: **Training loop (contexts)**:
0: **for** batch $\{(s_t, a_t, r_t, s_{t+1})_{t=1}^{H},\ a^\star\}$ **do**
0: Encode *[query; context]* with the transformer
0: Obtain logits π_θ, ensemble values $Q_{\phi_{1:K}}$; define $\bar{Q}$, σ_Q
0: Compute weights $\omega = \omega_{\text{IS}} \cdot \omega_{\text{adv}} \cdot \omega_{\text{epi}}$
0: Update policy with weighted cross-entropy $\mathcal{L}_\pi$
0: Update value heads with TD(n) regression + conjugate shrinkage + anchor regulariser
0: **end for**
0: **Test-time decision (query state s with context $\mathcal{C}$)**:
0: Run transformer to get prior $(\bar{Q}(a), \sigma_Q(a))$
0: Form state-weighted evidence $(c_a(s), \tilde{y}_a(s))$ via kernel weights
0: Fuse prior and evidence by precision additivity to get posterior $(m_a^{\text{post}}, v_a^{\text{post}})$
0: Select action $a^\star = \arg\max_a \left(m_a^{\text{post}} + \beta_t \sqrt{v_a^{\text{post}}}\right)$ (UCB) or $\arg\max_a m_a^{\text{post}}$ (greedy) =0

A.2 INTUITION AND DESIGN CHOICES

Our goal is to make the model act as if it had a task-specific Bayesian posterior over action values at the query state.

- **Learn a good prior from suboptimal data.** Rather than requiring optimal labels or learning histories, we attach a lightweight ensemble of Q-heads to a DPT-style Transformer trunk. We train this ensemble using TD(n) regression and Bayesian shrinkage to conjugate per-action means computed from the offline dataset, resulting in a calibrated per-action value prior (mean and variance).
- **Why an ensemble?** Diversity across heads (encouraged by randomised priors and anchoring) captures epistemic uncertainty in areas where the training data provides limited guidance. This uncertainty is needed to perform coherent exploration and for mitigating the effect of suboptimal or incomplete training data.
- **Why a Transformer trunk?** The causal trunk provides a shared representation that conditions on the entire in-task context (state, actions, rewards). This enables the value heads to output prior estimates that are task-aware at the query state, while preserving the simplicity and scalability of sequence modelling.
- **Why train a policy head if we act with the posterior?** We train a policy-head with a propensity-advantage-epistemic weighted cross-entropy loss. Although we do not use this head for control at test time, it corrects the behaviour-policy bias during representation learning, allocated learning capacity to high-value and high-uncertainty examples and co-trains the trunks so that the Q ensemble receives inputs that facilitate reliable value estimation. Decoupling learning (policy supervision improves the trunk) from acting (posterior-UCB uses value uncertainty) is key to achieve robustness from suboptimal training data.
- **Inference time control.** At test time we adapt by performing Bayesian context fusion: we treat the transitions in the context dataset as local evidence about the value of each action near the query state, weight them by similarity to the query (via a kernel) and combine this evidence with the learned value prior. The results is a closed-form posterior mean and variance for every action. This allows the agent to i) exploit the prior knowledge when the context is scarce or empty when interacting with a new environment, ii) update

flexibly as more task-specific evidence accumulates and iii) act either conservatively offline (greedy with respect to the posterior mean) or optimistically online (using a UCB rule for exploration). Thus, adaptation produces coherent exploration and strong offline choices entirely through inference, without any gradient updates.

A.3 Additional Related Work

Uncertainty for exploration in deep RL. Bootstrapped DQN Osband et al. (2016) and randomised prior functions Osband et al. (2018) introduce randomised value functions and explicit priors for deep exploration. Deep ensembles provide strong, simple uncertainty estimates Lakshminarayanan et al. (2017), and "anchored" ensembles justify ensembling as approximate Bayesian inference by regularising weights toward prior draws Pearce et al. (2018). SPICE adapts the randomised prior principle to the ICRL setting with an ensemble of value heads and uses a Normal–Normal fusion at test time to produce posterior estimates that feed a UCB-style controller.

