

Information Directed Sampling for Linear Partial Monitoring

Johannes Kirschner

ETH Zurich, Department of Computer Science

JKIRSCHNER@INF.ETHZ.CH

Tor Lattimore

DeepMind

LATTIMORE@GOOGLE.COM

Andreas Krause

ETH Zurich, Department of Computer Science

KRAUSEA@ETHZ.CH

Abstract

Partial monitoring is a rich framework for sequential decision making under uncertainty that generalizes many well known bandit models, including linear, combinatorial and dueling bandits. We introduce *information directed sampling* (IDS) for stochastic partial monitoring with a linear reward and observation structure. IDS achieves adaptive worst-case regret rates that depend on precise observability conditions of the game. Moreover, we prove lower bounds that classify the minimax regret of all finite games into four possible regimes. IDS achieves the optimal rate in all cases up to logarithmic factors, without tuning any hyper-parameters. We further extend our results to the contextual and the kernelized setting, which significantly increases the range of possible applications.

Keywords: Information Directed Sampling, Linear Partial Monitoring, Bandits

1. Introduction

Partial monitoring is an expressive framework for sequential decision making in which the learner does not directly observe the reward (Rustichini, 1999). Instead, the learner obtains observations from pre-specified observation distributions that are associated to the actions and may or may not provide direct information about the reward. In this work, we consider a stochastic version of the problem with a linear reward and observation model, which is sometimes referred to as combinatorial partial monitoring (Lin et al., 2014; Chaudhuri and Tewari, 2016). Among other settings as described in Section 4, the linear partial monitoring model strictly generalizes linear bandits (Abe and Long, 1999; Auer, 2003), combinatorial bandits (Cesa-Bianchi and Lugosi, 2012) both with bandit and semi-bandit feedback, and some variants of dueling bandits (Yue and Joachims, 2009).

Linear Partial Monitoring Let $\mathcal{X} \subset \mathbb{R}^d$ be a compact set of actions and $\theta \in \mathbb{R}^d$ be an unknown parameter. For each action $x \in \mathcal{X}$, let $A_x \in \mathbb{R}^{d \times m}$ be a known *linear observation operator*. The learner and environment interact over n rounds. In each round t , the learner chooses an action $x_t \in \mathcal{X}$ and receives an m -dimensional observation $a_t = A_{x_t}^\top \theta + \epsilon_t$ where $(\epsilon_t)_{t=1}^n$ is a sequence of independent ρ -subgaussian noise vectors such that $\epsilon_t \in \mathbb{R}^m$. The reward for the learner is $\langle x_t, \theta \rangle$ and is *not observed*. As usual, the aim is to minimize cumulative regret

$$R_n = \sum_{t=1}^n \langle x^* - x_t, \theta \rangle,$$

where $x^* = \arg \max_{x \in \mathcal{X}} \langle x, \theta \rangle$ is the optimal action, chosen arbitrarily whenever the choice is not unique. A slightly more general formulation of the setup is in Appendix A.1, which we will use for some applications. *Bandit games* are a special case where $A_x = x$. We discuss further applications in detail in Section 4. Readers seeking further motivation and intuition for the setup will benefit from skipping ahead to this section.

A linear partial monitoring game is called *finite* if it has finitely many actions. An action $x \in \mathcal{X}$ is called *Pareto optimal* if it is an extreme point of the convex hull of \mathcal{X} . The set of actions that are optimal for θ is

$$\mathcal{P}(\theta) = \{x \in \mathcal{X} : \langle x, \theta \rangle = \max_{y \in \mathcal{X}} \langle y, \theta \rangle\},$$

which is defined on sets $\mathcal{C} \subseteq \mathbb{R}^d$ by $\mathcal{P}(\mathcal{C}) = \cup_{\theta \in \mathcal{C}} \mathcal{P}(\theta)$. A game is called *globally observable* if

$$x - y \in \text{span}(A_z : z \in \mathcal{X}) \text{ for all } x, y \in \mathcal{X}. \quad (1)$$

A game is called *locally observable* if for every convex set $\mathcal{C} \subset \mathbb{R}^d$,

$$x - y \in \text{span}(A_z : z \in \mathcal{P}(\mathcal{C})) \text{ for all } x, y \in \mathcal{P}(\mathcal{C}). \quad (2)$$

Intuitively, in globally observable games, the learner has access to actions from which reward differences between different actions can be estimated. In locally observable games, the reward differences can be estimated in a local sense that greatly eases learning. Any locally observable game is also globally observable. Although it is not important for this work, readers familiar with partial monitoring will be pleased to know these definitions coincide with the usual definitions, as discussed in Appendix F. At least for finite games, we will see that global observability leads to a $\tilde{O}(n^{2/3})$ regret, while local observability leads to $\tilde{O}(n^{1/2})$ regret. The relation to finite partial monitoring is subtle, however, and the classification results do not imply each other (see Appendix A.2).

Information Directed Sampling We propose a new algorithm for stochastic linear partial monitoring based on the *information directed sampling (IDS)* principle. This strategy uses the observations $a_t = A_{x_t}^\top \theta + \epsilon_t$ to construct conservative estimates $\Delta_t(x) \geq \langle x^* - x, \theta \rangle$ of the true gaps and an associated *information gain* $I_t(x)$, detailed below. The information gain quantifies the uncertainty reduction in the parameter estimate when the learner chooses $x \in \mathcal{X}$ and observes $A_x^\top \theta + \epsilon$. IDS is the policy that samples action x_t from a distribution μ_t that minimizes the *information ratio*,

$$\mu_t = \arg \min_{\mu} \frac{\mathbb{E}_{\mu}[\Delta_t(x)]^2}{\mathbb{E}_{\mu}[I_t(x)]}.$$

Our Contributions Our main contribution is a new algorithm for linear stochastic partial monitoring. We show that the same algorithm achieves the optimal regret dependence on the horizon in all finite-action games, without tuning any hyper-parameters. Along the way, we prove a classification theorem showing that, up to logarithmic factors, the minimax regret of all finite-action games is either 0, $\tilde{O}(n^{1/2})$, $\tilde{O}(n^{2/3})$ or $\Omega(n)$. This result mirrors that for the standard setting (Lattimore and Szepesvári, 2019a), but neither result implies the other. Our upper bounds are general and apply beyond the finite case. For infinite actions, however, the classification theorem is no longer so straightforward: we show that the minimax regret depends on finer geometric properties of the action set and observation structure, such as curvature. We further consider a novel *contextual partial monitoring setting*, where IDS exhibits an elegant planning behavior to exploit the distribution over contexts. Lastly, our algorithm and analysis are easily *kernelized*, which enables utilizing practically important smoothness priors, with applications such as Bayesian optimization with gradient observations.

Related work Finite partial monitoring dates back to Rustichini (1999). The generality of partial monitoring yields a rich structure of games (Bartók et al., 2014; Lattimore and Szepesvári, 2019a) where the minimax regret rate depends on precise observability conditions. The complete classification of finite games is achieved in a line of work by Cesa-Bianchi et al. (2006); Bartók et al. (2014); Antos et al. (2013); Lattimore and Szepesvári (2019a), with a focus on the stochastic version of the problem in the work by Bartók et al. (2011); Bartók et al. (2012). Asymptotics for finite games are known as well (Komiyama et al., 2015). Partial monitoring with prior information was studied by Vanchinathan et al. (2014), and with side-information by Bartók and Szepesvári (2012). Latter setup is different from our contextual setting. The linear version of the problem that we study here is due to Lin et al. (2014); Chaudhuri and Tewari (2016). Both previous approaches rely on explicit exploration schemes and achieve a $\tilde{O}(n^{2/3})$ worst-case regret on globally observable games, but not the faster $\tilde{O}(n^{1/2})$ rate on locally observable games. Information directed sampling was proposed by Russo and Van Roy (2014) in the Bayesian setting to address short-comings of the UCB algorithm (Auer, 2003) and Thompson sampling (Agrawal and Goyal, 2013) on examples that capture the spirit of partial monitoring. Liu et al. (2018) tailored the regret analysis of Bayesian IDS to bandits with graph feedback, a setting that is captured in the linear partial monitoring framework as well. The frequentist version of the algorithm that we analyze here was proposed by Kirschner and Krause (2018) in a bandit setting with heteroscedastic noise, which we strictly generalize. Recently, information theoretic tools were also introduced in the partial monitoring literature, to obtain minimax rates (Lattimore and Szepesvári, 2019c) and to define the sampling distribution of an algorithm for finite (adversarial) partial monitoring (Lattimore and Szepesvári, 2019b). Bandits (Lattimore and Szepesvári, 2018) are perhaps the most prominent special case of partial monitoring. We discuss additional relevant work in the context of applications in Section 4.

Notation We write $\|\cdot\|$ for the standard euclidean norm. For positive semi-definite A let $\|x\|_A^2 = x^\top A x$. Given $\mathcal{X} \subset \mathbb{R}^d$ we let $\text{conv}(\mathcal{X})$, $\partial\mathcal{X}$ and $\text{diam}(\mathcal{X}) = \sup_{x,y \in \mathcal{X}} \|x - y\|$ denote its convex hull, boundary and diameter respectively. The smallest and largest eigenvalues of a matrix A are denoted by $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ respectively. The identity matrix of dimension d is denoted by $\mathbf{1}_d$. For two square matrices A, B , $A \preceq B$ means that $B - A$ is positive semi-definite. For a possibly non-square matrix C , $\|C\|_2 = \sqrt{\lambda_{\max}(C^\top C)}$ is the operator norm. Given an index set \mathcal{Z} and a collection of matrices $(A_z : z \in \mathcal{Z})$, all with the same number of rows, we define $\text{span}(A_z : z \in \mathcal{Z})$ to be the span of the collection of all columns of the matrices $(A_z : z \in \mathcal{Z})$. When \mathcal{X} is Borel measurable we let $\mathcal{P}(\mathcal{X})$ be the space of probability measures on \mathcal{X} with respect to the Borel σ -algebra. The Dirac probability measure at $x \in \mathcal{X}$ is denoted by $\delta_x \in \mathcal{P}(\mathcal{X})$. The optimal action given parameter θ is $x^*(\theta) = \arg \max_{x \in \mathcal{X}} \langle x, \theta \rangle$. Probability measures on subsets of \mathbb{R}^d are always defined over the Borel σ -algebra. Given a probability measure π on \mathcal{X} let $V_\pi = \int_{\mathcal{X}} A_x A_x^\top d\pi(x) \in \mathbb{R}^{d \times d}$. The filtration $\mathcal{F}_t = \sigma(x_1, a_1, \dots, x_t, a_t)$ contains the observed quantities at the end of round t . \tilde{O} is the Landau notation with logarithmic factors suppressed.

Assumptions Throughout, we make technical boundedness assumptions $\|\theta\|_2 \leq 1$, $\|A_x\|_2 \leq 1$ and $\text{diam}(\mathcal{X}) \leq 1$, which implies $\langle x - y, \theta \rangle \leq 1$ for all $x, y \in \mathcal{X}$. The noise vector ϵ_t is conditionally ρ -subgaussian, $\forall \lambda \in \mathbb{R}$, $\mathbb{E}[e^{\lambda \epsilon_t} | \mathcal{F}_{t-1}, x_t] \leq \exp(\rho^2 \lambda^2 / 2)$ understood coordinate wise. For continuous action sets \mathcal{X} , the map $x \mapsto A_x$ is assumed to be continuous.

2. Information Directed Sampling for Linear Partial Monitoring

Information directed sampling (IDS) was introduced by [Russo and Van Roy \(2014\)](#) for the Bayesian bandit setting. IDS samples actions from a distribution that minimizes the ratio of squared expected regret and mutual information. The *information ratio* appears in a sum under the square root in the regret bound and IDS is the policy that (greedily) minimizes this term. [Kirschner and Krause \(2018\)](#) introduced a frequentist analog of the algorithm that replaces the Bayesian expected suboptimality and information gain with frequentist counterparts and exhibits high-probability regret bounds on linear bandits. In the following we generalize the latter approach, which we simply refer to as IDS.

Formally, let $\Delta_t : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ be a *gap estimate* and $I_t : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ an *information gain* that we will define shortly. We linearly extend the functions Δ_t and I_t to probability distributions over \mathcal{X} so that for distributions $\mu \in \mathcal{P}(\mathcal{X})$ we have $\Delta_t(\mu) = \int_{\mathcal{X}} \Delta_t(x) d\mu(x)$ and $I_t(\mu) = \int_{\mathcal{X}} I_t(x) d\mu(x)$. Information directed sampling is the strategy that samples the action x_t at step t from a distribution $\mu_t \in \mathcal{P}(\mathcal{X})$ that minimizes the *information ratio* $\Psi_t(\mu)$:

$$\mu_t = \arg \min_{\mu \in \mathcal{P}(\mathcal{X})} \Psi_t(\mu), \quad \text{where} \quad \Psi_t(\mu) = \frac{\Delta_t(\mu)^2}{I_t(\mu)}.$$

The minimizing distribution is well defined and can always be chosen with a support of two actions. Furthermore, $\mu \mapsto \Psi_t(\mu)$ is convex. These results were previously shown for the bandit setting by [Kirschner and Krause \(2018, Lemma 4,16,17\)](#) and continue to hold in the more general setting.

Let $A_t = A_{x_t}$ be the observation operator for the action chosen in round t . To estimate the gap $\Delta_t(x)$, IDS uses the regularized least squares estimator, which after t rounds is

$$\hat{\theta}_t = \arg \min_{\theta \in \mathbb{R}^d} \sum_{s=1}^t \|A_s^\top \theta - a_s\|^2 + \|\theta\|^2 = V_t^{-1} \sum_{s=1}^t A_s a_s,$$

where $V_t = \sum_{s=1}^t A_s A_s^\top + \mathbf{1}_d$. Define a sequence of confidence sets $(\mathcal{C}_t)_{t=0}^n$ by

$$\mathcal{C}_t = \left\{ \theta' \in \mathbb{R}^d : \|\hat{\theta}_t - \theta'\|_{V_t} \leq \beta_t^{1/2} \right\}, \quad \text{where} \quad \beta_t^{1/2} = \sqrt{\log \det V_t + 2 \log \left(\frac{1}{\delta} \right) + 1}. \quad (3)$$

The concentration bound by [Abbasi-Yadkori et al. \(2011, Theorem 2\)](#) shows that with probability at least $1 - \delta$ it holds that $\theta \in \mathcal{C}_t$ for all t . Our estimate of the suboptimality gap is defined as

$$\Delta_t(x) = \min \left\{ 1, \max_{y \in \mathcal{X}} \langle y - x, \hat{\theta}_{t-1} \rangle + \beta_{t-1}^{1/2} \|x - y\|_{V_{t-1}^{-1}} \right\},$$

which is chosen so that with high probability $\langle x^* - x, \theta \rangle \leq \Delta_t(x)$ for all $x \in \mathcal{X}$ and all rounds t . Note that $\Delta_t(x) \geq 0$ for all $x \in \mathcal{X}$. For the information gain we use

$$I_t(x) = \log \det \left(\mathbf{1}_m + A_x^\top V_{t-1}^{-1} A_x \right). \quad (4)$$

The definition corresponds to the usual Shannon mutual information when using a Gaussian prior on the parameter and a Gaussian likelihood function. For the bandit setting, it was previously demonstrated by [Kirschner and Krause \(2018\)](#) that the choice of I_t can have a large impact on empirical performance. In Appendix B we discuss some alternative choices for both Δ_t and I_t .

Computationally, a naive implementation of IDS for finite games has a per-step complexity of $\mathcal{O}(d|\mathcal{X}|^2)$, which is required to compute the gap estimates. The IDS distribution can be found by iterating over pairs of actions (a solution supported on two actions always exists). With a weaker regret estimate (Appendix B.1), the computational complexity can be reduced to $\mathcal{O}(|\mathcal{X}|)$ in some cases, which matches the computational cost of index based approaches like UCB.

2.1. A General Regret Bound

The regret of any strategy can be bounded in terms of the cumulative sum of the information ratio $\sum_{t=1}^n \Psi_t(\mu_t)$ and the *total information gain*, $\gamma_n = \sum_{t=1}^n I_t(x_t)$. The following result is a generic regret bound that generalizes Theorem 1 of [Kirschner and Krause \(2018\)](#). Note that for deterministic policies the result can be simplified ([Kirschner and Krause, 2018](#), cf. Theorem 2).

Lemma 1 (IDS regret bound) *For $\eta \geq 0$, let $G_\eta = \{t \in [n] : \Delta_t(\mu_t) \leq \eta\}$. There exists a universal constant $C > 1$ such that for any $n \geq 1$ with probability at least $1 - \delta$ the regret of any (possibly) randomized policy $(\mu_t)_{t=1}^n$ is bounded by*

$$R_n \leq C \inf_{\eta \geq 0} \left[n\eta + \sqrt{\sum_{t \in [n] \setminus G_\eta} \Psi_t(\mu_t) (\gamma_n + \log \frac{1}{\delta})} \right] + 4 \log \left(\frac{4n + 4}{\delta} \right).$$

For the proof, note that $\langle x^* - x_t, \theta \rangle \leq \Delta_t(x_t)$ and consider the sum over the expected gap estimates,

$$\sum_{t=1}^n \Delta_t(\mu_t) = \sum_{t \in G_\eta} \Delta_t(\mu_t) + \sum_{t \in [n] \setminus G_\eta} \Delta_t(\mu_t) \leq n\eta + \sqrt{\sum_{t \in [n] \setminus G_\eta} \Psi_t(\mu_t) \sum_{t=1}^n I_t(\mu_t)},$$

where we used the definition of G_η and the Cauchy–Schwarz inequality. A variance-dependent martingale bound such as Freedman’s inequality shows that the regret R_n concentrates on the sum over (conditional) expected regret up to an additive $\mathcal{O}(\log n)$ term and the total expected information gain is bounded by $\sum_{t=1}^n I_t(\mu_t) \leq c(\gamma_n + \log \frac{1}{\delta})$. The complete proof is given in [Appendix A.3](#).