Our weighted pretraining objective is conceptually related to advantage-weighted policy learning. AWR performs supervised policy updates with exponentiated advantage weights Peng et al. (2019); AWAC extends this to offline-to-online settings Nair et al. (2020); IQL attains strong offline performance with expectile (upper-value) regression and advantage-weighted cloning Kostrikov et al. (2021). Propensity weighting and counterfactual risk minimisation (IPS/SNIPS/DR) provide a principled basis for importance-weighted objectives under covariate shift Swaminathan & Joachims (2015a;b); Jiang & Li (2016); Thomas & Brunskill (2016). These methods are single-task and do not yield a test-time value posterior for across-task in-context adaptation, which is our focus

RL via supervised learning and return conditioning. Beyond DT, the broader RL-via-supervised-learning literature includes return-conditioned supervised learning (RCSL) and analyses of when it recovers optimal policies Brandfonbrener et al. (2022). Implicit Offline RL via Supervised Learning Piche et al. (2022) unifies supervised formulations with implicit models and connects to return-aware objectives. These works motivate our supervised components but do not attach an explicit, calibrated posterior used for a principled controller at test time.

A.4 Practical Guidance

- Ensemble size. A small K (e.g., 5–10) already gives reliable uncertainty due to trunk sharing and randomised priors.
- Shrinkage. Moderate shrinkage stabilises training under weak supervision; too much shrinkage can understate uncertainty.
- TD(n). Larger n reduces bootstrap bias but increases variance; we found mid-range n helpful in sparse-reward MDPs.
- Kernels. Uniform kernels are sufficient for bandits; RBF or cosine kernels help in MDPs with structured state similarity.
- Exploration parameter. β_{ucb} tunes optimism; our theory motivates $\beta_t \propto \sqrt{\log t}$, with a fixed β working well in short-horizon evaluations.

A.5 Implementation Details

A.5.1 Bandit Algorithms

We follow the baselines and evaluation protocol of Lee et al. (2023). We report offline suboptimality and online cumulative regret, averaging over N tasks; for SPICE and DPT we additionally average over three seeds.

Empirical Mean (Emp). Greedy selection by empirical means: $\hat{a} \in \arg\max_a \hat{\mu}_a$, where $\hat{\mu}_a$ is the sample mean of rewards for arm a. Offline we restrict to arms observed at least once; online we initialise with one pull per arm (standard good-practice).

Upper Confidence Bound (UCB). Optimistic exploration using a Hoeffding bonus. At round t, pick $\hat{a} \in \arg\max_a \left(\hat{\mu}_{a,t} + \sqrt{1/n_{a,t}}\right)$, with $n_{a,t}$ pulls of arm a. UCB has logarithmic regret in stochastic bandits.

Lower Confidence Bound (LCB). Pessimistic selection for offline pick-one evaluation: $\hat{a} \in \arg\max_a \left(\hat{\mu}_a - \sqrt{1/n_a}\right)$. This favours well-sampled actions and is a strong offline baseline when datasets are expert-biased.

Thompson Sampling (TS) Bayesian sampling with Gaussian prior; we set prior mean $1/2$ and variance $1/12$ to match $\mu_a \sim \mathrm{Unif}[0, 1]$ in the DPT setup, and use the correct noise variance at test time.

DPT. Decision-Pretrained Transformer: a GPT-style model trained to predict the optimal action given a query state and an in-context dataset. Offline, DPT acts greedily; online, it samples actions from its policy (as in Lee et al. (2023)), which empirically yields UCB/TS-level exploration and robustness to reward-noise shifts, but only when trained on optimal data.

SPICE. Uncertainty-aware ICRL with a value-ensemble prior and Bayesian test-time fusion. At the query, SPICE forms a per-action posterior from (i) the ensemble prior mean/variance and (ii) state-weighted context statistics, then acts either greedily (offline) or with a posterior-UCB rule (online). The controller attains optimal $O(\log H)$ regret with any prior miscalibration entering only as a constant warm-start term.

A.5.2 RL Algorithms

We compare to the same meta-RL and sequence-model baselines used in Lee et al. (2023), and deploy SPICE/DPT in the same in-context fashion.

Proximal Policy Optimisation (PPO). Single-task RL trained from scratch (no pretraining); serves as an online-only point of reference for sample efficiency in our few-episode regimes. Hyperparameters follow common practice (SB3 defaults in our code) Schulman et al. (2017).

Algorithm Distillation (AD). A transformer trained via supervised learning on multi-episode learning traces of an RL algorithm; at test time, AD conditions on recent history to act in-context Laskin et al. (2022).

DPT. The same DPT model as described above but applied to MDPs: offline greedy; online sampling from the predicted action distribution each step Lee et al. (2023).