The next lemma is a standard result ([Abbasi-Yadkori et al., 2011](#), cf. Lemma 10) and shows that for fixed dimension, the total information gain depends only logarithmically on the horizon.

Lemma 2 *For the information gain I_t as defined in (4), $\gamma_n = \log \det(V_n) - \log \det(V_0) \leq d \log \left(1 + \frac{nm}{d} \right)$ and $\beta_n = \sqrt{\gamma_n + 2 \log \frac{1}{\delta}} + 1$.*

Importantly, γ_n and β_n have no dependence on the number of actions. Further, if $\theta \in \mathcal{H}$ is contained in a reproducing kernel Hilbert space \mathcal{H} with bounded Hilbert norm $\|\theta\|_{\mathcal{H}} \leq 1$ (corresponding to a Gaussian process), even the dependence on d can be avoided ([Srinivas et al., 2010](#)).

2.2. Regret Bound for Globally Observable Games

We first analyze globally observable games (see Eq. (1)). The condition implies that $\langle x - y, \theta \rangle$ can be estimated from data collected by the algorithm using appropriate actions. The game-dependent constants that appear in the analysis depend on the degree to which the learner can efficiently gain information, which roughly depends on how well the observation operators A_x are aligned with a direction $x - y$ in which we try to improve the accuracy of our estimation. We define the *worst-case alignment* constant as

$$\alpha = \max_{v \in \mathbb{R}^d} \max_{x, y \in \mathcal{X}} \min_{z \in \mathcal{X}} \frac{\langle x - y, v \rangle^2}{\|A_z^\top v\|^2}.$$

Note that for games that are globally observable, α is always bounded, independent of the number of actions ([Lemma 13](#), [Appendix A.6](#)).

Theorem 3 For any game that satisfies the global observability condition (1), there exists a universal constant $C > 0$ such that for any $n \geq 1$ with probability at least $1 - \delta$,

$$R_n \leq Cn^{2/3} (\alpha\beta_n(\gamma_n + \log \frac{1}{\delta}))^{1/3} + 4 \log \left(\frac{4n+4}{\delta} \right).$$

For fixed feature dimension, β_n and γ_n depend only logarithmically on the horizon, and therefore the regret for globally observable games is $R_n \leq \tilde{O}(n^{2/3})$. We show in Appendix G that all globally observable games that are not locally observable have $R_n = \Omega(n^{2/3})$ in the worst case. The first step in the proof of Theorem 3 is to establish the existence of an action for which the information gain is large relative to the regret of the greedy action.

Lemma 4 The greedy action $\hat{x}_t^* = \arg \max_{x \in \mathcal{X}} \langle x, \hat{\theta}_{t-1} \rangle$ satisfies $\Delta_t(\hat{x}_t^*)^2 \leq 2\alpha\beta_{t-1} \max_z I_t(z)$.

Proof Let $w_t = \arg \max_{w \in \{x-y: x, y \in \mathcal{X}\}} \|w\|_{V_{t-1}^{-1}}^2$ be the most uncertain direction. Then,

$$\Delta_t(\hat{x}_t^*) = \max_{y \in \mathcal{X}} (y - \hat{x}_t^*)^\top \hat{\theta}_{t-1} + \beta_{t-1}^{1/2} \|y - \hat{x}_t^*\|_{V_{t-1}^{-1}} \leq \beta_{t-1}^{1/2} \|w_t\|_{V_{t-1}^{-1}}.$$

Note three ways to write $\|w\|_{V_{t-1}^{-1}}^2 = \|V_{t-1}^{-1/2} w\|^2 = \langle w_t, V_{t-1}^{-1} w_t \rangle$. Basic linear algebra shows that

$$\begin{aligned} \frac{\|A_z^\top V_{t-1}^{-1} w_t\|^2}{\|V_{t-1}^{-1/2} w_t\|^2} &\leq \max_{v \in \mathbb{R}^d} \frac{\|A_z^\top V_{t-1}^{-1/2} v\|^2}{\|v\|^2} = \lambda_{\max}(A_z^\top V_{t-1}^{-1} A_z) \\ &\leq 2 \log \det(\mathbf{1}_d + A_z^\top V_{t-1}^{-1} A_z) = 2I_t(z). \end{aligned}$$

For the last step we used the inequality $a \leq 2 \log(1 + a)$ for $a \in [0, 1]$ and that the eigenvalues of $A_z^\top V_{t-1}^{-1} A_z$ are bounded in $[0, 1]$ by the assumption $\|A_z\| \leq 1$ and $V_{t-1}^{-1} \preceq \mathbf{1}_d$. With the most informative action $z_t = \arg \max_{x \in \mathcal{X}} I_t(x)$, it follows that

$$\frac{\Delta_t(\hat{x}_t^*)^2}{I_t(z_t)} \leq 2\beta_{t-1} \min_z \frac{\langle w_t, V_{t-1}^{-1} w_t \rangle^2}{\|A_z^\top V_{t-1}^{-1} w_t\|^2} \leq 2\beta_{t-1} \max_{v \in \mathbb{R}^d} \max_{x, y \in \mathcal{X}} \min_z \frac{\langle x - y, v \rangle^2}{\|A_z^\top v\|^2} = 2\alpha\beta_{t-1}.$$

Rearranging completes the proof. ■

The following lemma shows that IDS never plays a distribution that is too far from greedy. The proof is deferred to Appendix A.4.

Lemma 5 Let μ_t be the IDS distribution at time t . Then $\Delta_t(\mu_t) \leq 2 \min_{x \in \mathcal{X}} \Delta_t(x)$.

Proof of Theorem 3 Let $z_t = \arg \max_{x \in \mathcal{X}} I_t(x)$ be the informative action. For $p \in [0, 1]$, let $\mu(p) = (1-p)\delta_{\hat{x}_t^*} + p\delta_{z_t}$ be the distribution that randomizes between the greedy and the informative action. By definition, the information ratio of IDS is bounded by the ratio of $\mu(p)$,

$$\Psi_t(\mu_t) \leq \min_{p \in [0, 1]} \frac{\Delta_t(\mu(p))^2}{I_t(\mu(p))} \leq 2\alpha\beta_{t-1} \min_{p \in [0, 1]} \frac{((1-p)\Delta_t(\hat{x}_t^*) + p)^2}{p\Delta_t(\hat{x}_t^*)^2} \leq \frac{8\alpha\beta_{t-1}}{\Delta_t(\hat{x}_t^*)} \leq \frac{16\alpha\beta_{t-1}}{\Delta_t(\mu_t)}.$$

The second inequality uses $\Delta_t(z) \leq 1$, $I_t(\hat{x}_t^*) \geq 0$ and Lemma 4 to bound $\Delta_t(\hat{x}_t^*)^2 \leq 2\alpha\beta_{t-1}I_t(z)$. The third inequality follows by choosing $p = \Delta_t(\hat{x}_t^*) \in [0, 1]$ and the last follows from Lemma 5. Next, Lemma 1 shows that with probability at least $1 - \delta$,

$$R_n \leq C \inf_{\eta \geq 0} \left[n\eta + \sqrt{\frac{16n\alpha\beta_n(\gamma_n + \log \frac{1}{\delta})}{\eta}} \right] + 4 \log \left(\frac{4n+4}{\delta} \right),$$

where we used the fact that $(\beta_t)_{t=0}^n$ is non-decreasing. Optimizing η completes the proof. \blacksquare

2.3. Regret Bound for Locally Observable Games

In globally observable games, the learner can estimate the gaps for all actions, but may need to play actions that are known to be suboptimal. The definition of local observability (see Eq. (2)) means that the learner can gain information while playing only actions that appear plausibly optimal.

Recall the definition of the confidence set \mathcal{C}_t in Eq. (3) and let $\mathcal{P}_t = \mathcal{P}(\mathcal{C}_{t-1})$ be the set of actions that are plausibly optimal in round t . Again, our bound depends on the signal to noise ratio when exploring. For a set of (plausible optimal) actions $\mathcal{Y} \subset \mathcal{X}$, define the *worst-case alignment* for \mathcal{Y} ,

$$\alpha(\mathcal{Y}) = \max_{v \in \mathbb{R}^d} \max_{x, y \in \mathcal{Y}} \min_{z \in \mathcal{Y}} \frac{\langle x - y, v \rangle^2}{\|A_z^\top v\|^2}. \quad (5)$$

Globally observable games satisfy $\alpha = \alpha(\mathcal{X}) < \infty$. The local observability condition implies that this remains true if we restrict actions to $\mathcal{P}(\mathcal{C})$. All games with bandit feedback ($A_x = x$) satisfy $\alpha(\mathcal{Y}) \leq 4$. We refer to Lemma 13 in Appendix A.6 for details. We say a game is *uniformly locally observable* if $\alpha(\mathcal{P}(\mathcal{C})) \leq \alpha_0$ for all $\mathcal{C} \subset \mathbb{R}^d$ convex. All finite locally observable games are uniformly locally observable because there are only finitely many subsets. The definition of the alignment constant can be tightened with a more careful analysis, to obtain improved bounds on model parameters such as the dimension in some cases. We refer to Appendix C.5 for details.

Theorem 6 *For locally observable games denote $\alpha_t = \alpha(\mathcal{P}_t)$. There exists a universal constant $C > 0$ such that for any $n \geq 1$ with probability at least $1 - \delta$,*

$$R_n \leq C \sqrt{\sum_{t=1}^n \alpha_t \beta_t (\gamma_n + \log \frac{1}{\delta})} + 4 \log \left(\frac{4n+4}{\delta} \right).$$

For games that are uniformly locally observable, the regret bound is $R_n \leq \tilde{O}(\sqrt{\alpha_0 \beta_n \gamma_n n})$.

We show in Appendix G that on locally observable games with more than one Pareto optimal action, any algorithm suffers $\Omega(n^{1/2})$ regret in the worst case. To prove the upper bound, the first step is to construct an exploration distribution that is supported on the plausible maximizers \mathcal{P}_t and has a constant information ratio. Note that IDS is not restricted to playing actions within \mathcal{P}_t , nor is it required to explicitly compute this set. In fact, actions that are not plausible maximizers can have a better trade-off between regret and information.

Lemma 7 *For locally observable games, there exists an exploration action in $z \in \mathcal{P}_t$ such that,*

$$\forall x \in \mathcal{P}_t, \quad \Delta_t(x)^2 \leq 8\alpha_t \beta_{t-1} I_t(z).$$

The complete proof is in Appendix A.5. The argument shows that for plausible maximizers $x \in \mathcal{P}_t$, $\Delta_t(x)^2 \leq 4\beta_{t-1} \max_{y \in \mathcal{P}_t} \|y - x\|_{V_{t-1}^{-1}}^2$ and is otherwise similar to the proof of Lemma 4.

Proof of Theorem 6 Let $z_t = \arg \max_{x \in \mathcal{P}_t} I_t(x)$ be the most informative action in the current plausible maximizer set \mathcal{P}_t . By Lemma 7,

$$\Psi_t(\mu_t) \leq \frac{\Delta_t(z_t)^2}{I_t(z_t)} \leq 8\alpha_t\beta_{t-1}.$$

Invoking the general IDS bound (Lemma 1) with $\eta = 0$ completes the proof. \blacksquare

The proof shows that randomization is not necessary to achieve a bounded information ratio in locally observable games. Deterministic IDS (Kirschner and Krause, 2018), which optimizes the ratio over a deterministic action choice $x_t = \arg \min_{x \in \mathcal{X}} \Psi_t(\delta_x)$, achieves the same upper bound with our analysis. Moreover, the bound shows how IDS adapts towards the current instance of the partial monitoring game. Consider a globally observable game where after some finite time n_0 , the plausible maximizer sets \mathcal{P}_t are locally observable in sense that $\alpha(\mathcal{P}_t) \leq \alpha_0, \forall t \geq n_0$. In this case the regret bound is $R_n \leq \tilde{O}(\alpha^{1/3}n_0^{3/2} + (\alpha_0n)^{1/2})$. We have not yet identified non-artificial conditions that ensure this behavior, however. The gold standard would be to prove finite-time, instance-dependent regret bounds with small constants. At present such results are more or less restricted to finite-armed bandits, however, and remain open even for linear bandits (Hao et al., 2019).

2.4. Smooth Convex Action Sets

The observability conditions are more ambiguous when $\text{conv}(\mathcal{X})$ is not a polytope. Here we prove that when \mathcal{X} has strictly positive principle curvature, then IDS enjoys $\tilde{O}(\sqrt{n})$ regret on globally observable games. Curvature of the action set has been exploited in online learning (Huang et al., 2017) and bandits (Bubeck et al., 2018). The latter article considers the starved adversarial linear bandit, where the learner only observes the rewards when sampling an action from a pre-specified distribution. They consider the case where the action set is the unit ball with respect to $\|\cdot\|_p$ and prove that for $p = 2$ one can obtain $O(n^{1/2})$ regret, but not for $p > 2$. This setting is close to a special case of linear partial monitoring (see Appendix C.3). Let $h_{\mathcal{X}} : \mathbb{R}^d \rightarrow \mathbb{R}$ be the support function of \mathcal{X} , which is defined by $h_{\mathcal{X}}(u) = \sup_{x \in \mathcal{X}} \langle x, u \rangle$.

Theorem 8 *Assume that \mathcal{X} is closed, convex, has a non-empty interior and that $h_{\mathcal{X}}$ is twice differentiable. Suppose furthermore that the game is globally observable according to Eq. (1) and has strictly positive principle curvature everywhere:*

$$\kappa_o = \left(\max_{\eta \in \mathbb{R}^d: \|\eta\|_2=1} \lambda_{\max}(\nabla^2 h_{\mathcal{X}}(\eta)) \right)^{-1} > 0.$$

Then, with probability at least $1 - \delta$, for any $n \geq 1$,

$$R_n \leq C \sqrt{\max \left(1, \frac{1}{\kappa_o} \right) \beta_n \left(\gamma_n + \log \frac{1}{\delta} \right) n + 4 \log \left(\frac{4n+4}{\delta} \right)},$$

where C is a constant depending only on $(A_z : z \in \mathcal{X})$.

The proof is given in Appendix D. The key argument shows that Δ_t applied to the empirically optimal action scales like the square of the diameter of the confidence set. This compares favorably with the case without curvature, where the error is about linear in the diameter of the confidence set.

2.5. Contextual Partial Monitoring Games

The contextual bandit problem is a well known extension of the bandit setting where the learner receives a context before choosing the action (Woodroffe, 1979; Langford and Zhang, 2008). We introduce a novel contextual variant of linear partial monitoring, that strictly generalizes the linear contextual bandit setting. Let \mathcal{Z} be a compact set of contexts. Each context $z \in \mathcal{Z}$ defines a partial monitoring game with action set \mathcal{X}_z and the observation operators $\{A_x^z\}$, where the map $(x, z) \mapsto A_x^z$ is assumed to be continuous. At time t , the learner receives a context $z_t \in \mathcal{Z}$ and chooses an action $x_t \in \mathcal{X}_{z_t}$. The reward is $\langle x_t, \theta \rangle$ and the observation is $a_t = A_{x_t}^{z_t \top} \theta + \epsilon_t$ where the parameter θ is the same in every context. The objective is to compete with the best in-hindsight policy that maps context to actions. Regret is defined with respect to the context-dependent solution $x_t^* = \arg \max_{x \in \mathcal{X}_{z_t}} \langle x, \theta \rangle$:

$$R_n = \sum_{t=1}^n \langle x_t^* - x_t, \theta \rangle.$$

The regret of the learner depends on the sequence of contexts observed and the corresponding sequence of partial monitoring games which share the common parameter θ . All our notions extend with the contextual argument,

$$\Delta_t(x, z) = \max_{y \in \mathcal{X}_z} \langle y - x, \hat{\theta}_t \rangle + \beta_{t-1}^{1/2} \|y - x\|_{V_{t-1}^{-1}}, \quad I_t(x, z) = \log \det(\mathbf{1} + A_x^z \top V_{t-1}^{-1} A_x^z).$$

Conditional IDS is the policy that minimizes the $\mu_t = \arg \min_{\mu \in \mathcal{P}(\mathcal{X})} \Psi(\mu, z_t)$ conditioned on the observed context. The next result extends the regret guarantees for locally and globally observable games to the contextual setting by making strong assumptions on the sequence of games defined by the context. We refer to Appendix E.1 for our formal result.

Corollary 9 (Informal) *If the sequence of games defined by the observed contexts z_1, \dots, z_n are globally observable, conditional IDS achieves $R_n \leq \mathcal{O}(n^{2/3}(\alpha\beta_n(\gamma_n + \log \frac{1}{\delta}))^{1/3})$ regret with high probability. If the sequence of games is uniformly locally observable, then conditional IDS achieves $R_n \leq \mathcal{O}((\alpha_0\beta_n(\gamma_n + \log \frac{1}{\delta})n)^{1/2})$.*

Perhaps surprisingly, the contextual case allows for much weaker conditions under which no-regret is possible if the learner exploits the distribution of contexts. Here we study the case where the context follows a known distribution $\nu \in \mathcal{P}(\mathcal{Z})$; the case where the distribution is unknown or the learner tries to adapt her behaviour towards an arbitrary sequence is left as an interesting direction for future work. It is instructive to think about some examples:

- An extreme case is where for some $z \in \mathcal{Z}$ the learner obtains no information ($A_x^z = 0$ for all $x \in \mathcal{X}_z$). In such rounds the only sensible choice is the greedy action. Exploration needs to happen in rounds where information is available and needs to be sufficiently diverse to account for rounds where the learner is forced to play greedily. Note that while there can be vanishing information gain in *some* rounds, the *expected* information gain, that takes the distribution ν over the context into account, is non-zero.
- Since also the greedy action depends on the random context, there can be cases where the learner incurs sufficient exploration by playing mostly greedy. This effect has been studied in the bandit literature before (Bastani et al., 2017; Hao et al., 2019).

Conditional IDS does not depend on the distribution ν and it is easy to see that it can behave suboptimally in both examples. To include the randomness of the context within the IDS framework, consider a joint distribution $\xi \in \mathcal{P}(\mathcal{X} \times \mathcal{Z})$ over context and actions with marginal $\xi_z \in \mathcal{P}(\mathcal{Z})$. As before, $\Delta_t(\xi)$ and $I_t(\xi)$ extend linearly. At time t , *contextual IDS* computes a distribution ξ with marginal $\xi_z = \nu$, that minimizes the joint ratio,

$$\xi_t = \arg \min_{\xi \in \mathcal{P}(\mathcal{X} \times \mathcal{Z}), \xi_z = \nu} \Psi(\xi).$$

The action is sampled from $x_t \sim \xi_t(x|z_t)$ after observing z_t . In the joint minimization of the information ratio the contextual distribution ν contributes to exploration and a smaller information ratio. The intuition is that to estimate along a direction $x - y$ in a contextual action set $x, y \in \mathcal{X}_z$, the learner can wait for a different context z' to be realized where $x - y$ can easily be estimated and at low cost. This leads to the following condition that defines *globally observable contextual games*:

$$\forall z \in \mathcal{Z} \text{ and } x, y \in \mathcal{X}_z \Rightarrow \exists z' \in \mathcal{Z} \text{ s.t. } x - y \in \text{span}(A_x : x \in \mathcal{X}_{z'}). \quad (6)$$

The regret bound depends on the probability that a context occurs where estimation is possible. The *expected worst-case alignment* $\alpha(\nu)$ is defined in Appendix E.2, Eq. (14). It satisfies the intuitive upper bound $\alpha(\nu) \leq \mathbb{E}_\nu[\alpha(z)]$ and recovers the previous definition for Dirac delta distributions (Lemma 21, Appendix E.2). For finite games with finite context set, it further holds that

$$\alpha(\nu) \leq \max_{v \in \mathbb{R}^d} \max_{z \in \mathcal{Z}} \max_{x, y \in \mathcal{X}_z} \min_{z' \in \mathcal{Z}} \min_{u \in \mathcal{X}_{z'}} \frac{\langle v, x - y \rangle^2}{\nu(z') \|A_u^{z'} v\|^2},$$

thus it suffices that a direction $x - y$ in \mathcal{X}_z can be estimated under *some* context $z' \in \mathcal{Z}$ that appears with non-zero probability $\nu(z')$. The next result quantifies the rate in globally observable games.

Theorem 10 *For globally observable contextual games with bounded expected worst-case alignment $\alpha(\nu)$, for any $n \geq 1$, the regret is bounded with probability at least $1 - \delta$,*

$$R_n \leq Cn^{2/3} \left(\alpha(\nu) \beta_n (\gamma_n + \log \frac{1}{\delta}) \right)^{1/3} + 4 \log \left(\frac{4n + 4}{\delta} \right).$$

All proofs and details for this result can be found in Appendix E.2 and the analogous result for the locally observable case is in Appendix E.3.

3. Classification of Finite Games

The upper bounds show that for globally observable games the regret is $\tilde{O}(n^{2/3})$, while for locally observable games it is $\tilde{O}(n^{1/2})$. Of course, if there is only one Pareto optimal action, then the regret vanishes for any algorithm that just plays this action. The classification theorem follows by proving that for games that are not globally observable, the regret is linear in the worst case, that for globally observable games that are not locally observable the regret is $\Omega(n^{2/3})$ and that for locally observable games with more than one Pareto optimal action it is $\Omega(n^{1/2})$. These lower bounds are supplied in Appendix G. For simplicity, our results are for the *expected minimax regret*, which is

$$R_n^* = \inf_{\pi} \sup_{\theta} \mathbb{E}[R_n(\pi, \theta)].$$

The infimum is over policies $\pi = (\pi_t)_{t=1}^n$ defined by a sequence of \mathcal{F}_t -measurable random variables on \mathcal{X} and $R_n(\pi, \theta) = \sum_{t=1}^n \langle x^* - x_t, \theta \rangle$ is the regret for parameter $\theta \in \mathbb{R}^d$ when the actions are sampled from the policy π .

Theorem 11 *The minimax regret for any finite linear partial monitoring game satisfies*

$$R_n^* = \begin{cases} 0 & \text{if there is only one Pareto optimal action,} \\ \tilde{\Theta}(n^{1/2}) & \text{for locally observable games,} \\ \tilde{\Theta}(n^{2/3}) & \text{for globally observable games,} \\ \Omega(n) & \text{otherwise.} \end{cases}$$

The classification theorem is proven by combining upper and lower bounds, carefully checking that all cases have been covered. We further show in Appendix F that our definitions of local and global observability coincide with the standard notions in finite partial monitoring that are based on the neighborhood graph, as well as the notion of a *global observer set* used by Lin et al. (2014).

4. Applications and Extensions

The framework of linear partial monitoring captures many applications and models for sequential decision making that were previously studied in the literature. We outline some of them below and provide additional details in Appendix C.

Semi-Bandits and Full Information Feedback The observation operators can be defined to yield *more* information than in the bandit case, up to revealing the parameter in each round ($A_x = \mathbf{1}_d$). Naturally, additional information should only improve performance, but in our analysis, the total information gain degrades logarithmically with the observation dimension m . For the case of full information feedback, we show in Appendix C.1 how to improve the bounds to get $R_n \leq \tilde{O}(\sqrt{dn})$. Bandits with graph feedback are another type of semi-bandits, where the learner observes the reward of neighbouring nodes in a graph (Alon et al., 2017). In fact, Bayesian IDS has been applied in this setting before (Liu et al., 2018).

Linear Bandits The *stochastic linear bandit* setting is a special case of our setup with $A_x = x$ (Abe and Long, 1999; Auer, 2003; Dani et al., 2008; Abbasi-Yadkori et al., 2011). Our analysis achieves the optimal $R_n \leq \tilde{O}(d\sqrt{n})$ dependency for the regret and generalizes the results for heteroscedastic bandits by Kirschner and Krause (2018). The UCB algorithm (Auer et al., 2002) has a distinct relation to the IDS framework, as we explain in Appendix B.3.

Dueling Bandits In *dueling bandits*, the learner chooses a pair of actions and receives binary feedback indicating which action has higher reward (Yue and Joachims, 2009). This feedback model can be cast as partial monitoring game (Gajane and Urvoy, 2015). Let $\mathcal{X}_0 \subset \mathbb{R}^d$ be a ground set and $\mathcal{X} = \mathcal{X}_0 \times \mathcal{X}_0$. The relative feedback for $(x_1, x_2) \in \mathcal{X}$ is defined with $A_{x_1, x_2} = x_1 - x_2$ and noise is added such that $a_t \in \{-1, 1\}$ is binary with expectation $A_{x_1, x_2}^\top \theta$. Note that bounded noise is subgaussian and our analysis applies. A possible reward model is to use averaged features $x_{x_1, x_2} = (x_1 + x_2)/2$. Dueling bandits are locally observable if the learner can compare any pair of actions, and globally observable if comparisons are restricted to ‘adjacent’ actions. See Appendix C.5 for details.

Combinatorial Bandits This is the original motivation for the linear partial monitoring setting by [Lin et al. \(2014\)](#) and [Chaudhuri and Tewari \(2016\)](#) and leads to games that are either locally or globally observable. We refer to the previous works for further applications. Our formulation covers combinatorial bandits both with bandit and semi-bandit feedback. An important special case is the batch setting ([Appendix C.4](#)).

Transductive and Starved Bandits The *transductive linear bandit* setting was recently proposed by [Fiez et al. \(2019\)](#). The learner has access to a set of actions that is dedicated for exploration, while the objective is to achieve low regret on a different, target set of actions. It was open to find an approach that minimizes cumulative regret, which we effectively resolve ([Appendix C.2](#)). Similar in spirit are *starved bandits* ([Bubeck et al., 2018](#)), where the learner only obtains information when sampling actions from a pre-defined distribution. This setting is closely connected to our contextual setting (see [Appendix C.3](#)) and the regret bounds on convex action sets in [Section 2.4](#).

Product Testing and Invasive Measurements An early toy example for a globally, but not locally observable game is that of “*apple tasting*” ([Cesa-Bianchi et al., 2006](#)). In this task, the learner optimizes a production chain with the option to remove a product for inspection (and destroying it in the process). Other applications include parameter tuning of experimental facilities such as particle accelerators ([Kirschner et al., 2019](#)), where invasive measurement devices provide a very rich signal at the expense of voiding any downstream measurements (for a stylized version of this problem and a numerical demonstration of IDS, see [Appendix C.6](#)).

Kernelized Partial Monitoring Our approach and the analysis extend to the kernelized setting, where the reward function is in a known reproducing kernel Hilbert space (RKHS). This includes kernelized bandits ([Srinivas et al., 2010](#); [Abbasi-Yadkori, 2012](#); [Chowdhury and Gopalan, 2017](#)), also known as *Bayesian optimization*, as a special case. Interesting applications beyond the bandit setting include Bayesian optimization with gradients ([Wu et al., 2017b](#)) or even Hessian evaluations ([Wu et al., 2017a](#)). Unlike previous results, our approach leverages all available information and achieves a strong finite time convergence guarantee. We refer the reader to [Appendix C.6](#) for a detailed introduction and formal statements. In the limit with continuous action sets, dueling bandits can be understood as *global optimization* where the learner has access *only* to the gradient.

5. Discussion

We introduced *information directed sampling* for stochastic *linear partial monitoring*, which – to the best of our knowledge – is the first approach that achieves the optimal regret rate in all finite linear games. Our classification theorem provides a complete picture of the achievable worst-case regret rates in finite linear games. Nevertheless, many directions are left for future work. Proving non-trivial instance-dependent regret bounds for IDS is an important open question, even for the standard linear bandit setting. Another challenge is to find precise observability conditions that capture the rate achievable on continuous action sets. On the computational side, some previous approaches rely on oracle solvers ([Lin et al., 2014](#); [Chaudhuri and Tewari, 2016](#)) and for the bandit setting, Thompson sampling is a well-known oracle efficient method ([Agrawal and Goyal, 2013](#); [Abeille and Lazaric, 2017](#)). Given the generality of our results, finding an oracle-efficient approximation of IDS is an important task for future work.

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme grant agreement No 815943.

References

- Yasin Abbasi-Yadkori. *Online Learning for Linearly Parametrized Control Problems*. PhD thesis, 2012.
- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- Naoki Abe and Philip M. Long. Associative reinforcement learning using linear probabilistic concepts. In *Proceedings of the Sixteenth International Conference on Machine Learning, ICML ’99*, pages 3–11, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc. ISBN 1-55860-612-2.
- Marc Abeille and Alessandro Lazaric. Linear thompson sampling revisited. In *Artificial Intelligence and Statistics*, pages 176–184, 2017.
- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.
- Noga Alon, Nicolo Cesa-Bianchi, Claudio Gentile, Shie Mannor, Yishay Mansour, and Ohad Shamir. Nonstochastic multi-armed bandits with graph-structured feedback. *SIAM Journal on Computing*, 46(6):1785–1826, 2017.
- A. Antos, G. Bartók, D. Pál, and Cs. Szepesvári. Toward a classification of finite partial-monitoring games. *Theoretical Computer Science*, 473:77–99, 2013.
- P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *J. Mach. Learn. Res.*, 3:397–422, March 2003. ISSN 1532-4435.
- G. Bartók, D. Pál, and Cs. Szepesvári. Minimax regret of finite partial-monitoring games in stochastic environments. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 133–154, 2011.
- G. Bartók, N. Zolghadr, and Cs. Szepesvári. An adaptive algorithm for finite stochastic partial monitoring. In *Proceedings of the 29th International Conference on International Conference on Machine Learning, ICML, pages 1779–1786, USA, 2012. Omnipress.*
- G. Bartók, D. P. Foster, D. Pál, A. Rakhlin, and Cs. Szepesvári. Partial monitoring—classification, regret bounds, and algorithms. *Mathematics of Operations Research*, 39(4):967–997, 2014.
- Gábor Bartók and Csaba Szepesvári. Partial monitoring with side information. In *International Conference on Algorithmic Learning Theory*, pages 305–319. Springer, 2012.

- Hamsa Bastani, Mohsen Bayati, and Khashayar Khosravi. Mostly exploration-free algorithms for contextual bandits. *arXiv preprint arXiv:1704.09011*, 2017.
- J Bretagnolle and C Huber. Estimation des densités: risque minimax. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 47(2):119–137, 1979.
- S. Bubeck, M. Cohen, and Y. Li. Sparsity, variance and curvature in multi-armed bandits. In F. Janoos, M. Mohri, and K. Sridharan, editors, *Proceedings of Algorithmic Learning Theory*, volume 83 of *Proceedings of Machine Learning Research*, pages 111–127. PMLR, 07–09 Apr 2018.
- N. Cesa-Bianchi, G. Lugosi, and G. Stoltz. Regret minimization under partial monitoring. *Mathematics of Operations Research*, 31:562–580, 2006.
- Nicolo Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- Sougata Chaudhuri and Ambuj Tewari. Phased exploration with greedy exploitation in stochastic combinatorial partial monitoring games. In *Advances in Neural Information Processing Systems*, pages 2433–2441, 2016.
- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, 2017.
- Varsha Dani, Thomas P. Hayes, and Sham M. Kakade. Stochastic Linear Optimization under Bandit Feedback. In *COLT*, pages 355–366. Omnipress, 2008.
- Tanner Fiez, Lalit Jain, Kevin G Jamieson, and Lillian Ratliff. Sequential experimental design for transductive linear bandits. In *Advances in Neural Information Processing Systems 32*, pages 10666–10676. Curran Associates, Inc., 2019.
- Pratik Gajane and Tanguy Urvoy. Utility-based dueling bandits as a partial monitoring game. *arXiv preprint arXiv:1507.02750*, 2015.
- S. Gerchinovitz and T. Lattimore. Refined lower bounds for adversarial bandits. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, NIPS, pages 1198–1206. Curran Associates, Inc., 2016.
- Javier González, Zhenwen Dai, Andreas Damianou, and Neil D Lawrence. Preferential bayesian optimization. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1282–1291. JMLR. org, 2017.
- Botao Hao, Tor Lattimore, and Csaba Szepesvari. Adaptive exploration in linear contextual bandit. *arXiv preprint arXiv:1910.06996*, 2019.
- R. Huang, T. Lattimore, A. György, and Cs. Szepesvári. Following the leader and fast rates in online linear prediction: Curved constraint sets and other regularities. *Journal of Machine Learning Research*, 18:1–31, 2017.

- Motonobu Kanagawa, Philipp Hennig, Dino Sejdinovic, and Bharath K Sriperumbudur. Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv preprint arXiv:1807.02582*, 2018.
- Johannes Kirschner and Andreas Krause. Information directed sampling and bandits with heteroscedastic noise. *arXiv preprint arXiv:1801.09667*, 2018.
- Johannes Kirschner, Mojmir Mutny, Nicole Hiller, Rasmus Ischebeck, and Andreas Krause. Adaptive and safe bayesian optimization in high dimensions via one-dimensional subspaces. In *International Conference on Machine Learning*, pages 3429–3438, 2019.
- J. Komiyama, J. Honda, and H. Nakagawa. Regret lower bound and optimal algorithm in finite stochastic partial monitoring. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, NIPS, pages 1792–1800. Curran Associates, Inc., 2015.
- J. Langford and T. Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, NIPS, pages 817–824. Curran Associates, Inc., 2008.
- Tor Lattimore and Csaba Szepesvári. Bandit algorithms. 2018.
- Tor Lattimore and Csaba Szepesvári. Cleaning up the neighborhood: A full classification for adversarial partial monitoring. In *Algorithmic Learning Theory*, pages 529–556, 2019a.
- Tor Lattimore and Csaba Szepesvári. Exploration by optimisation in partial monitoring. *arXiv preprint arXiv:1907.05772*, 2019b.
- Tor Lattimore and Csaba Szepesvári. An information-theoretic approach to minimax regret in partial monitoring. *arXiv preprint arXiv:1902.00470*, 2019c.
- Tian Lin, Bruno Abrahao, Robert Kleinberg, John Lui, and Wei Chen. Combinatorial partial monitoring game with linear feedback and its applications. In *International Conference on Machine Learning*, pages 901–909, 2014.
- Fang Liu, Swapna Buccapatnam, and Ness Shroff. Information directed sampling for stochastic bandits with graph feedback. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. In *Advances in Neural Information Processing Systems*, pages 1583–1591, 2014.
- A. Rustichini. Minimizing regret: The general case. *Games and Economic Behavior*, 29(1):224–243, 1999.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *International Conference on Machine Learning*, 2010.
- Yanan Sui, Yisong Yue, and Joel W Burdick. Correlational dueling bandits with application to clinical treatment in large decision spaces. *arXiv preprint arXiv:1707.02375*, 2017a.