SPICE. The same SPICE controller: posterior-mean (offline) and posterior-UCB (online) built from an ensemble value prior and Bayesian context fusion at test time.

A.5.3 Bandit Pretraining and Testing

Task generator and evaluation. Each task is a stochastic A-armed bandit with $\mu_a \sim \mathrm{Unif}[0, 1]$ and rewards $r \sim \mathcal{N}(\mu_a, \sigma^2)$. Default: A=5, H=500, σ=0.3. We report offline suboptimality $\mu^\star - \mu_{\hat{a}}$ vs. context length h and online cumulative regret $\sum_{t=1}^{H}(\mu^\star - \mu_{a_t})$, averaging across N=200 test environments; for SPICE/DPT we additionally average across 3 seeds and plot $\pm$ SEM bands. For robustness we fix arm means and sweep $\sigma \in \{0.0, 0.3, 0.5\}$.

Pretraining. **DPT:** 100,000 training bandits; trunk n_{layer}=6, n_{emb}=64, n_{head}=1, dropout 0, AdamW (lr $= 10^{-4}$), 300 epochs, shuffle, seeds $\{0, 1, 2\}$. **SPICE:** same trunk; K=7 Q-heads with randomised priors and a small anchor penalty. We optimise a combined objective (policy cross-entropy with propensity/advantage/epistemic weighting for trunk shaping, plus value loss with TD(n) regression and shrinkage). Unless noted, we use uniform kernel weights for bandits at test time.

Controllers and deployment. Offline: given a fixed context, each method outputs a single arm; SPICE uses $\arg\max_a m_a^{\text{post}}$. Online: methods interact for H steps from empty context; SPICE uses $\arg\max_a \left(m_a^{\text{post}} + \beta\sqrt{v_a^{\text{post}}}\right)$. We match the DPT evaluation by using the same dataset generator, the same number of environments, and identical horizon and noise settings Lee et al. (2023).

Why SPICE succeeds under weak supervision (intuition). The value ensemble provides a calibrated prior that behaves like a small virtual sample count for each arm. Bayesian fusion then combines this prior with weighted empirical evidence, so the posterior rapidly concentrates as data accrues, shrinking any pretraining bias. Our theory shows this yields $O(\log H)$ **regret with only a constant warm-start penalty** from prior miscalibration; the curves in Fig. 2c–2b mirror this behaviour.

A.5.4 Darkroom Pretraining and Testing

Environment and data. We use a continuous darkroom navigation task in which rewards are smooth and peaked around a latent goal location. Each state is represented by a d-dimensional feature vector (default d=10). Actions are discrete with cardinality A; dynamics are deterministic given the current state and a one-hot action. For evaluation we generate N=100 held-out tasks of horizon H=100 and form an in-context dataset per task consisting of tuples $(s_t, a_t, r_t, s_{t+1})_{t=1}^{H}$. Unless stated otherwise, we use the "weak-last" split from our data generator (the same split is used for all methods).

Pretraining. Both SPICE and DPT share the same GPT-style trunk (n_{layer}=6, n_{emb}=64, n_{head}=1, dropout 0), trained with AdamW at learning rate 10^{-4} for 50 epochs.[3] DPT is trained with the standard DPT objective on 100,000 darkroom tasks (shuffled mini-batches). SPICE attaches an ensemble of K=7 value heads with randomised priors and trains them via TD(n) regression with n=5 and γ=0.95, plus conjugate shrinkage and a small anchor penalty (see Alg. 1). All hyperparameters used by the test-time Bayesian fusion are fixed a priori: RBF kernel with scale τ=0.5, evidence noise σ^2=0.09, and prior-variance floor $v_{\min}$=10^{-2}.

Controllers at test time. For SPICE we evaluate posterior-UCB with three optimism levels, $\beta \in \{0.5, 1.0, 2.0\}$; the offline analogue uses the posterior mean (greedy). For DPT we use the greedy controller that selects $\arg\max_a$ of the policy logits at the query state. When averaging across seeds, we first average per task across the three checkpoints and then aggregate across tasks; error bands report $\pm$ SEM.

Evaluation protocol. We report two metrics: (i) *Online return*: (ii) *Online cumulative regret*: starting from an empty context, a controller interacts for H steps; at each step we compare the reward of the chosen action to the reward of the environment's optimal action at the same state. To ensure a fair comparison, for each held-out task we draw a single initial state s_0 and use it for all controllers and seeds before averaging. For each metric we average across the N=100 held-out tasks. We average over three seeds. Shaded regions denote $\pm$ SEM across tasks.