- Yanan Sui, Vincent Zhuang, Joel W Burdick, and Yisong Yue. Multi-dueling bandits with dependent arms. *arXiv preprint arXiv:1705.00253*, 2017b.
- Yanan Sui, Masrour Zoghi, Katja Hofmann, and Yisong Yue. Advancements in dueling bandits. In *IJCAI*, pages 5502–5510, 2018.
- Hastagiri Vanchinathan, Gábor Bartók, and Andreas Krause. Efficient partial monitoring with prior information. In *Neural Information Processing Systems (NIPS)*, 2014.
- Zi Wang, Bolei Zhou, and Stefanie Jegelka. Optimization as estimation with gaussian processes in bandit settings. In *Artificial Intelligence and Statistics*, pages 1022–1031, 2016.
- Michael Woodroffe. A one-armed bandit problem with a concomitant variable. *Journal of the American Statistical Association*, 74(368):799–806, 1979.
- Anqi Wu, Mikio C Aoi, and Jonathan W Pillow. Exploiting gradients and Hessians in Bayesian optimization and Bayesian quadrature. *arXiv preprint arXiv:1704.00060*, 2017a.
- Jian Wu, Matthias Poloczek, Andrew G Wilson, and Peter Frazier. Bayesian optimization with gradients. In *Advances in Neural Information Processing Systems*, pages 5267–5278, 2017b.
- Yisong Yue and Thorsten Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1201–1208, 2009.

Appendix A. Additional Lemmas and Proofs

A.1. Linear Partial Monitoring: General Setup

Our setting can be formulated more generally, to allow applications where the learner can choose between different observation maps that are associated to the same action. Let \mathcal{I} be a compact index set. Each $i \in \mathcal{I}$ indexes an action-observation tuple (x_i, A_i) and the collection of such tuples represents a game $\mathcal{G} = \{(x_i, A_i) : i \in \mathcal{I}\} \subset \mathbb{R}^d \times \mathbb{R}^{d \times m}$. At step t , the learner chooses an action index i_t and observes the outcome $A_{i_t}^\top \theta + \epsilon_t$. The unobserved reward is $\langle x_{i_t}, \theta \rangle$. We assume that the map $i \mapsto (x_i, A_i)$ is continuous to guarantees that the IDS distribution exists. The dimension m of the observation can also depend on the action i in general, but for simplicity, we set $m_i = m$. We overload

A.2. Finite Partial Monitoring

Unlike the standard finite and linear bandit frameworks, finite partial monitoring is not quite a special case of the linear setting. On the one hand, our setting permits infinite observation (and action) spaces, which are not usually covered by existing results. On the other hand, the assumptions of our setting mean the algorithm does not recover known bounds for algorithms in the finite unstructured setting. The main reason is that we do not restrict θ except in terms of $\|\theta\| \leq 1$, while in the finite setting the θ is effectively constrained to the probability simplex. Consider the following finite game, characterized by reward and signal matrices

$$\mathcal{R} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}, \quad \Sigma = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

The signal matrix is such that the learner observes no information. Meanwhile, however, the rewards are such that the learner knows immediately that the first action is optimal, so in the finite partial monitoring literature this game is trivial and good algorithms suffer zero regret. Our algorithm, however, does not assume that θ lies in the probability simplex, and when $\theta = (-\sqrt{2}, -\sqrt{2})$, the second action is clearly optimal. The different assumptions on θ mean that this game is now hopeless and algorithms consequentially suffer linear regret.

A.3. Proof of Lemma 1

Proof Using Freedman's inequality one can get the following concentration result on the regret (Kirschner and Krause, 2018, Lemma 13). For any fixed n , with probability at least $1 - \delta/2$,

$$R_n \leq \frac{5}{4} \sum_{t=1}^n \Delta_t(\mu_t) + 4 \log \left(\frac{4n+4}{\delta} \right).$$

The first sum is bounded by

$$\sum_{t=1}^n \Delta_t(\mu_t) = \sum_{t \in G_\eta} \Delta_t(\mu_t) + \sum_{t \in [n] \setminus G_\eta} \Delta_t(\mu_t) \leq n\eta + \sqrt{\sum_{t \in [n] \setminus G_\eta} \Psi_t(\mu_t) \sum_{t=1}^n I_t(\mu_t)}.$$

The inequality follows from the definition of G_η and we use Cauchy-Schwarz to bound

$$\sum_{t \in [n] \setminus G_\eta} \Delta_t(\mu_t) \leq \sum_{t \in [n] \setminus G_\eta} \sqrt{\Psi_t(\mu_t) I_t(\mu_t)} \leq \sqrt{\sum_{t \in [n] \setminus G_\eta} \Psi_t(\mu_t) \sum_{t \in [n]} I_t(\mu_t)}.$$

In the last step, we also used the non-negativity of $I_t(\mu_t)$. Finally, the sum over expected information gain $\sum_{t=1}^n I_t(\mu_t)$ is close to the realized information gain $\sum_{t=1}^n I_t(x_t)$ with high probability. This is made precise in Lemma 3 of [Kirschner and Krause \(2018\)](#), which shows that if $I_t(x) \leq 1$, then with probability at least $1 - \delta/2$, for any $n \geq 1$,

$$\sum_{t=1}^n I_t(\mu_t) \leq 2 \sum_{t=1}^n I_t(x_t) + 4 \log \left(\frac{1}{\delta} \right) + 8$$

Note that our boundedness assumptions $\|A_x\|_2 \leq 1$ and the fact that $V_t^{-1} \preceq \mathbf{1}_d$ imply the required assumption $I_t(x) \leq 1$. By definition $\gamma_n = \sum_{t=1}^n I_t(x_t)$. A union bound over the previous displays completes the proof. \blacksquare

A.4. Proof of Lemma 5

Proof By assumption, for any $p \in [0, 1)$ and any $x \in \mathcal{X}$,

$$\Psi_t(\mu_t) = \frac{\Delta_t(\mu_t)^2}{I_t(\mu_t)} \leq \frac{\Delta_t((1-p)\mu_t + p\delta_x)^2}{(1-p)I_t(\mu_t)} =: \Psi_t(p)$$

Since $\Psi_t(0) = \Psi_t(\mu_t)$ and $p \mapsto \Psi_t(p)$ is differentiable at $p = 0$ it follows that

$$0 \leq \Psi_t'(0) = \frac{2\Delta_t(\mu_t)\Delta_t(x) - \Delta_t(\mu_t)^2}{I_t(\mu_t)}.$$

The claim follows by rearranging. \blacksquare

A.5. Proof of Lemma 7

For the analysis it is useful to define a lower bound on the regret,

$$\delta_t(x) = \min_{\theta \in \mathcal{C}_{t-1}} \max_{y \in \mathcal{X}} \langle y - x, \theta \rangle. \quad (7)$$

By definition, with probability at least $1 - \delta$ it holds that $\delta_t(x) \leq \langle x^* - x, \theta \rangle$. The set of plausible maximizers is equivalently described by $\mathcal{P}_t = \{x \in \mathcal{X} : \delta_t(x) = 0\}$ and by continuity \mathcal{P}_t is a compact set. We further define the relaxed bound $\tilde{\delta}_t(x) = \max_{x' \in \mathcal{X}} \min_{\theta \in \mathcal{C}_{t-1}} \langle x' - x, \theta \rangle$. By the minimax inequality it holds that $\tilde{\delta}_t(x) \leq \delta_t(x)$. For $\tilde{\delta}_t$, we can explicitly solve the inner maximization to get $\tilde{\delta}_t(x) = \max_{y \in \mathcal{X}} \langle y - x, \hat{\theta}_{t-1} \rangle - \beta_{t-1}^{1/2} \|y - x\|_{V_{t-1}^{-1}}$.

Lemma 12 *For the upper bound on the regret $\Delta_t(x)$, it holds that*

$$\Delta_t(x) = \max_{y \in \mathcal{P}_t} (y - x)^\top \hat{\theta}_{t-1} + \beta_{t-1}^{1/2} \|y - x\|_{V_{t-1}^{-1}},$$

where we restricted the maximum to plausible maximizers.

Proof [of Lemma 12] Assume that y is not a plausible maximizer, i.e. $\delta_t(y) > 0$. Then for any $\theta \in \mathcal{C}_t$, there exists a $z = z(\theta) \in \mathcal{X}$ s.t. $(z - y)^\top \theta > 0$. For fixed $x \in \mathcal{X}$ we find,

$$\max_{\theta \in \mathcal{C}_t} (y - x)^\top \theta < \max_{\theta \in \mathcal{C}_t} (z_\theta - x)^\top \theta \leq \max_{\theta \in \mathcal{C}_t} \max_z (z - x)^\top \theta = \Delta_t(x). \quad (8)$$

Hence, the left-hand side is maximized only if $y \in \mathcal{P}_t$ is a plausible maximizer. \blacksquare

Proof [of Lemma 7] Lemma 12 shows that we can write $\Delta_{t+1}(x)$ as follows:

$$\Delta_t(x) = \max_{y \in \mathcal{P}_t} \langle y - x, \hat{\theta}_{t-1} \rangle + \beta_{t-1}^{1/2} \|y - x\|_{V_{t-1}^{-1}}.$$

Further, for any plausible action $x \in \mathcal{P}_t$, we can bound the estimated gap by the associated uncertainty, $\langle y - x, \hat{\theta}_{t-1} \rangle \leq \beta_{t-1}^{1/2} \|y - x\|_{V_{t-1}^{-1}}$, which follows from the fact that $0 = \delta_t(x) \geq \tilde{\delta}_t(x)$.

This implies that for all $x \in \mathcal{P}_t$, $\Delta_t(x) \leq 2\beta_{t-1}^{1/2} \max_{y \in \mathcal{P}_t} \|x - y\|_{V_{t-1}^{-1}}$. Specifically, let w_t be the most uncertain direction in the set of plausible maximizers $w_t = \arg \max_{w \in \{x - y : x, y \in \mathcal{P}_t\}} \|w\|_{V_{t-1}^{-1}}^2$.

Then, for $x \in \mathcal{P}_t$,

$$\Delta_t(x)^2 \leq 4\beta_{t-1} \|w_t\|_{V_{t-1}^{-1}}^2 \leq 8\alpha_t \beta_{t-1} \max_{z \in \mathcal{P}_t} I_t(z).$$

The last step follows from the same argument as in the proof of Lemma 4, where we restrict x, y to \mathcal{P}_t and use the definition $\alpha_t = \alpha(\mathcal{C}_{t-1})$. \blacksquare

A.6. Bounds for the Alignment Constant

Lemma 13 Let $\mathcal{P} \subset \mathcal{X}$ finite with $p = |\mathcal{P}|$ such that for all $x, y \in \mathcal{P}$, $x - y \in \text{span}(A_z : z \in \mathcal{P})$. Let $A \in \mathbb{R}^{d \times pm}$ be the matrix formed by concatenating $(A_z : z \in \mathcal{P})$ and let B be a subset of at most d columns of A such that $\text{span}(B) = \text{span}(A)$. Then

$$\alpha(\mathcal{P}) = \max_{v \in \mathbb{R}^d} \max_{x, y \in \mathcal{P}, x \neq y} \frac{\langle x - y, v \rangle^2}{\max_{z \in \mathcal{P}} \|A_z^\top v\|^2} \leq \min_{w: Aw = x - y} \left(\sum_{z \in \mathcal{P}} \|w_z\| \right)^2 \leq d \lambda_{\min}(BB^\top)^{-1}.$$

Further, in the bandit game (where $A_x = x$), $\alpha(\mathcal{P}) \leq 4$.

Proof Let $x, y \in \mathcal{P}$ with $x \neq y$. By assumption, there exists a w such that $x - y = Aw$ with $w \neq 0$. Then,

$$\langle x - y, v \rangle^2 = \langle Aw, v \rangle^2 = \langle w, A^\top v \rangle^2 = \left(\sum_{z \in \mathcal{P}} \langle w_z, A_z^\top v \rangle \right)^2,$$

where we denote by $w_z \in \mathbb{R}^m$ the weights corresponding to A_z . An application of Cauchy-Schwarz proves the first inequality,

$$\alpha(\mathcal{P}) \leq \frac{\langle x - y, v \rangle^2}{\max_{z \in \mathcal{P}} \|A_z^\top v\|^2} \leq \frac{(\sum_{z \in \mathcal{P}} \|w_z\| \|A_z^\top v\|)^2}{\max_{z \in \mathcal{P}} \|A_z^\top v\|^2} \leq \left(\sum_{z \in \mathcal{P}} \|w_z\| \right)^2.$$

In the bandit game we can choose $w_x = 1$, $w_y = -1$ and $w_z = 0, \forall z \in \mathcal{P} \setminus \{x, y\}$, hence $\alpha(\mathcal{P}) \leq 4$. In general, we can choose $\mathcal{Q} \subset \mathcal{P}$ with $|\mathcal{Q}| = d$ s.t. $w_z = 0$ for $z \in \mathcal{P} \setminus \mathcal{Q}$. Therefore (we reuse the symbol w in a different dimension),

$$\alpha(\mathcal{P}) \leq \left(\sum_{z \in \mathcal{Q}} \|w_z\| \right)^2 \leq d \|w\|^2.$$

Denote $B = (A_z : z \in \mathcal{Q})$. The solution that minimizes the right-hand side is the ordinary least-squares solution $w^* = (B^\top B)^\dagger B^\top (x - y)$ where \dagger denotes the pseudo inverse. Therefore, using the properties of the pseudo inverse and $\|x - y\| \leq 1$,

$$\alpha(\mathcal{P}) \leq d \|w^*\|^2 \leq d \lambda_{\max} \left(B(B^\top B)^\dagger (B^\top B)^\dagger B^\top \right) = d \lambda_{\max} \left((BB^\top)^{-1} \right) = d \lambda_{\min} (BB^\top)^{-1}.$$

■

Appendix B. Regret Estimators and Information Gain Functions

B.1. Regret estimate

Our regret estimate $\Delta_t(x) = \max_{\theta \in \mathcal{C}_{t-1}} \max_{y \in \mathcal{X}} (y - x)^\top \theta$ is defined the tightest way for the given confidence bounds (up to truncation for bounded gaps). An interesting fact is that $\Delta_t(x)$ is a convex function because the maximum is over convex functions. The estimate can be relaxed to

$$\tilde{\Delta}_t(x) = \max_{y \in \mathcal{X}} \langle y, \hat{\theta}_{t-1} \rangle + \beta_{t-1}^{1/2} \|y\|_{V_{t-1}^{-1}} - \left(\langle x, \hat{\theta}_{t-1} \rangle - \beta_{t-1}^{1/2} \|x\|_{V_{t-1}^{-1}} \right).$$

It holds that $\Delta_t(x) \leq \tilde{\Delta}_t(x)$. For $\tilde{\Delta}_t$, the maximum over \mathcal{X} is independent of x , which reduces the computational complexity to compute the regret estimate from $\mathcal{O}(|\mathcal{X}|^2)$ to $\mathcal{O}(|\mathcal{X}|)$. The estimate $\tilde{\Delta}_t$ relies on directly estimating the value of $x^\top \theta$ for all actions $x \in \mathcal{X}$, which is not always possible in the general partial monitoring setting. The bandit game is an example where this is possible.

B.2. Directed Information Gain

Various ways of defining the information gain $I_t(x)$ are discussed in (Kirschner and Krause, 2018). The choice $I_t(x) = \log \det(\mathbf{1}_m + A_x^\top V_t^{-1} A_x)$ that we use in our main exposition is perhaps the most natural starting point, as it corresponds to the mutual information $\mathbb{I}(x, A_x \theta; \theta | \mathcal{F}_{t-1})$ if we define a corresponding Gaussian prior and likelihood. We denote $V_t | A_x = V_t + A_x A_x^\top$. For a fixed $w \in \mathbb{R}^d$ the *directed information gain* is

$$I_t(x; w) := \log \left(\|w\|_{V_{t-1}^{-1}}^2 \right) - \log \left(\|w\|_{(V_{t-1} | A_x)^{-1}}^2 \right). \quad (9)$$

The definition corresponds to the Shannon mutual information $\mathbb{I}(x, a_t; \langle w, \theta \rangle)$ which measures the Gaussian entropy reduction of θ projected onto the subspace spanned by w . The next lemma shows the information processing inequality $I_t(x; w) \leq I_t(x)$.

Lemma 14 (Information processing inequality) *For all $w, x \in \mathbb{R}^d$, $I_t(x; w) \leq I_t(x)$.*

Proof The proof is an exercises in linear algebra and makes use of the Sherman-Morrison formula and the matrix determinant lemma.

$$\begin{aligned}
 I_t(x; w) &= \log \left(\frac{\|w\|_{V_{t-1}^{-1}}^2}{\|w\|_{(V_{t-1}|A_x)^{-1}}^2} \right) \\
 &= -\log \left(1 - \frac{w^\top V_{t-1}^{-1} A_x (\mathbf{1}_m + A_x^\top V_{t-1}^{-1} A_x)^{-1} A_x^\top V_{t-1}^{-1} w}{\|w\|_{V_{t-1}^{-1}}^2} \right) \\
 &\leq \max_{v \in \mathbb{R}^d: \|v\|_2=1} -\log \left(v^\top v - v^\top V_{t-1}^{-1/2} A_x (\mathbf{1}_m + A_x^\top V_{t-1}^{-1} A_x)^{-1} A_x^\top V_{t-1}^{-1/2} v \right) \\
 &= \max_{v \in \mathbb{R}^d: \|v\|_2=1} -\log \left(v^\top \left(\mathbf{1}_d - V_{t-1}^{-1/2} A_x (\mathbf{1}_m + A_x^\top V_{t-1}^{-1} A_x)^{-1} A_x^\top V_{t-1}^{-1/2} \right) v \right) \\
 &= \log \left(\lambda_{\max} \left(\left(\mathbf{1}_d - V_{t-1}^{-1/2} A_x (\mathbf{1}_m + A_x^\top V_{t-1}^{-1} A_x)^{-1} A_x^\top V_{t-1}^{-1/2} \right)^{-1} \right) \right).
 \end{aligned}$$

We first used Sherman-Morrison to compute $(V_{t-1}|A_x)^{-1}$ and then maximize over $v = \frac{V_{t-1}^{-1/2} w}{\|w\|_{V_{t-1}^{-1}}}$.

$$\begin{aligned}
 &\leq \log \left(\det \left(\mathbf{1}_d - V_{t-1}^{-1/2} A_x (\mathbf{1}_m + A_x^\top V_{t-1}^{-1} A_x)^{-1} A_x^\top V_{t-1}^{-1/2} \right)^{-1} \right) \\
 &= \log \left(\det \left(\mathbf{1}_m + A_x^\top V_{t-1}^{-1} A_x \right) \det \left(\mathbf{1}_m + A_x^\top V_{t-1}^{-1} A_x - A_x^\top V_{t-1}^{-1/2} V_{t-1}^{-1/2} A_x \right)^{-1} \right) \\
 &= \log \det \left(\mathbf{1}_m + A_x^\top V_{t-1}^{-1} A_x \right) = I_t(x)
 \end{aligned}$$

The inequality follows because all eigenvalues of the matrix inside the determinant are not smaller than 1, and then the generalized matrix determinant lemma to rewrite the expression. \blacksquare

Lemma 15 *Let $\mathcal{Y} \subset \mathcal{X}$ be a subset of actions and let $w = x - y$ for $x, y \in \mathcal{Y}$ such that $w \in \text{span}(\{A_x : x \in \mathcal{Y}\})$. Then the most informative action in the set \mathcal{Y} satisfies*

$$\|w\|_{V_{t-1}^{-1}}^2 \leq 2\alpha(\mathcal{Y}) \max_{z \in \mathcal{Y}} I_t(z; w).$$

Proof First, note that $\frac{1}{2}(\mathbf{1}_d + A_z^\top V_{t-1}^{-1} A_z) \preceq \mathbf{1}_d$ by our assumption that $\|A_z\|_2 \leq 1$, hence

$$\|A_z^\top V_{t-1}^{-1} w\|^2 \leq 2w^\top V_{t-1}^{-1} A_z (\mathbf{1}_m + A_z^\top V_{t-1}^{-1} A_z)^{-1} A_z^\top V_{t-1}^{-1} w$$

We further bound the following fraction:

$$\min_{z \in \mathcal{Y}} \frac{(w^\top V_{t-1}^{-1} w)^2}{\|A_z^\top V_{t-1}^{-1} w\|^2} \leq \max_{x, y \in \mathcal{Y}} \max_{v \in \mathbb{R}^d} \min_{z \in \mathcal{Y}} \frac{\langle x - y, v \rangle^2}{\|A_z^\top v\|^2} = \alpha(\mathcal{Y}).$$

Since $x \leq -\log(1-x)$ for all $x \in [0, 1]$,

$$\begin{aligned}
 \|w\|_{V_{t-1}^{-1}}^2 &= \min_{z \in \mathcal{Y}} \frac{(w^\top V_{t-1}^{-1} w)^2}{\|A_z^\top V_{t-1}^{-1} w\|^2} \max_{z \in \mathcal{Y}} \frac{\|A_z^\top V_{t-1}^{-1} w\|^2}{w^\top V_{t-1}^{-1} w} \\
 &\leq 2\alpha(\mathcal{Y}) \max_{z \in \mathcal{Y}} \frac{w^\top V_{t-1}^{-1} A_z (\mathbf{1}_m + A_z^\top V_{t-1}^{-1} A_z)^{-1} A_z^\top V_{t-1}^{-1} w}{\|w\|_{V_{t-1}^{-1}}^2} \\
 &\leq -2\alpha(\mathcal{Y}) \max_{z \in \mathcal{Y}} \log \left(1 - \frac{w^\top V_{t-1}^{-1} A_z (\mathbf{1}_m + A_z^\top V_{t-1}^{-1} A_z)^{-1} A_z^\top V_{t-1}^{-1} w}{\|w\|_{V_{t-1}^{-1}}^2} \right) \\
 &= 2\alpha(\mathcal{Y}) \max_{z \in \mathcal{Y}} \log \left(\frac{\|w\|_{V_{t-1}^{-1}}^2}{\|w\|_{(V_{t-1}^{-1} A_z)^{-1}}^2} \right) \\
 &= 2\alpha(\mathcal{Y}) \max_{z \in \mathcal{Y}} I_t(z; w)
 \end{aligned}$$

This completes the proof. ■

Define the most uncertain direction in the set of plausible maximisers,

$$w_t = \arg \max_{w \in \{w=x-y: x, y \in \mathcal{P}_t\}} \|w\|_{V_{t-1}^{-1}}^2. \quad (10)$$

Our next results extends the regret bounds to the variant of IDS that uses $I_t(x, w_t)$ as information function. Note that the information processing inequality (Lemma 14) implies that $\sum_{t=1}^n I_t(x_t; w_t) \leq \sum_{t=1}^n I_t(x_t)$, and therefore the bound in Lemma 2 on the total information gain γ_n continues to hold.

Theorem 16 *IDS, defined with the directed information gain $I_t(x; w_t)$, achieves for any $n \geq 1$, with probability at least $1 - \delta$, $R_n \leq \mathcal{O}(n^{2/3}(\alpha\beta_n(\gamma_n + \log \frac{1}{\delta}))^{1/3})$ on globally observable games, and $R_n \leq \mathcal{O}((\alpha_0\beta_n(\gamma_n + \log \frac{1}{\delta})n)^{1/2})$ on uniformly locally observable games.*

Proof The proof is the same as for Theorem 3 and Theorem 6, but uses the stronger inequality of Lemma 15 to bound $\max_{x, y \in \mathcal{P}_t} \|x - y\|_{V_{t-1}^{-1}}^2 \leq \alpha(\mathcal{P}_t) \max_{z \in \mathcal{P}_t} I_t(x; w_t)$. ■

Unlike for IDS defined with $I_t(x)$, the information gain $I_t(x; w_t)$ requires to compute the set of plausible maximizers $\mathcal{P}_t = \{x \in \mathcal{X} : \delta_t(x) = 0\}$. This can be done by computing $\delta_t(x) = \min_{\theta \in \mathcal{C}_t} \max_{y \in \mathcal{X}} \langle y - x, \theta \rangle$ for each $x \in \mathcal{X}$. Note that the minimization over θ is on a convex function and therefore can be solved efficiently.

B.3. Relation to the UCB algorithm

Kirschner and Krause (2018) refer to the algorithm that chooses $x_t = \arg \min_{x \in \mathcal{X}} \Psi_t(\delta_x)$ as *deterministic IDS*. Optimizing over a deterministic action choice is computationally cheaper and sufficient to obtain $\tilde{\mathcal{O}}(\sqrt{n})$ regret on locally observable games as evident by Lemma 7. We draw a connection to the UCB algorithm. For $m = 1$ and $\|A_x\|_{V_{t-1}^{-1}}^2 \ll 1$ we have

$$I_t(x) = \log(1 + \|A_x\|_{V_{t-1}^{-1}}^2) \approx \|A_x\|_{V_{t-1}^{-1}}^2.$$

Define $\tilde{I}_t(x) = \|A_x\|_{V_{t-1}}^2$ and $\tilde{\Delta}_t(x) = \max_y \langle y, \hat{\theta}_{t-1} \rangle + \beta_{t-1}^{1/2} \|y\|_{V_{t-1}} - (\langle x, \hat{\theta}_{t-1} \rangle - \beta_{t-1}^{1/2} \|x\|_{V_{t-1}})$.

The next lemma shows that in bandit games ($A_x = x$), deterministic IDS with $\tilde{\Delta}_t$ and \tilde{I}_t as gap estimate and information gain, is equivalent to the UCB algorithm.

Lemma 17 *For a bandit game, let $x_t^{ucb} = \arg \max_{x \in \mathcal{X}} \langle x, \hat{\theta}_{t-1} \rangle + \beta_{t-1}^{1/2} \|x\|_{V_{t-1}}$ be the UCB action. Then,*

$$x_t^{ucb} \in \arg \min_{x \in \mathcal{X}} \frac{\tilde{\Delta}_t(x)^2}{\tilde{I}_t(x)}.$$

Proof A related result appears in (Wang et al., 2016, Lemma 2.1). The information-ratio of the UCB action is

$$\frac{\tilde{\Delta}_t(x_t^{ucb})^2}{\tilde{I}_t(x_t^{ucb})} = \frac{\left(\langle x_t^{ucb}, \hat{\theta}_{t-1} \rangle + \beta_{t-1}^{1/2} \|x_t^{ucb}\|_{V_{t-1}} - \left(\langle x_t^{ucb}, \hat{\theta}_{t-1} \rangle - \beta_{t-1}^{1/2} \|x_t^{ucb}\|_{V_{t-1}} \right) \right)^2}{\|x_t^{ucb}\|_{V_{t-1}}^2} = 4\beta_{t-1}.$$

Further, for any $x \in \mathcal{X}$, $\max_y \langle y, \hat{\theta}_{t-1} \rangle + \beta_{t-1}^{1/2} \|y\|_{V_{t-1}} \geq \langle x, \hat{\theta}_{t-1} \rangle + \beta_{t-1}^{1/2} \|x\|_{V_{t-1}}$, therefore

$$\frac{\tilde{\Delta}_t(x)^2}{\tilde{I}_t(x)} \geq \frac{\left(\langle x, \hat{\theta}_{t-1} \rangle + \beta_{t-1}^{1/2} \|x\|_{V_{t-1}} - \left(\langle x, \hat{\theta}_{t-1} \rangle - \beta_{t-1}^{1/2} \|x\|_{V_{t-1}} \right) \right)^2}{\|x\|_{V_{t-1}}^2} = 4\beta_{t-1}.$$

This shows that the UCB action minimizes the deterministic information ratio. \blacksquare

Appendix C. Applications and Extensions

We discuss applications and extensions. Note that we make use of the generalized setup (Appendix A.1) where necessary. In this case $\mathcal{P}(\mathcal{C}) \subset \mathcal{I}$ is defined to contain indexes plausible actions.

C.1. Full Information

The *full information setting* is perhaps not the most interesting case to study in the stochastic setting, because IDS reduces to the naive algorithm that aggregates the information and always plays greedy. Nevertheless, we demonstrate that the regret bounds improve given the additional information. Two natural settings are $A_x = \mathbf{1}_d$ and $A_x = X$ where $X = (x \in \mathcal{X})$ collects the actions as columns. In games where the information gain does not dependent on the action, IDS simply picks a regret minimizing action, $x_t = \arg \min_{x \in \mathcal{X}} \Delta_t(x)$. We show that IDS achieves $R_n \leq \tilde{O}(\sqrt{dn})$, which improves a factor \sqrt{d} compared to the bandit setting. For simplicity, let $A_x = \mathbf{1}_d$ and therefore $V_t = (t+1)\mathbf{1}_d$. The information gain is

$$I_t(x) = \log \det(\mathbf{1}_d + (t+1)^{-1}\mathbf{1}_d) = d \log(1 + \frac{1}{t}).$$

Hence $\gamma_n = d \log(n)$, but the ratio for the greedy action \hat{x}_t^* is

$$\Psi_t(\hat{x}_t^*) \leq \frac{\beta_t \max_{x,y} \|x - y\|_{V_{t-1}}^2}{d \log(1 + \frac{1}{t})} \approx \beta_t d^{-1}.$$

Given that $\beta_n \approx d \log(n)$, this means the overall bound is $R_n \leq \tilde{O}(\sqrt{dn})$. The same holds true for the directed information gain $I_t(x; w_t)$. Interestingly, here the improvements stem from a reduced total information gain $\gamma_n \approx \log(n)$, and the ratio remains $\Psi_t(x_t^*) \approx \beta_t \approx d \log(t)$.

C.2. Transductive Bandits

In the *transductive bandit setting* (Fiez et al., 2019) the learner has access to a set of informative actions $\mathcal{S} \subset \mathbb{R}^d$ for exploration and a set of actions $\mathcal{V} \subset \mathbb{R}^d$ that, when played, return reward. The sets are allowed to overlap or be contained in the other. In the original formulation the objective is to minimize the simple regret of a final recommendation on the target set \mathcal{V} by choosing actions only from \mathcal{S} . When the objective is to minimize cumulative regret, we can model this setting as a partial monitoring game by defining action-observation tuples $\mathcal{G}_1 = \{(0, x) : x \in \mathcal{S} \setminus \mathcal{V}\}$, $\mathcal{G}_2 = \{(x, 0) : x \in \mathcal{V} \setminus \mathcal{S}\}$ and $\mathcal{G}_3 = \{(x, x) : x \in \mathcal{S} \cap \mathcal{V}\}$, corresponding to informative actions with zero reward, actions that return reward but no information, and actions with the usual bandit information. The game is defined by $\mathcal{G} = \mathcal{G}_1 \cup \mathcal{G}_2 \cup \mathcal{G}_3$. Depending on the sets \mathcal{S} and \mathcal{V} , the game can be either locally observable or globally observable (or even infeasible).

C.3. Starved Bandits

In the *starved bandit setting* (Bubeck et al., 2018) the learner only receives information if the action is sampled from a predefined distribution. Let \mathcal{X}_0 be a ground set of actions that, when played, yield no information ($A_x = 0$). Denote by $\nu \in \mathcal{P}(\mathcal{X})$ the distribution that the learner can use for exploration and $z_t \sim \nu$ is a sample from the distribution in round t . The starved bandit setting is closely related to the contextual partial monitoring game with $\mathcal{G} = \mathcal{G}_0 \cup \{(z_t, A_{z_t} = z_t)\}$ added to the set of action-observation tuples. This game is globally observable if the distribution ν is sufficiently diverse such that the samples $(z_t)_{t=1}^n$ span the set of differences $\{x - y : x, y \in \mathcal{X}\}$. Note that on a curved actions set, the rate can still be $\mathcal{O}(\sqrt{n})$ as shown by Bubeck et al. (2018) (also compare our results on curved action sets in Section 2.4).

C.4. Batch Setting

In the batch setting, the learner commits to choosing B actions before observing the associated outcomes. This is important for applications where querying the objective for a number of actions in parallel is cheaper (or faster) than obtaining individual evaluations. This setting can be naturally formulated as a combinatorial partial monitoring game with semi-bandit feedback. Let $\mathcal{X}_0 \subset \mathbb{R}^d$ be a ground set of actions. The learner chooses a batch $(x_1, \dots, x_B) \in \mathcal{X}_0^B$. In the special case of a bandit feedback game, the reward is $\langle x_1 + \dots + x_B, \theta \rangle$ and the observation operator is $A_{x_1, \dots, x_B} = (x_1, \dots, x_B)$. With general feedback matrices, the batch game is

$$\mathcal{G} = \left\{ \left(\sum_{i=1}^B x_i, (A_{x_1}, \dots, A_{x_B}) \right) : (x_1, \dots, x_B) \in \mathcal{X}_0^B \right\}.$$

The bandit batch game is locally observable with $\alpha_0 \leq 4B^2$ (see Lemma 13). The disadvantage of this formulation is, however, that the action space is exponentially large. Finding an efficient approximation of the IDS distribution is an interesting direction for future work.

C.5. Dueling Bandits with Average Reward

Let $\mathcal{X}_0 \subset \mathbb{R}^d$ be a ground set of actions. The dueling bandit with average reward is the following game with index set $\mathcal{I} = \mathcal{X}_0 \times \mathcal{X}_0$:

$$\mathcal{G} = \left\{ \left(\frac{x_1 + x_2}{2}, x_1 - x_2 \right) : (x_1, x_2) \in \mathcal{I} \right\}.$$

In words, the learner can pick any pair of actions $x_1, x_2 \in \mathcal{X}_0$, obtains the average reward $(x_1 + x_2)^\top \theta / 2$ and a noisy observation of the reward difference $(x_1 - x_2)^\top \theta$. Note that the learner can also choose (x_1, x_1) with reward $x_1^\top \theta$ and no observation. Let $\mathcal{P}(\mathcal{C})$ be a plausible set of actions. The first observation is that if $(x_1, x_2) \in \mathcal{P}(\mathcal{C})$ then $(x_1, x_1) \in \mathcal{P}(\mathcal{C})$ and $(x_2, x_2) \in \mathcal{P}(\mathcal{C})$, because $(x_1 + x_2)/2 \in [x_1, x_2]$ lays on the line segment between x_1 and x_2 . Let $(x_1, x_2), (y_1, y_2) \in \mathcal{P}(\mathcal{C})$ be two plausible actions. We can choose a path (z_1, \dots, z_l) with $z_1 = x_1, z_l = y_1$ and $(z_i, z_{i+1}) \in \mathcal{P}(\mathcal{C})$. Therefore we can write

$$x_1 - y_1 = \sum_{i=1}^{l-1} x_i - x_{i+1} = \sum_{i=1}^{l-1} A_{(x_i, x_{i+1})}.$$

The difference $x_2 - y_2$ can be written similarly, which shows that $\frac{x_1 + x_2}{2} - \frac{y_1 + y_2}{2} \in \text{span}(A_i : i \in \mathcal{P}(\mathcal{C}))$. This shows that the game is locally observable. Turning to the local alignment constant

$$\alpha(\mathcal{P}) = \max_{v \in \mathbb{R}^d} \max_{(x_1, x_2), (y_1, y_2) \in \mathcal{P}} \min_{(z_1, z_2) \in \mathcal{P}} \frac{\langle (x_1 + x_2)/2 - (y_1 + y_2)/2, v \rangle^2}{\|(z_1 - z_2)^\top v\|^2}.$$

Using Lemma 13 and the path construction above we can bound $\alpha \leq l$ or $\alpha \leq Cd$.