B Proof of Theorem 1

Proof Overview. We analyse the posterior-UCB controller by (i) treating the ensemble prior at the query as a Normal prior with mean μ_a^{pri} and variance v_a^{pri}, resulting in a posterior with pseudo-count $N_a^{\text{pri}} = \sigma^2/v_a^{\text{pri}}$ under Normal-Normal conjugacy (Murphy, 2007; 2012); (ii) showing that the posterior mean is a convex combination of the empirical and prior means and the posterior variance shrinks at least as $O(1/n_{a,t})$ (Lemma 1); and (iii) combining sub-Gaussian concentration (Hoeffding-style) with a UCB schedule $\beta_t = \sqrt{2\log t}$ (Hoeffding, 1963; Auer et al., 2002) to upper-bound pulls of suboptimal arms. This results in $O(\log K)$ regret plus a constant warm-start term proportional to $N_a^{\text{pri}}|\mu_a^{\text{pri}} - \mu_a|$ (Lemma 2), recovering classical UCB when the prior is uninformative or well calibrated.

[3] We train three seeds for each method; checkpoints are averaged only at evaluation time.

For completeness, we re-state the theorem here: *Under the assumption of σ^2-sub-Gaussian reward distributions, the SPICE inference controller satisfies*

$$\mathbb{E}\Big[\sum_{t=1}^{K}(\mu_\star - \mu_{a_t})\Big] \le \sum_{a\neq\star}\Big(\frac{32\sigma^2 \log K}{\Delta_a} + 4N_a^{\text{pri}}\left|\mu_a^{\text{pri}} - \mu_a\right|\Big) + O(1).$$

First, we consider the following lemmas

Lemma 1 (Bias-variance decomposition)**.** *With the posterior defined in Eq. 19, for all a, t it holds that*

$$\left|m_{a,t}^{post} - \mu_a\right| \le \left|\widehat{\mu}_{a,t} - \mu_a\right| + \frac{N_a^{pri}}{N_a^{pri} + n_{a,t}}\left|\mu_a^{pri} - \mu_a\right|, \quad v_{a,t}^{post} \le \frac{\sigma^2}{n_{a,t}} \tag{23}$$

This lemma shows that the posterior mean forms a weighted average of the empirical and prior means, with relative error decomposing into two components: a variance term $\left|\widehat{\mu}_{a,t} - \mu_a\right|$ capturing finite-sample noise in the empirical mean, and a bias term $\frac{N_a^{\text{pri}}}{N_a^{\text{pri}}+n_{a,t}}\left|\mu_a^{\text{pri}} - \mu_a\right|$ reflecting prior miscalibration. As $n_{a,t} \to \infty$, the bias term vanishes, eliminating prior miscalibration, while the posterior variance shrinks at least as fast as the frequentist variance $\frac{\sigma^2}{n_{a,t}}$ (see Eq. 23).

Proof. The posterior mean for arm a at round t from equation Eq. 19 can be rewritten as a convex combination

$$m_{a,t}^{\text{post}} = \alpha_{a,t}\widehat{\mu}_{a,t} + (1-\alpha_{a,t})\mu_a^{\text{pri}}, \quad \alpha_{a,t} = \frac{n_{a,t}}{N_a^{\text{pri}} + n_{a,t}}$$

Subtracting the true mean μ_a gives

$$m_{a,t}^{\text{post}} - \mu_a = \alpha_{a,t}(\widehat{\mu}_{a,t} - \mu_a) + (1-\alpha_{a,t})(\mu_a^{\text{pri}} - \mu_a)$$

Taking absolute values and applying the triangle inequality gives

$$\left|m_{a,t}^{\text{post}} - \mu_a\right| \le \alpha_{a,t}\left|(\widehat{\mu}_{a,t} - \mu_a)\right| + (1-\alpha_{a,t})\left|(\mu_a^{\text{pri}} - \mu_a)\right|$$

Since $\alpha_{a,t} \le 1$ we can drop the factor and since $1 - \alpha_{a,t} = \frac{N_a^{\text{pri}}}{N_a^{\text{pri}}+n_{a,t}}$, we obtain

$$\left|m_{a,t}^{\text{post}} - \mu_a\right| \le \left|(\widehat{\mu}_{a,t} - \mu_a)\right| + \frac{N_a^{\text{pri}}}{N_a^{\text{pri}} + n_{a,t}}\left|(\mu_a^{\text{pri}} - \mu_a)\right|$$