Tightening the Alignment Constant Define the sets

$$\begin{aligned} \mathcal{Q}_1(\mathcal{C}) &= \{i \in \mathcal{I} : x_i \in \text{conv}(x_j : j \in \mathcal{P}(\mathcal{C}))\} \\ \mathcal{Q}_2(\mathcal{C}) &= \{i \in \mathcal{I} : \Delta_{\mathcal{C}}(x_i) \leq \max_{j \in \mathcal{P}(\mathcal{C})} \Delta_{\mathcal{C}}(x_j)\} \end{aligned}$$

with the regret estimate $\Delta_{\mathcal{C}}(x) = \max_{\theta \in \mathcal{C}} \max_{j \in \mathcal{P}(\mathcal{C})} (x_j - x)^\top \theta$. Note that $\Delta_{\mathcal{C}}(x)$ is a convex function which implies that $\mathcal{P}(\mathcal{C}) \subset \mathcal{Q}_1(\mathcal{C}) \subset \mathcal{Q}_2(\mathcal{C})$, but equality is not true in general. The observation is that in locally observable games, we can play actions in $\mathcal{Q}_2(\mathcal{P}_t)$ without worsening the regret bound. Consequently, the local alignment constant can be tightened to

$$\bar{\alpha}(\mathcal{C}) = \max_{v \in \mathbb{R}^d} \max_{i, j \in \mathcal{P}(\mathcal{C})} \min_{k \in \mathcal{Q}_2(\mathcal{C})} \frac{\langle x_i - x_j, v \rangle^2}{\|A_k^\top v\|^2}.$$

Clearly, $\bar{\alpha}(\mathcal{C}) \leq \alpha(\mathcal{C})$ and all regret bounds hold true with α replaced by $\bar{\alpha}$. For the dueling bandit game with average reward, recall that $x_1, y_1 \in \mathcal{P}(\mathcal{C})$ and therefore $(x_1 + y_1)/2 \in \text{conv}(x_1, y_1)$, and the same holds true for x_2, y_2 . This means we can now choose $A_{(x_1, y_1)} = x_1 - y_1$ and $A_{(x_2, y_2)} = x_2 - y_2$ as a response to estimate along the direction $(x_1 + x_2)/2 - (y_1 + y_2)/2$. We then write

$$\frac{x_1 + x_2}{2} - \frac{y_1 + y_2}{2} = \frac{1}{2}A_{(x_1, y_1)} + \frac{1}{2}A_{(x_2, y_2)},$$

and therefore, using the argument of Lemma 13, $\bar{\alpha}(\mathcal{C}) \leq 1$.

C.6. Partial Monitoring in Reproducing Kernel Hilbert Spaces

The *kernelized* setting is a practically relevant extension of the linear setting, where the feature dimension can be infinite. Let \mathcal{X}_0 be a ground set of actions, not to be confused with the features. This is often a subset of \mathbb{R}^d but can be defined on other structures (e.g. graphs) as well. The actions $x \in \mathcal{X}_0$ exhibit a non-linear dependence on the features through a positive-definite kernel map $k : \mathcal{X}_0 \times \mathcal{X}_0 \rightarrow \mathbb{R}$. Let \mathcal{H} be the reproducing kernel Hilbert space (RKHS) corresponding to the given kernel k and Hilbert norm $\|\cdot\|_{\mathcal{H}}$. Vectors in \mathcal{H} represent functions over \mathcal{X}_0 , so we denote the unknown parameter by $f \in \mathcal{H}$ (instead of θ). The standard boundedness assumption is that the unknown function has bounded Hilbert norm $\|f\|_{\mathcal{H}} \leq 1$. The kernel features $k_x = k(x, \cdot) \in \mathcal{H}$ satisfy $f(x) = \langle k_x, f \rangle$ according to the reproducing property and the set of kernel features associated to the actions is $\mathcal{X} = \{k_x : x \in \mathcal{X}_0\}$. The best action is $x^* = \arg \max_{x \in \mathcal{X}_0} f(x)$, and the regret is

$$R_n = \sum_{t=1}^n f(x^*) - f(x_t).$$

The linear observation functions are linear operators $A_x : \mathcal{H} \rightarrow \mathbb{R}^m$. As before, the observations when choosing x_t are $a_t = A_{x_t} f + \epsilon_t$. The regularized kernel least squares estimator is

$$\hat{f}_t = \arg \min_{f \in \mathcal{H}} \sum_{s=1}^t \|A_{x_s} f - a_s\|^2 + \|f\|_{\mathcal{H}}^2. \quad (11)$$

In the bandit setting, the *kernel trick* allows to express all quantities of interest in terms of the inner product $\langle k_x, k_y \rangle = k(x, y)$ evaluated on observed data points. In the general case where observations are generated from the observation operators A_x , we will need a slightly stronger assumption. Denote the adjoint map of A_x by $A_x^* : \mathbb{R}^m \rightarrow \mathcal{H}$. The requirement is that the matrix $M_{x,y} = A_x A_y^* \in \mathbb{R}^{m \times m}$ and the vectors $k_x A_y^* \in \mathbb{R}^m$ can be computed for any $x, y \in \mathcal{X}_0$ (the theory also holds without the assumption, but it is needed to implement the algorithm if the feature dimension is infinite). We detail such a computation in examples below. By (a slight modification of) the representer theorem, we can write the solution to (11) as $\hat{f}_t = \sum_{s=1}^t A_{x_s}^* \varphi_s$ for weights $\varphi_s \in \mathbb{R}^m$. Denote $\mathbf{a}_t \in \mathbb{R}^{mt}$ the vector that collects all observations $\mathbf{a}_t = (a_1^\top, \dots, a_t^\top)^\top$, $K_t \in \mathbb{R}^{mt \times mt}$ the kernel matrix that collects the matrices $(K_t)_{ij} = M_{x_i, y_j}$ and $k_t(x) \in \mathbb{R}^{mt}$ the evaluation vector $k_t(x) = (k_x A_{x_1}^*, \dots, k_x A_{x_t}^*)^\top$. The solution $\hat{f}_t(x) = \langle k_x, \hat{f}_t \rangle$ to the least squares problem evaluated at $x \in \mathcal{X}_0$ is

$$\hat{f}_t(x) = k_t(x)^\top (K_t + \mathbf{1}_{mt})^{-1} \mathbf{a}_t.$$

The estimate corresponds to the posterior mean of a Gaussian process (GP) model with kernel k and Gaussian likelihood (c.f. Kanagawa et al., 2018). The gap estimate at time $t + 1$ is defined as

$$\Delta_{t+1}(x) = \max_{y \in \mathcal{X}_0} \hat{f}_t(y) - \hat{f}_t(x) + \beta_t \sqrt{\sigma_t(x)^2 + \sigma_t(y)^2 - 2k_t(x, y)},$$

where

$$\begin{aligned} k_t(x, y) &= k(x, y) - k_t(x)^\top (K_t + \mathbf{1}_{mt})^{-1} k_t(y), \\ \sigma_t(x) &= \sqrt{k_t(x, x)}, \\ \beta_t^{1/2} &= \sqrt{\log \det(K_t + \mathbf{1}_{mt}) + 2 \log \frac{1}{\delta} + 1}. \end{aligned}$$

The estimate is chosen such that with probability at least $1 - \delta$, $f(x^*) - f(x) \leq \Delta_t(x)$ for any $x \in \mathcal{X}$ and $t \geq 1$ (Abbasi-Yadkori, 2012, Theorem 3.11).

To compute the information gain, define $M_t(x) = (M_{x,x_1}^\top, \dots, M_{x,x_t}^\top)^\top \in \mathbb{R}^{mt \times m}$. The kernelized information gain (4) is given by

$$I_t(x) = \log \det \left(\mathbf{1}_m + M_{x,x} - M_t(x)^\top K_t^{-1} M_t(x) \right).$$

Denote by $k_{t|A_z}$ and $\sigma_{t|A_z}$ the uncertainty estimates that are (tentatively) updated with an observation generated from A_z . Such an update does not require the observation outcome y_t , similar to the linear case, where we can update the precision matrix $V_t|A_z = V_t + A_z A_z^\top$. Further, let $w_{x,y} = k_x - k_y$ be the difference of kernel features for the gap difference that we want to estimate. The kernelized directed information gain is

$$I_t(z; w_{x,y}) = \log \left(\frac{\sigma_t(x)^2 + \sigma_t(y)^2 - 2k_t(x,y)}{\sigma_{t|A_z}(x)^2 + \sigma_{t|A_z}(y)^2 - 2k_{t|A_z}(x,y)} \right).$$

As before the information processing inequality (Lemma 14) implies that $I_t(z; w_{x,y}) \leq I_t(z)$. The bound in Lemma 2 on the total information gain $\gamma_n = \sum_{t=1}^n I_t(x_t) = \log \det(K_n + \mathbf{1}_{mn})$ for finite feature dimension can be replaced by bounds that depend on the eigenspectrum of the kernel (Srinivas et al., 2010), for example $\gamma_n = \mathcal{O}(\log(n)^{d+1})$ for the squared-exponential kernel on \mathbb{R}^d . We remark that in the kernelized setting, only the computation of the estimator and information gain are different compared to the linear setting. The regret analysis remains the same with the appropriate constants β_n and γ_n , defined above. We therefore summarize our result:

Corollary 18 *The kernelized variant of IDS achieves $R_n \leq \mathcal{O}(n^{2/3}(\alpha\beta_n(\gamma_n + \log \frac{1}{\delta}))^{1/3})$ on globally observable games and $R_n \leq \mathcal{O}((\alpha_0\beta_n(\gamma_n + \log \frac{1}{\delta})n)^{1/2})$ on uniformly locally observable games for any $n \geq 1$ with probability at least $1 - \delta$.*

Example: Kernelized Dueling Bandits We illustrate a dueling bandit setting, where the learner chooses two actions $(x, x') \in \mathcal{X}_0^2$ and observes binary feedback on $f(x) \geq f(x')$. In the partial monitoring formulation, the observation operator is $A_{x,x'} = k_x - k_{x'}$, which means that the learner observes $\langle k_x - k_{x'}, f \rangle = f(x) - f(x')$ up to noise. The learner obtains the reward of the first action (other reward models are possible), so the set of action-observation tuples is

$$\mathcal{G} = \{(k_x, A_{x,x'} = k_x - k_{x'}) : (x, x') \in \mathcal{X}_0 \times \mathcal{X}_0\}.$$

The noise on the observation $a_t = f(x_t) - f(x'_t) + \epsilon_t$ is such that $a_t \in \{-1, 1\}$ and $\mathbb{E}[a_t] = f(x_t) - f(x'_t)$ (i.e. $\mathbb{P}[a_t = 1] = 1 - \mathbb{P}[a_t = -1] = (1 + f(x) - f(x'))/2$). The quantities that are required to compute the estimator are

$$\begin{aligned} M_{(x,x'),(z,z')} &= (k_x - k_{x'})(k_z - k_{z'})^* = k(x, z) - k(x, z') - k(x', z) + k(x', z'), \\ k_x A_{(z,z')}^* &= k_x(k_z + k_{z'})^* = k(x, z) - k(x, z'). \end{aligned}$$

Kernelized dueling bandits have been studied in the literature (González et al., 2017; Sui et al., 2017a, 2018) as well as extensions with multi-point comparisons (Sui et al., 2017b). Assuming that the learner can compare any pair of actions (x, x') , the setting is locally-observable by nature with $\alpha \leq 1$. Therefore, IDS achieves a $\tilde{\mathcal{O}}(\sqrt{\beta_n \gamma_n n})$ regret bound. The same holds true for the deterministic variant that simply chooses the action which minimizes the information ratio.

Example: Bayesian Optimization with Gradients While Bayesian optimization (or kernelized bandits) is typically phrased for the noisy, zero-order oracle, previous work also incorporates gradient information where it is available (Wu et al., 2017b). We illustrate a setting where the learner *only* observes the gradient, which can be understood as a type of dueling bandit. Let $\mathcal{X}_0 \subset \mathbb{R}^d$ be a compact, connected domain and $f : \mathcal{X}_0 \rightarrow \mathbb{R}^d$ be an element in \mathcal{H} with a kernel that guarantees that any $f \in \mathcal{H}$ is continuously differentiable. The gradient ∇_x acts linearly on the function f and therefore is a valid choice for the observation operator $A_x : \mathcal{H} \rightarrow \mathbb{R}^d$ with $m = d$. The key step is to compute the quantities required for the estimation,

$$\begin{aligned} (A_x A_y^*)_{ij} &= \langle e_i, A_x A_y^* e_j \rangle = \langle \nabla_x \langle A_y k_x, e_j \rangle \rangle_i = \frac{\partial}{\partial y_i} \frac{\partial}{\partial x_j} k(x, y), \\ (k_x A_y^*)_i &= \langle k_x, A_y^* e_i \rangle = \langle \nabla_y k_x, e_i \rangle = \frac{\partial}{\partial y_i} k(x, y). \end{aligned}$$

The game where the learner observes only the gradient is globally observable, which means that for all $x, y \in \mathcal{X}_0$, $k_x - k_y \in \text{span}(A_x^* : x \in \mathcal{X}_0)$. To see this, let $\tau : [0, 1] \rightarrow \mathcal{X}_0$ be a differentiable path with $\tau(0) = x$ and $\tau(1) = y$. We claim that

$$k_x - k_y = \int_0^1 A_{\alpha(t)}^* \dot{\alpha}(t) dt.$$

This is verified, because for any $f \in \mathcal{H}$ by the fundamental theorem of calculus,

$$\begin{aligned} \left\langle \int_0^1 A_{\alpha(t)}^* \dot{\alpha}(t) dt, f \right\rangle &= \int_0^1 \langle A_{\alpha(t)}^* \dot{\alpha}(t), f \rangle dt \\ &= \int_0^1 \langle \dot{\alpha}(t), A_x f \rangle dt \\ &= \int_0^1 \langle \dot{\alpha}(t), \nabla_x f \rangle dt \\ &= f(x) - f(y) = \langle k_x - k_y, f \rangle. \end{aligned}$$

If the learner observes both the function value and the gradient, the game is locally observable.

Example: Invasive Laser Alignment Consider a simplistic setup, where an experimenter wishes to align a laser on a squared target using two parameters $(x_1, x_2) \in \mathcal{X}_0 = [-1, 1]^2$ that correspond to a vertical and horizontal shift of the device (see Figure 1 for an illustration). The power of the laser on a two-dimensional plane is given by an (initially) unknown function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$. In the illustrated example it is set to $f(z_1, z_2) = \exp(-((z_1 - 0.5)^2 + (z_2 - 0.5)^2))$. The objective is to find a parameter setting that maximize the integrated intensity on the (e.g. 1×1) target, $I_t(x_1, x_2) = \int_{x_1}^{x_1+1} \int_{x_2}^{x_2+1} f(z_1, z_2) dz_1 dz_2$. At any step, the experimenter can choose to evaluate a setting (x_1, x_2) and observes the corresponding intensity $I_t(x_1, x_2)$ up to noise. Since the intensity is the reward, this action has standard bandit feedback. Alternatively, the experimenter can drive a screen into the laser beam to measure the laser power on a $m \times m$ -grid $G(x_1, x_2)$ centered at (x_1, x_2) , which yields m^2 noisy measurements $(f(z_1, z_2) : (z_1, z_2) \in G(x_1, x_2))$, possibly at a lower noise level than the integrated intensity measurement. As the screen blocks of the beam, there is no reward in such rounds (hence the term ‘invasive measurement’). The learner therefore has the choice between a direct measurement of the objective and a more informative action that yields no

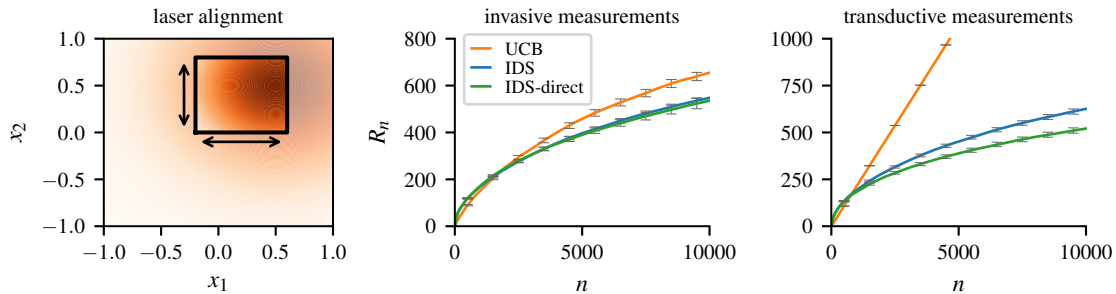


Figure 1: A demonstration of the stylized laser example. The left plot shows the energy of the laser on the two dimensional plane. The objective is to shift the square target such that the integrated intensity within the square is maximized. The learner chooses to either observe a noisy measurement of the intensity, or alternatively, the energy function directly, evaluated on a measurement grid within the square (invasive measurements). The latter feedback is obtained from a screen that is put in the line of the laser, which blocks the beam and voids the reward signal. In the second variant (transductive measurements), the learner obtains information *only* through the grid measurements. To solve the task, the learner needs to estimate the function and ‘blindly’ move the target to the position with maximum integrated intensity. The plots on the right show the regret of IDS (with directed and undirected information gain) compared to the UCB algorithm. Note that UCB *never* chooses the informative actions and therefore suffers linear regret on the second task.

reward. Clearly, the game is locally observable as each action contains the bandit feedback. We remark that the UCB algorithm *never* chooses the invasive measurements, because the UCB score for these actions is always zero. On the other hand, IDS naturally trades off between the informative actions and those that lead to reward. In a (transductive) variant of the setup, the signal can *only* be observed through the invasive measurements and the integrated signal is *not* observed. In this case, the game is globally, but not locally observable.

We present a numerical simulation of this setup in Figure 1. Our set \mathcal{X}_0 is discrete with 9 actions corresponding to a unit shift in any direction (or no shift). We use 25-dimensional features computed from a radial basis function kernel. In the setup where the reward signal can be observed directly, UCB outperforms IDS for the first ~ 1000 steps; but then IDS gains an advantage from choosing the more informative measurements from time to time. Without the direct reward observation, UCB continues to play actions that yield the integrated reward, but no longer receives any information. The parameter estimate is therefore never updated, and the UCB algorithm suffers linear regret. On the other hand, IDS still achieves no-regret through trading off the informative measurements with parameter settings that yield reward.

Appendix D. Convex Action Sets

The proof of Theorem 8 follows by using the curvature to bound the information ratio. We will show the following:

$$\Psi_t(\mu_t) \leq C\beta_{t-1} \max\left(\frac{\text{diam}(\mathcal{X})}{\kappa_o}, \text{diam}(\mathcal{X})^2\right).$$

where $C > 0$ is a constant depending only on $(A_z : z \in \mathcal{X})$. Recall the definition of the support function $h_{\mathcal{X}}(u) = \sup_{x \in \mathcal{X}} \langle x, u \rangle$. A simple calculation shows that $\nabla h_{\mathcal{X}}(u) = \arg \max_{x \in \mathcal{X}} \langle x, u \rangle$, and $\hat{x}_t^* = \nabla h_{\mathcal{X}}(\hat{\theta}_t)$ is the greedy action. Before the proof of the theorem we need a simple lemma bounding the regret in terms of the curvature.

Lemma 19 *Suppose that $1/\kappa_o = \max_{x \in \mathcal{X}} \lambda_{\max}(\nabla^2 h_{\mathcal{X}}(x))$. Then for any $\theta, \theta' \in \mathbb{R}^d$,*

$$\langle \nabla h_{\mathcal{X}}(\theta) - \nabla h_{\mathcal{X}}(\theta'), \theta \rangle \leq \frac{2\|\theta - \theta'\|^2}{\kappa_o \|\theta\|}.$$

Proof Abbreviate $\eta = \theta / \|\theta\|$ and $\eta' = \theta' / \|\theta'\|$. Note that for $c > 0$, $h_{\mathcal{X}}(c \cdot u) = c \cdot h_{\mathcal{X}}(u)$, which implies that $\nabla h_{\mathcal{X}}(c \cdot u) = \nabla h_{\mathcal{X}}(u)$. Using the definitions,

$$\begin{aligned} \langle \nabla h_{\mathcal{X}}(\theta) - \nabla h_{\mathcal{X}}(\theta'), \theta \rangle &= \langle \nabla h_{\mathcal{X}}(\eta) - \nabla h_{\mathcal{X}}(\eta'), \theta \rangle \\ &= \|\theta\| \langle \nabla h_{\mathcal{X}}(\eta) - \nabla h_{\mathcal{X}}(\eta'), \eta \rangle \\ &\stackrel{(i)}{\leq} \|\theta\| \langle \nabla h_{\mathcal{X}}(\eta) - \nabla h_{\mathcal{X}}(\eta'), \eta - \eta' \rangle \\ &\stackrel{(ii)}{\leq} \|\theta\| \frac{\|\eta - \eta'\|^2}{\kappa_o} \\ &\stackrel{(iii)}{\leq} \frac{2\|\theta - \theta'\|^2}{\kappa_o \|\theta\|}, \end{aligned}$$

where inequality (i) follows because $\nabla h_{\mathcal{X}}(\eta') = \arg \max_{x \in \mathcal{X}} \langle x, \eta' \rangle$. The second inequality (ii) follows from the definition of κ_o and because

$$\langle \nabla h_{\mathcal{X}}(\eta) - \nabla h_{\mathcal{X}}(\eta'), \eta - \eta' \rangle = \int_0^1 \|\eta - \eta'\|_{\nabla^2 h_{\mathcal{X}}((1-t)\eta + t\eta')}^2 dt \leq \frac{\|\eta - \eta'\|^2}{\kappa_o}.$$

The last inequality (iii) follows from the following geometric inequality:

$$\forall x, y \in \mathbb{R}^d \quad \left\| x - \frac{\|x\|}{\|y\|} y \right\|^2 \leq 2\|x - y\|^2.$$

■

Proof of Theorem 8 Let $\hat{x}_t^* = \nabla h_{\mathcal{X}}(\hat{\theta}_{t-1})$. Then, by Lemma 19,

$$\Delta_t(x_t^*) \leq \max_{y \in \mathcal{X}, \phi \in \mathcal{C}_{t-1}} \langle y - x, \phi \rangle \leq \max_{\phi \in \mathcal{C}_{t-1}} \frac{2\|\phi - \hat{\theta}_{t-1}\|^2}{\kappa_o \|\phi\|}.$$

Let $z_t = \arg \max_{z \in \mathcal{X}} I_t(z)$. By the assumption that \mathcal{X} spans \mathbb{R}^d and the definition of global observability, it follows that $\text{span}(A_z : z \in \mathcal{X}) = \mathbb{R}^d$, which means there exists a constant c depending only on $(A_z : z \in \mathcal{X})$ such that

$$\lambda_{\max}(V_{t-1}^{-1}) \leq c \lambda_{\max}(A_{z_t}^\top V_{t-1}^{-1} A_{z_t}) \leq c I_t(z_t),$$

where the second inequality follows from the same argument as in Lemma 4. Hence,

$$\begin{aligned} \text{diam}(\mathcal{C}_{t-1})^2 &= \max_{\theta, \phi \in \mathcal{C}_{t-1}} \|\theta - \phi\|^2 \\ &\leq \max_{\theta, \phi \in \mathcal{C}_{t-1}} \lambda_{\max}(V_{t-1}^{-1}) \|\theta - \phi\|_{V_{t-1}}^2 \\ &\leq \lambda_{\max}(V_{t-1}^{-1}) \beta_{t-1} \\ &\leq c \beta_{t-1} I_t(z_t). \end{aligned}$$

The analysis of the information ratio is decomposed into two cases. The first case is when \mathcal{C}_{t-1} has a large diameter, in which case the information ratio is well controlled without using curvature, and by only exploration. Suppose that

$$2 \max \left(1, \sqrt{\frac{1}{\kappa_\circ \text{diam}(\mathcal{X})}} \right) \text{diam}(\mathcal{C}_{t-1}) \geq \max_{\phi \in \mathcal{C}_{t-1}} \|\phi\| \quad (12)$$

Then, using Cauchy–Schwarz inequality and the definition of Δ_t ,

$$\begin{aligned} \Delta_t(z_t)^2 &\leq \text{diam}(\mathcal{X})^2 \max_{\phi \in \mathcal{C}_{t-1}} \|\phi\|^2 \\ &\leq 4 \max \left(1, \frac{1}{\kappa_\circ \text{diam}(\mathcal{X})} \right) \text{diam}(\mathcal{X})^2 \text{diam}(\mathcal{C}_{t-1})^2 \\ &\leq 4c \beta_{t-1} \max \left(1, \frac{1}{\kappa_\circ \text{diam}(\mathcal{X})} \right) \text{diam}(\mathcal{X})^2 I_t(z_t), \end{aligned}$$

which implies that

$$\Psi_t(z_t) \leq 4c \beta_{t-1} \max \left(1, \frac{1}{\kappa_\circ \text{diam}(\mathcal{X})} \right) \text{diam}(\mathcal{X})^2.$$

Moving to the second case where Eq. (12) does not hold. Let

$$p = \frac{2 \text{diam}(\mathcal{C}_{t-1})^2}{\kappa_\circ \min_{\phi \in \mathcal{C}_{t-1}} \|\phi\| \max_{\phi \in \mathcal{C}_{t-1}} \|\phi\| \text{diam}(\mathcal{X})}.$$

That $p \in [0, 1]$ follows by virtue of the fact that

$$\min_{\phi \in \mathcal{C}_{t-1}} \|\phi\| \geq \max_{\phi \in \mathcal{C}_{t-1}} \|\phi\| - \text{diam}(\mathcal{C}_{t-1}) \geq \frac{1}{2} \max_{\phi \in \mathcal{C}_{t-1}} \|\phi\|.$$

Hence, $\mu = (1 - p)\delta_{\hat{x}_t^*} + p\delta_{z_t} \in \mathcal{P}(\mathcal{X})$. Using that $\hat{x}_t^* = \nabla h_{\mathcal{X}}(\hat{\theta}_{t-1})$ and Lemma 19,

$$\begin{aligned} \Delta_t(\mu) &\leq p \operatorname{diam}(\mathcal{X}) \max_{\phi \in \mathcal{C}_{t-1}} \|\phi\| + \max_{\phi \in \mathcal{C}_{t-1}} \langle \nabla h_{\mathcal{X}}(\phi) - \nabla h_{\mathcal{X}}(\hat{\theta}_{t-1}), \phi \rangle \\ &\leq p \operatorname{diam}(\mathcal{X}) \max_{\phi \in \mathcal{C}_{t-1}} \|\phi\| + \frac{2}{\kappa_o} \max_{\phi \in \mathcal{C}_{t-1}} \frac{\|\hat{\theta}_{t-1} - \phi\|^2}{\|\phi\|} \\ &\leq p \operatorname{diam}(\mathcal{X}) \max_{\phi \in \mathcal{C}_{t-1}} \|\phi\| + \frac{2 \operatorname{diam}(\mathcal{C}_{t-1})^2}{\kappa_o \min_{\phi \in \mathcal{C}_{t-1}} \|\phi\|} \\ &= \frac{4 \operatorname{diam}(\mathcal{C}_{t-1})^2}{\kappa_o \min_{\phi \in \mathcal{C}_{t-1}} \|\phi\|}. \end{aligned}$$

Therefore, using the fact that $c\beta_{t-1}I_t(\mu) \geq c\beta_{t-1}pI_t(z_t) \geq p \operatorname{diam}(\mathcal{C}_{t-1})^2$,

$$\begin{aligned} \frac{\Delta_t(\mu)^2}{I_t(\mu)} &\leq \frac{16c\beta_{t-1} \operatorname{diam}(\mathcal{C}_{t-1})^2}{p\kappa_o^2 \min_{\phi \in \mathcal{C}_{t-1}} \|\phi\|^2} \\ &= \frac{8c\beta_{t-1} \operatorname{diam}(\mathcal{X}) \max_{\phi \in \mathcal{C}_{t-1}} \|\phi\|}{\kappa_o \min_{\phi \in \mathcal{C}_{t-1}} \|\phi\|} \\ &\leq \frac{16c\beta_{t-1} \operatorname{diam}(\mathcal{X})}{\kappa_o}. \end{aligned}$$

Combining the two parts shows that

$$\Psi_t(\mu_t) \leq c\beta_{t-1} \max \left(\frac{16 \operatorname{diam}(\mathcal{X})}{\kappa_o}, 4 \max \left(1, \frac{1}{\kappa_o \operatorname{diam}(\mathcal{X})} \right) \operatorname{diam}(\mathcal{X})^2 \right).$$

With the bound on the information ratio and Lemma 1, the proof of Theorem 8 follows now immediately. \blacksquare

Appendix E. Contextual Partial Monitoring

E.1. Conditional IDS for Contextual Games

Conditional IDS optimizes the sampling distribution for the given context z_t ,

$$\mu_t = \arg \min_{\mu} \Psi_t(\mu; z_t).$$

The computational complexity required to find the minimizer of the information ratio is the same as in the non-contextual case. We extend the notion of the alignment-constant with the contextual argument,

$$\alpha(\mathcal{C}, z) = \max_{v \in \mathbb{R}^d} \max_{x, y \in \mathcal{P}(\mathcal{C}, z)} \min_{u \in \mathcal{P}(\mathcal{C}, z)} \frac{\langle x - y, v \rangle^2}{\|A_u^{z^\top} v\|^2}, \quad (13)$$

where $\mathcal{P}(\mathcal{C}, z) = \cup_{\theta \in \mathcal{C}} \{x \in \mathcal{X}_z : \langle x, \theta \rangle \geq \max_{y \in \mathcal{X}_z} \langle y, \theta \rangle\}$ is the plausible maximizer set for context z . For globally observable games, we denote $\alpha(z) = \alpha(\mathbb{R}^d, z)$.

The next result is an immediate upper bound for the regret of conditional IDS under the assumption that for any context $z \in \mathcal{Z}$, each game $(\mathcal{X}_z, \mathcal{A}_z)$ is globally or locally observable, respectively.

Corollary 20 *If a contextual game is globally observable in the sense that for any context z , the game $(\mathcal{X}_z, \mathcal{A}_z)$ is globally observable with uniformly bounded alignment constant $\alpha(z) \leq \alpha$, then for any $n \geq 1$ with probability at least $1 - \delta$, conditional IDS achieves*

$$R_n \leq Cn^{2/3} (\alpha\beta_n(\gamma_n + \log \frac{1}{\delta}))^{1/3} + 4 \log \left(\frac{4n+4}{\delta} \right).$$

If the contextual game is locally observable in the sense that for any $z \in \mathcal{Z}$, the game $(\mathcal{X}_z, \mathcal{A}_z)$ is locally observable with uniformly bounded alignment constant $\alpha(\mathcal{C}; z) \leq \alpha_0$ for all convex $\mathcal{C} \subset \mathbb{R}^d$, then for any $n \geq 1$ with probability at least $1 - \delta$, conditional IDS achieves

$$R_n \leq C\sqrt{n\alpha_0\beta_n(\gamma_n + \log \frac{1}{\delta})} + 4 \log \left(\frac{4n+4}{\delta} \right),$$

where C is a universal constant.

The corollary follows along the lines of our main results, Theorem 3 & 6. The assumptions of Corollary 20 imply that the information ratio is bounded for any context z in the respective regimes. One can achieve a slightly stronger result by replacing the alignment constant α_0 with the average observed alignment $\frac{1}{n} \sum_{t=1}^n \alpha(\mathcal{C}_t; z_t)$. In this case the bound explicitly depends on the sequence of observed contexts $(z_t)_{t=1}^n$ and the confidence sets $(\mathcal{C}_t)_{t=1}^n$, which can lead to improved bounds in benign cases.

E.2. Regret Bounds for Contextual IDS

In this section we summarize results for *contextual IDS*, which minimizes

$$\xi_t = \arg \min_{\xi \in \mathcal{P}(\mathcal{X} \times \mathcal{Z}), \xi_z = \nu} \Psi(\xi),$$

where $\nu \in \mathcal{P}(\mathcal{Z})$ is a known distribution over the set of contexts. For general compact $\mathcal{X} \times \mathcal{Z}$, Prokhorov's theorem guarantees the existence of a minimizer (cf. [Kirschner and Krause, 2018](#)).

We overload the notation and let $\mathcal{X} = \times_{z \in \mathcal{Z}} \mathcal{X}_z$ be the joint action space over all contexts and $\mathcal{P}_t = \times_{z \in \mathcal{Z}} \mathcal{P}(\mathcal{C}_t; z)$ the joined set of plausible maximisers. For a function $g : \mathcal{Z} \rightarrow \mathcal{Z}$ and a vector $x \in \mathcal{X}$, we define $x_g \in \mathbb{R}^{|\mathcal{X}|}$ by $(x_g)_z = x_{g(z)}$. The *expected alignment constant* is defined as

$$\alpha(\mathcal{C}, \nu) = \max_{v \in (\mathbb{R}^d)^{\mathcal{Z}}} \max_{x, y \in \mathcal{P}(\mathcal{C})} \min_{u \in \mathcal{P}(\mathcal{C})} \min_{g: \mathcal{Z} \rightarrow \mathcal{Z}} \mathbb{E}_{\nu} \left[\langle v, x - y \rangle^{1/2} \right]^2 \mathbb{E}_{\nu} \left[\frac{\|A_u v_g\|^2}{\langle v_g, x_g - y_g \rangle} \right]^{-1}. \quad (14)$$

As before, this corresponds to the signal-to-noise ratio that can be achieved by choosing the best aligned observation operator \mathcal{A}_u^z in context z , with the additional twist that learner can choose to estimate along a direction $x_{z'} - y_{z'}$ of a different context $z' = g(z)$. In the next lemma, we show that the definition satisfies more intuitive upper bounds. We will see that in the finite case, the definition of the alignment constant relates to natural conditions for local and global observability.

Lemma 21 *Let $\alpha(\mathcal{C}, \nu)$ be the expected alignment (14) and $\alpha(\mathcal{C}, z)$ the conditional alignment (13).*

i) *For any convex $\mathcal{C} \subset \mathbb{R}^d$ and distribution $\nu \in \mathcal{P}(\mathcal{Z})$ it holds that,*

$$\alpha(\mathcal{C}; \nu) \leq \mathbb{E}_{z \sim \nu} [\alpha(\mathcal{C}; z)].$$

ii) For finite context sets,

$$\alpha(\mathcal{C}, \nu) \leq \max_{z \in \mathcal{Z}} \max_{x, y \in \mathcal{P}(\mathcal{C}; z), v \in \mathbb{R}^d} \min_{z' \in \mathcal{Z}} \min_{u \in \mathcal{P}(\mathcal{C}, z')} \frac{\langle v, x - y \rangle^2}{\nu(z') \|A_{u'}^{z'} v\|^2}.$$

iii) For Dirac-delta distributions $\nu = \delta_z$,

$$\alpha(\mathcal{C}, \delta_z) = \alpha(\mathcal{C}, z).$$

The first inequality implies that the information ratio of *contextual IDS* is never worse than for *conditional IDS*. The second inequality captures the intuition that for every direction $x - y \in \mathcal{X}_z$ in a context z , there needs to be a context z' that appears with positive probability $\nu(z') > 0$ where $x - y$ can be estimated. The last inequality is a sanity check which shows that for a constant context, we recover the previous definitions.