Since $N_a^{\text{pri}} \ge 0$ we get

$$v_{a,t}^{\text{post}} = \frac{\sigma^2}{N_a^{\text{pri}} + n_{a,t}} \le \frac{\sigma^2}{n_{a,t}}.$$

$\square$

Lemma 2 (Posterior concentration)**.** *We fix a horizon $K \ge 2$. Under the assumption of σ^2-sub-Gaussian reward distributions, the following inequality holds simultaneously for all arms a and all rounds $t \in \{1, ..., K\}$ with probability at least $1 - O(\frac{1}{K})$*

$$\mu_a \le m_{a,t}^{post} + \beta_t\sqrt{v_{a,t}^{post}} + \frac{N_a^{pri}}{N_a^{pri} + n_{a,t}}\left|\mu^{pri} - \mu_a\right|, \quad \beta_t = \sqrt{2\log t}$$

Note that this also yields a symmetric lower bound with the last two terms negated.

Proof. Using Eq. 18 and a union bound over all a and $t \le K$ we obtain the following bound with probability at least $1 - O(1/K)$

$$|\widehat{\mu}_{a,t} - \mu_a| \le \sigma\sqrt{\frac{2\log t}{n_{a,t}}} \quad \text{for all } a \text{ and } t \le K$$

Since $v_{a,t}^{\text{post}} = \frac{\sigma^2}{N_a^{\text{pri}}+n_{a,t}} \le \frac{\sigma^2}{n_{a,t}}$ we get

$$\sigma\sqrt{\frac{2\log t}{n_{a,t}}} \le \sqrt{2\log t}\sqrt{v_{a,t}^{\text{post}}} = \beta_t\sqrt{v_{a,t}^{\text{post}}}$$

and thus

$$|\widehat{\mu}_{a,t} - \mu_a| \le \beta_t\sqrt{v_{a,t}^{\text{post}}}.$$

Combining that with Lemma 1 gives

$$\left|m_{a,t}^{\text{post}} - \mu_a\right| \le \beta_t\sqrt{v_{a,t}^{\text{post}}} + \frac{N_a^{\text{pri}}}{N_a^{\text{pri}} + n_{a,t}}\left|\mu_a^{\text{pri}} - \mu_a\right|.$$

Expanding this absolute value bound into one-sided inequalities yields the result. □

Using these lemmas, the proof of Theorem 1 follows.

Proof. Let $N_a(K) := \sum_{t=1}^{K} \mathbf{1}\{a_t = a\}$ be the pull count of arm a up to horizon K. We can decompose the regret as $\mathbb{E}\big[\sum_{t=1}^{K}(\mu_\star - \mu_{a,t})\big] = \sum_{a\neq\star}\Delta_a\mathbb{E}\big[N_a(K)\big]$. We derive an upper bound for $N_a(K)$ for each suboptimal arm a.

Consider an horizon $K \ge 2$, define the good event for each arm $a \in [A]$ and step $t \in \{1, ..., K\}$

$$G_{a,t} := \left\{\left|\widehat{\mu}_{a,t} - \mu_a\right| \le \sigma\sqrt{\tfrac{2\log t}{n_{a,t}}}\right\}$$

$G_{a,t}$ is the event that the empirical mean of arm a at time t lies within its confidence interval. We define the event that concentration holds for all arms and times simultaneously

$$\mathcal{E} := \bigcap_{a=1}^{A}\bigcap_{t=1}^{K} G_{a,t}.$$

The complement corresponds to the event that concentration fails for at least one (a, t)

$$\mathcal{E}^c := \bigcup_{a=1}^{A}\bigcup_{t=1}^{K} G_{a,t}^c.$$

Using the union bound, we get

$$\Pr(\mathcal{E}^c) \le \sum_{a=1}^{A}\sum_{t=1}^{K}\Pr(G_{a,t}^c) \le \sum_{a=1}^{A}\sum_{t=1}^{K}\frac{2}{t^2} \le 2A\sum_{t=1}^{\infty}\frac{1}{t^2} = \frac{\pi^2}{3}A.$$

If we instead define $G_{a,t}$ using an inflated radius $\sigma\sqrt{\frac{2\log(cAK^2)}{n_{a,t}}}$ we similarly get $\Pr(\mathcal{E}^c) \le O(\frac{1}{K})$.

We decompose the regret as

$$\mathbb{E}[R_K] = \mathbb{E}[R_K \mid \mathcal{E}]\Pr(\mathcal{E}) + \mathbb{E}[R_K \mid \mathcal{E}^c]\ \Pr(\mathcal{E}^c)\ \le \mathbb{E}[R_K \mid \mathcal{E}] + K\Pr(\mathcal{E}^c) \le \mathbb{E}[R_K \mid \mathcal{E}] + O(1),$$

so it is sufficient to bound the regret on $\mathcal{E}$.

Using Lemma 1, knowing that $\sigma\sqrt{2\log t/n_{a,t}} \le \beta_t\sqrt{v_{a,t}^{\text{post}}}$ and that $|\widehat{\mu}_{a,t} - \mu_a| \le \sigma\sqrt{2\log t/n_{a,t}}$ for $\mathcal{E}$, we get

$$\mu_a \le m_{a,t}^{\text{post}} + \beta_t\sqrt{v_{a,t}^{\text{post}}} + \frac{N_a^{\text{pri}}}{N_a^{\text{pri}} + n_{a,t}}|\mu_a^{\text{pri}} - \mu_a|, \quad \beta_t = \sqrt{2\log t}. \tag{24}$$

Suppose we pick a suboptimal arm $a \neq \star$ at round t. Using Eq. 24 for a and $\star$ as well as the SPICE selection rule $m^{\text{post}}_{a,t-1} + \beta_t \sqrt{v^{\text{post}}_{a,t-1}} \geq m^{\text{post}}_{\star,t-1} + \beta_t \sqrt{v^{\text{post}}_{\star,t-1}}$ we get

$$\Delta_a \leq 2\beta_t \sqrt{v^{\text{post}}_{a,t-1}} + \frac{N^{\text{pri}}_a}{N^{\text{pri}}_a + n_{a,t}} |\mu^{\text{pri}}_a - \mu_a| + \frac{N^{\text{pri}}_\star}{N^{\text{pri}}_\star + n_{\star,t}} |\mu^{\text{pri}}_\star - \mu_\star|. \tag{25}$$

First, we derive a threshold for the variance term in Eq. 25. Using $\beta_t = \sqrt{2\log t}$, Lemma 1 and Eq. 23) we obtain

$$2\beta_t \sqrt{v^{\text{post}}_{a,t-1}} \leq 2\sqrt{2\log t}\frac{\sigma}{\sqrt{n_{a,t-1}}}.$$

Using a similar technique as in the classical UCB1 proof Auer et al. (2002) we make the variance term smaller than half the gap $\Delta_a/2$

$$2\sqrt{2\log t}\frac{\sigma}{\sqrt{n_{a,t-1}}} \leq \frac{\Delta_a}{2} \quad \Rightarrow \quad n_{a,t-1} \geq \frac{32\sigma^2 \log t}{\Delta_a^2}$$

As $t \leq K$ we can replace $\log t$ with the worst-case $logK$ to ensure that the condition holds for all rounds up to horizon K. The variance threshold is therefore

$$n^\dagger_a := \left\lceil \frac{32\sigma^2 \log K}{\Delta_a^2} \right\rceil.$$

Once arm a has been pulled at least $n^\dagger_a$ times, the variance term $2\beta_t \sqrt{v^{\text{post}}_{a,t-1}}$ in Eq. 25 is guaranteed to be at most $\Delta_a/2$ for every $t \leq K$.

Second, we derive a threshold for the prior bias terms. To force the prior bias term below $\Delta_a/4$ we define $\delta_a := |\mu^{\text{pri}}_a - \mu_a|$ and solve

$$\frac{N^{\text{pri}}_a}{N^{\text{pri}}_a + n_{a,t}}\delta_a \leq \frac{\Delta_a}{4} \quad \Rightarrow \quad n_{a,t-1} \geq \frac{4N^{\text{pri}}_a \delta_a}{\Delta_a} - N^{\text{pri}}_a \leq \frac{4N^{\text{pri}}_a \delta_a}{\Delta_a}.$$

Thus after about

$$n^{pri}_a := \left\lceil \frac{4N^{\text{pri}}_a \delta_a}{\Delta_a} \right\rceil$$

pulls of arm a its prior bias term is guaranteed to be below $\Delta_a/4$. The same argument applies to the optimal arm $\star$: its prior bias terms decreases as $n_{\star,t}$ grows and since $\star$ is selected frequently, only a constant number of pulls ins needed before its prior bias term is below $\Delta_A/4$.

By combining the bias and variance thresholds, we can derive the following bound for $N_a(K)$ under the event $\mathcal{E}$ for some constant C_a (independent of K)

$$\mathbb{E}[N_a(K)\mathbf{1}_\mathcal{E}] \leq n^\dagger_a + n^{\text{pri}}_a + C_a \leq \frac{32\sigma^2 \log K}{\Delta_a^2} + \frac{4N^{\text{pri}}_a \delta_a}{\Delta_a} + C_a.$$

By multiplying by Δ_a and summing over all arms $a \neq \star$ we obtain

$$\mathbb{E}[R_K \mathbf{1}_\mathcal{E}] \leq \sum_{a \neq \star} \frac{32\sigma^2 \log K}{\Delta_a^2} + \sum_{a \neq \star} 4N^{\text{pri}}_a \delta_a + O(1).$$

To conclude, we collect all bounded terms and include the contribution of the event $\mathcal{E}^c$ into an $O(1)$ term to obtain

$$\mathbb{E}\Big[\sum_{t=1}^{K} (\mu_\star - \mu_{a,t})\Big] \leq \sum_{a \neq \star} \frac{32\sigma^2 \log K}{\Delta_a^2} + \sum_{a \neq \star} 4N^{\text{pri}}_a |\mu^{pri}_a - \mu_a| + O(1).$$

$\square$

C Additional Experimental Results

C.1 Bandit setting

Setup and data. Each task is a stochastic A-armed bandit with i.i.d. arm means $\mu_a \sim \mathrm{Unif}[0,1]$ and Gaussian rewards $r \sim \mathcal{N}(\mu_a, \sigma^2)$. Unless noted, A=5, horizon H=500, and the default test noise is σ=0.3. We evaluate on N=200 held-out environments; for robustness we fix the means and sweep $\sigma \in \{0.0, 0.3, 0.5\}$ at test time. For SPICE/DPT we report the mean over 3 seeds. Offline we measure suboptimality $\mu^\star - \mu_{\hat{a}}$ as a function of context length h; online we report cumulative regret $\sum_{t=1}^{H}(\mu^\star - \mu_{a_t})$.

C.2 Ablation: Quality of training data

Setup. Each task is a stochastic A-armed bandit with i.i.d. arm means $\mu_a \sim \mathrm{Unif}[0,1]$ and rewards $r \sim \mathcal{N}(\mu_a, \sigma^2)$. We use A=20, horizon H=500, and default test noise σ=0.3. We evaluate on N=200 held-out environments.

Data generation (weakmix80). Following the DPT protocol, contexts are collected by a behaviour policy that mixes broad exploration with concentrated exploitation on one arm. Concretely, for each environment we form a per-arm distribution

$$p \;=\; (1-\omega)\,\mathrm{Dirichlet}(\mathbf{1}) \;+\; \omega\,\delta_{i^\star},$$

where $\delta_{i^\star}$ is a point mass on a single arm $i^\star$ (chosen uniformly at random for this experiment), and we fix the mix strength to ω=0.5. At each context step an action is drawn from p. Supervision is *weak*: the training label is generated in `mix` mode with probability q=0.8 using the true optimal arm, and with probability $1-q$ by sampling an arm from p. We denote this setting by **weakmix80**. We generate 100k training tasks and 200 evaluation tasks with the above roll-in and labels.

Models and deployment. DPT and SPICE share the same transformer trunk (6 layers, 64 hidden units, single head, no dropout). For this ablation we pretrain both for 100 epochs on the weakmix80 dataset with A=20. Offline, all methods select a single arm from a fixed context. Online, they interact for H steps starting from an empty context; SPICE acts with a posterior-UCB controller, DPT samples from its predicted action distribution, and classical bandits (Emp, UCB, TS) use standard update rules.

Results.

- **Offline .** DPT is competitive offline under weakmix80 (80% optimal labels), but still converges more slowly than TS/SPICE as h grows (Fig. 5a).
- **Online .** SPICE attains the lowest regret among learned methods and closely tracks UCB, while TS is slightly worse and Emp is clearly worse (Fig. 5b). In contrast, DPT exhibits near-linear growth in regret: it improves little with additional interaction despite 80% optimal labels.
- **Robustness to reward-noise shift.** SPICE, TS, UCB and Emp degrade smoothly as σ increases, with small absolute changes. DPT's final regret remains orders of magnitude larger is and essentially insensitive to σ, indicating failure to adapt from weak training data (Fig. 5c).

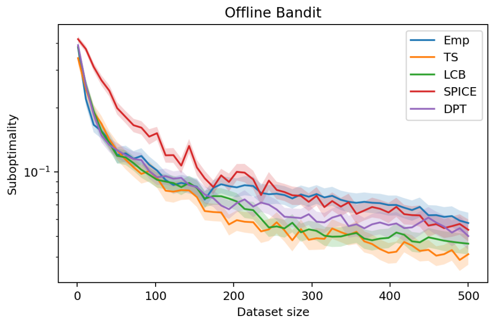

(a) Offline suboptimality vs. context size h.

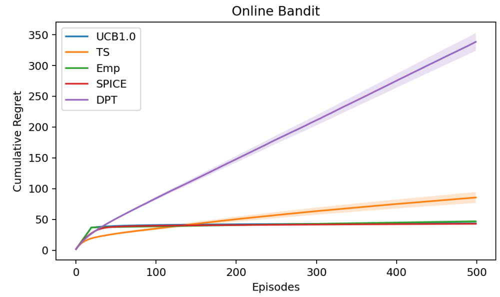

(b) Online cumulative regret.

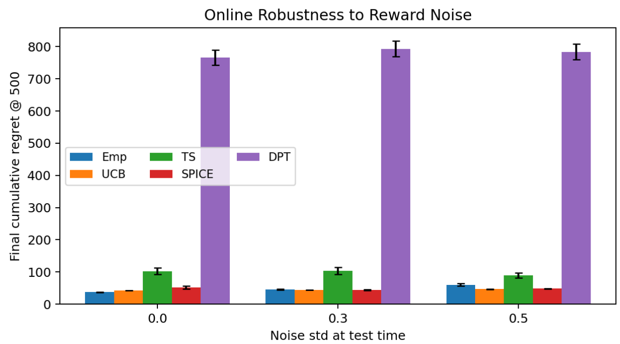

(c) Robustness: final regret at $H{=}500$ vs. σ.

Figure 5: **20-arm weak supervision (weakmix80).** Shaded regions/bars are $\pm$ s.e.m. over $N{=}200$ environments; SPICE/DPT averaged over 3 seeds.

C.3 Ablation: weight terms in training objective

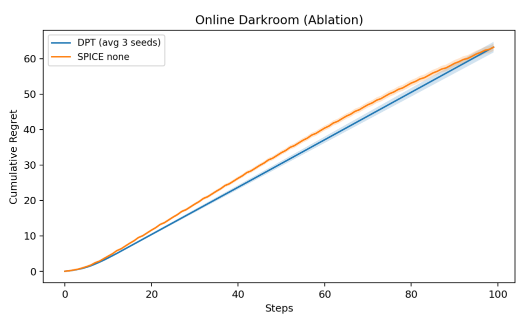

Figure 6: Online Darkroom ablation on weighting terms. We compare DPT against SPICE trained with no weights in the pretraining objective (both averaged over 3 seeds).

In this ablation, we studied the effect of the weighting terms in the SPICE objective (importance, advantage, epistemic). The results show that removing these terms degrades performance, highlighting their role in shaping better representations under weak data and reducing online regret.

D Use of LLMs.

ChatGPT was employed as a general-purpose assistant for enhancing writing clarity, conciseness, and tone, and providing technical coding support for plotting utilities and minor debugging tasks.

All outputs were verified by the authors, who retain full responsibility for research conception, algorithmic contributions, implementation, experimental findings, and manuscript writing.

E Ethics statement.

All authors have read and adhere to the ICLR Code of Ethics. This work does not involve human subjects, personally identifiable data, or sensitive attributes. We evaluate solely on synthetic bandit and control benchmarks and do not deploy in safety-critical settings. We discuss limitations (kernel choice, sub-Gaussian noise assumption, misspecified priors) and avoid claims beyond our experimental scope (Section 7). Our method could, in principle, be applied to high-stakes domains; we therefore emphasise the need for rigorous safety evaluation and domain-appropriate oversight before any real-world use.