Proof For i), note that

$$\begin{aligned} \alpha(\mathcal{C}; \nu) &= \max_{v \in (\mathbb{R}^d)^{\mathcal{Z}}} \max_{x, y \in \mathcal{P}(\mathcal{C})} \min_{u \in \mathcal{P}(\mathcal{C})} \min_{g: \mathcal{Z} \rightarrow \mathcal{Z}} \mathbb{E}_\nu \left[\langle v_z, x_z - y_z \rangle^{1/2} \right]^2 \mathbb{E}_\nu \left[\frac{\|A_{u_z}^z v_{g(z)}\|^2}{\langle v_{g(z)}, x_{g(z)} - y_{g(z)} \rangle} \right]^{-1} \\ &\leq \max_{v \in (\mathbb{R}^d)^{\mathcal{Z}}} \max_{x, y \in \mathcal{P}(\mathcal{C})} \min_{u \in \mathcal{P}(\mathcal{C})} \mathbb{E}_\nu \left[\langle v_z, x_z - y_z \rangle^{1/2} \right]^2 \mathbb{E}_\nu \left[\frac{\|A_{u_z}^z v_z\|^2}{\langle v_z, x_z - y_z \rangle} \right]^{-1} \\ &\leq \max_{v \in (\mathbb{R}^d)^{\mathcal{Z}}} \max_{x, y \in \mathcal{P}(\mathcal{C})} \min_{u \in \mathcal{P}(\mathcal{C})} \mathbb{E}_\nu \left[\frac{\langle v_z, x_z - y_z \rangle^2}{\|A_{u_z}^z v_z\|^2} \right] \\ &= \mathbb{E}_\nu \alpha(\mathcal{C}; z). \end{aligned}$$

The first inequality follows by choosing the identity function $g(z) = z$. The second inequality uses the fact that $(a, b) \mapsto a^2/b$ is convex on $\mathbb{R} \times \mathbb{R}_{>0}$ and Jensen's inequality.

For ii), denote $z^* = z^*(v, x, y) = \arg \max_{z \in \mathcal{Z}} \langle v_z, x_z - y_z \rangle$ and define $g(z) = z^*$ for all $z \in \mathcal{Z}$. Then

$$\begin{aligned} \alpha(\mathcal{C}; \nu) &= \max_{v \in (\mathbb{R}^d)^{\mathcal{Z}}} \max_{x, y \in \mathcal{P}(\mathcal{C})} \min_{u \in \mathcal{P}(\mathcal{C})} \min_{g: \mathcal{Z} \rightarrow \mathcal{Z}} \mathbb{E}_\nu \left[\langle v_z, x_z - y_z \rangle^{1/2} \right]^2 \mathbb{E}_\nu \left[\frac{\|A_{u_z}^z v_{g(z)}\|^2}{\langle v_{g(z)}, x_{g(z)} - y_{g(z)} \rangle} \right]^{-1} \\ &\leq \max_{v \in (\mathbb{R}^d)^{\mathcal{Z}}} \max_{x, y \in \mathcal{P}(\mathcal{C})} \min_{u \in \mathcal{P}(\mathcal{C})} \langle v_{z^*}, x_{z^*} - y_{z^*} \rangle^2 \mathbb{E}_\nu \left[\|A_{u_z}^z v_{z^*}\|^2 \right]^{-1} \\ &\leq \max_{v \in (\mathbb{R}^d)^{\mathcal{Z}}} \max_{x, y \in \mathcal{P}(\mathcal{C})} \min_{z' \in \mathcal{Z}} \min_{u \in \mathcal{P}(\mathcal{C}; z')} \langle v_{z^*}, x_{z^*} - y_{z^*} \rangle^2 \left(\nu(z') \|A_{u'}^{z'} v_{z^*}\|^2 \right)^{-1}, \end{aligned}$$

which proves the claim. We first used the definition of z^* and lower-bounded the expectation in the last step. The last equality iii) is immediate. \blacksquare

Our next result extends Lemma 15 to account for the contextual distribution ν in the information ratio. We provide the proof for the tighter information gain $I_t(x, z; w)$ as defined in (9). The information processing inequality (Lemma 14) implies the result for $I_t(x, z)$.

Lemma 22 For a convex set $\mathcal{C} \subset \mathbb{R}^d$ let $w \in \{x - y : x, y \in \mathcal{P}(\mathcal{C})\}$ be a difference in the plausible action set $\mathcal{P}(\mathcal{C}) = \times_{z \in \mathcal{Z}} \mathcal{P}(\mathcal{C}; z)$. Then

$$\mathbb{E}_\nu [\|w\|_{V_{t-1}^{-1}}]^2 \leq \alpha(\mathcal{C}, \nu) \max_{u \in \mathcal{P}(\mathcal{C})} \max_{g: \mathcal{Z} \rightarrow \mathcal{Z}} 2I_t(u, \nu; w_g) \leq \alpha(\mathcal{C}, \nu) \max_{u \in \mathcal{P}(\mathcal{C})} 2I_t(u, \nu).$$

Proof The proof is along the lines of Lemma 15, but keeps the expectation over ν . Let $g : \mathcal{Z} \rightarrow \mathcal{Z}$ be any function and $u \in \mathcal{X}$. From the proof of the mentioned lemma, we find

$$\frac{\|A_{u_z}^{\top} V_{t-1}^{-1} w_{g(z)}\|^2}{\|w_{g(z)}\|_{V_{t-1}^{-1}}} \leq 2I_t(u, z; w_{g(z)}).$$

Therefore, in expectation

$$\mathbb{E}_{\nu} \left[\frac{\|A_u^{\top} V_{t-1}^{-1} w_g\|^2}{\|w_g\|_{V_{t-1}^{-1}}} \right] \leq 2I_t(u, \nu; w_g).$$

Let $u^*, g^* = \arg \max_{u \in \mathcal{P}(\mathcal{C}), g: \mathcal{Z} \rightarrow \mathcal{Z}} I_t(u, \nu; w_g)$. With this we find

$$\begin{aligned} \frac{\mathbb{E}_{\nu} [\|w\|_{V_{t-1}^{-1}}]^2}{I_t(u^*, \nu; w_{g^*})} &\leq 2 \min_{u \in \mathcal{P}(\mathcal{C})} \min_{g: \mathcal{Z} \rightarrow \mathcal{Z}} \mathbb{E}_{\nu} [\|w\|_{V_{t-1}^{-1}}]^2 \mathbb{E}_{\nu} \left[\frac{\|A_u^{\top} V_{t-1}^{-1} w_g\|^2}{\|w_g\|_{V_{t-1}^{-1}}} \right]^{-1} \\ &\leq 2 \max_{v \in \mathcal{X}^{\mathcal{Z}}} \max_{x, y \in \mathcal{P}(\mathcal{C})} \min_{u \in \mathcal{P}(\mathcal{C})} \min_{g: \mathcal{Z} \rightarrow \mathcal{Z}} \mathbb{E}_{\nu} [\langle v, x - y \rangle^{1/2}]^2 \mathbb{E}_{\nu} \left[\frac{\|A_u^{\top} v_g\|^2}{\langle v_g, x_g - y_g \rangle} \right]^{-1} \\ &= 2\alpha(\mathcal{C}, \nu) \end{aligned}$$

Rearranging completes the proof. ■

With these results the regret bounds for contextual IDS follow. For simplicity, the proof is given for IDS with full information gain (4), but similar results can be obtained for the directed information gain. First, the globally observable case (Theorem 10, Section 2.5).

Proof [of Theorem 10] Let $x_t^* \in \mathcal{X}$ be the greedy action for each context, $x_t^*(z) = \arg \max_{x \in \mathcal{X}_z} x^{\top} \hat{\theta}_t$. Define the least accurate direction $w_t \in \{x - y : x, y \in \mathcal{X}\}$ in the set \mathcal{P}_t as

$$w_t(z) = \arg \max_{w=x-y: x, y \in \mathcal{P}_t(z)} \|w\|_{V_{t-1}^{-1}}^2. \quad (15)$$

Recall that $x_t^*(z)^{\top} \hat{\theta}_t \leq \beta_{t-1}^{1/2} \|w_t(z)\|_{V_{t-1}^{-1}}$. Lemma 22 implies

$$\Delta_t(x_t^*; \nu)^2 \leq \beta_{t-1} \mathbb{E}_{\nu} [\|w_t\|_{V_{t-1}^{-1}}]^2 \leq 2\beta_{t-1} \alpha(\nu) \max_{u \in \mathcal{X}} I_t(u; \nu).$$

The rest of argument is analogous to the proof of Lemma 4. Consider a sampling distribution $\mu(p) \in \mathcal{P}(\mathcal{X})$ that chooses $x_t^*(z)$ with probability $(1-p)$, and the most informative action $u(z) = \arg \max_{u \in \mathcal{X}_z} I_t(u, z)$ in context z with probability p . By definition of the IDS policy,

$$\begin{aligned} \Psi_t(\xi_t) &\leq \min_{p \in [0,1]} \frac{\Delta_t(\mu(p); \nu)^2}{I_t(\mu(p); \nu)} \\ &\leq \min_{p \in [0,1]} \frac{((1-p)\Delta_t(x_t^*; \nu) + p)^2}{pI_t(\mu(p); \nu)} \\ &\leq 2\alpha\beta_{t-1} \min_{p \in [0,1]} \frac{((1-p)\Delta_t(x_t^*; \nu) + p)^2}{p\Delta_t(x_t^*; \nu)} \\ &\leq \frac{4\alpha\beta_{t-1}}{\Delta_t(x_t^*; \nu)^2} \end{aligned}$$

We first used that $I_t(x_t^*) \geq 0$, $\Delta_t(u) \leq 1$ and the inequality $2\alpha\beta_{t-1} \max_u I_t(u; \nu) \geq \Delta_t(x_t^*; \nu)^2$ that we derived above. Then we optimized over $p \in [0, 1]$ in the last step. Similar to Lemma 5, one can show that IDS plays greedy most of the time. For any $x \in \mathcal{X}$,

$$\Delta_t(\mu_t; \nu) \leq 2\Delta(x; \nu).$$

With this we find

$$\Psi_t(\xi_t) \leq \frac{4\alpha\beta_{t-1}}{\Delta_t(x_t^*; \nu)} \leq \frac{8\alpha\beta_{t-1}}{\Delta_t(\mu_t; \nu)}.$$

Invoking the general bound (Lemma 1) and balancing the terms completes the proof. \blacksquare

E.3. Locally Observable Contextual Games

The condition for locally observable games is that any difference in the plausible action set $\mathcal{P}(\mathcal{C}, z)$ for a context $z \in \mathcal{Z}$ can be estimated under possibly different context $z' \in \mathcal{Z}$ by playing only actions that appear plausible optimal in the context z' . Formally,

$$\forall \mathcal{C} \subset \mathbb{R}^d \text{ convex}, z \in \mathcal{Z}, x, y \in \mathcal{P}_t(\mathcal{C}; z) \Rightarrow \exists z' \in \mathcal{Z} \text{ s.t. } x - y \in \text{span}(A_x : x \in \mathcal{P}_t(\mathcal{C}; z')).$$

If the condition holds true, Lemma 21 implies that $\alpha(\mathcal{C}, \nu) \leq \mathcal{O}((\min_{z \in \mathcal{Z}} \nu(z))^{-1})$ for finite action sets.

Theorem 23 *In contextual games that are uniformly locally observable in the sense that $\alpha(\mathcal{C}, \nu) \leq \alpha_0$ for any convex set $\mathcal{C} \subset \mathbb{R}^d$, the regret is bounded for any $n \geq 1$ with probability at least $1 - \delta$,*

$$R_n \leq C\sqrt{\alpha_0\beta_n n(\gamma_n + \log \frac{1}{\delta})} + 4 \log \left(\frac{4n + 4}{\delta} \right).$$

Proof Let w_t be the least accurate direction in the current set of plausible maximisers defined in Eq. (15). For any plausible maximiser $x \in \mathcal{P}_t$ it holds that

$$\Delta_t(x, c) \leq 2\beta_{t-1}^{1/2} \|w_t(c)\|_{V_{t-1}^{-1}},$$

Therefore $\Delta_t(x; \nu)^2 \leq 4\beta_{t-1} \mathbb{E}_\nu[\|w_t\|_{V_{t-1}^{-1}}]^2$. By Lemma 22,

$$\Delta_t(x; \nu)^2 \leq 8\beta_{t-1} \alpha(\mathcal{C}_t; \nu) \max_{z \in \mathcal{P}_t} I_t(z, \nu).$$

Finally, let $z_t = \arg \max_{x \in \mathcal{P}_t} I_t(z, \nu)$ be the most informative action that appears plausible optimal for each context. This action has bounded information ratio:

$$\Psi_t(\xi) \leq \frac{\Delta(z_t, \nu)^2}{I_t(z_t, \nu)} \leq 8\alpha_0\beta_{t-1},$$

where we also used that $\alpha(\mathcal{C}_t, \nu) \leq \alpha_0$ by assumption. The result follows from Lemma 1. \blacksquare

Appendix F. Proof of the Classification Theorem

The classification theorem is proven by combining upper and lower bounds, carefully checking that all cases have been covered. To begin, we introduce the classification of actions that is now standard in finite partial monitoring games. The lower bounds then follow using standard techniques and are sketched in Appendix G.

Notation We denote by $\text{relint}(C) = \{x \in C : \forall y \in C \exists \lambda > 1 : \lambda x + (1 - \lambda)y \in C\}$ the relative interior of a convex set C , and the set of extreme points by $\text{ext}(C)$. The closure operator on subsets of a metric space is $\text{cl}(\cdot)$ and $\text{dim}(\cdot)$ is the Hausdorff dimension. For points $x, y \in \mathbb{R}^d$, let $[x, y] = \{tx + (1 - t)y : t \in [0, 1]\}$.

Assumption 1 *For the remainder of this section we assume that \mathcal{X} is finite.*

The set of *Pareto optimal* actions is the set of extreme points $\text{ext}(\text{conv}(\mathcal{X}))$ of the convex hull of \mathcal{X} . An action is *degenerate* if it is on the boundary of $\text{conv}(\mathcal{X})$, but not an extreme point. Actions in the interior of $\text{conv}(\mathcal{X})$ are called *dominated*. The situation is illustrated in Figure 2. Finite partial monitoring games can be completely classified by considering a graph structure known as the *neighborhood graph* (Lattimore and Szepesvári, 2019a). Given an action $x \in \mathcal{X}$, the *cell* of x is the set of parameters for which action x is optimal:

$$C_x = \{\theta \in \mathbb{R}^d : x \in \mathcal{P}(\theta)\} = \{\theta \in \mathbb{R}^d : \max_{y \in \mathcal{X}} \langle y - x, \theta \rangle = 0\}.$$

Since \mathcal{X} is finite, $\text{conv}(\mathcal{X})$ is a polytope and C_x is either the singleton $\{0\}$ or an unbounded polytope. An action x is Pareto optimal if $\text{dim}(C_x) = d$ and degenerate otherwise, which can be seen by observing that C_x is the normal cone of x with respect to the convex body $\text{conv}(\mathcal{X})$. Pareto optimal actions x and y are called neighbours if $\text{dim}(C_x \cap C_y) = d - 1$, where the dimension of a polytope is defined as the dimension of the smallest affine space containing it. The neighbourhood relation defines a connected graph on the set of Pareto optimal actions. For neighboring actions x and y let $\mathcal{N}_{xy} = \{z : C_x \cap C_y \subseteq C_z\}$. Note that, besides x and y , \mathcal{N}_{xy} contains only actions z with $\text{dim}(C_z) = d - 1$. Lin et al. (2014) and Chaudhuri and Tewari (2016) use a different notion to ensure global observability and to construct an explicit exploration distribution. A *global observer set* is a set of actions $\mathcal{Y} \subset \mathcal{X}$ such that $\text{span}(A_x : x \in \mathcal{Y}) = \text{span}(x - y : x, y \in \mathcal{X})$.

Lemma 24 *The following conditions equivalently characterize globally observable games:*

- i) *For all actions $x, y \in \mathcal{X}$ it holds $x - y \in \text{span}\{A_z : z \in \mathcal{X}\}$.*
- ii) *For all Pareto optimal actions x, y , it holds $x - y \in \text{span}\{A_z : z \in \mathcal{X}\}$.*
- iii) *There exists a global observer set.*

Proof For the implication (ii \Rightarrow i), note that Pareto optimal actions are the extreme points of $\text{conv}(\mathcal{X})$, therefore any $x \in \mathcal{X}$ can be written as a convex combination of Pareto optimal actions. (i \Rightarrow iii) follows by taking \mathcal{X} as global observer set. (iii \Rightarrow ii) immediately follows from the definition of a global observer set. ■

The next lemma shows the relation of neighboring actions and local observability.




Figure 2: Green vectors are Pareto optimal, blue ones are degenerated and red are dominated. The light blue cones are the normal cones associated with each Pareto optimal action indicating the direction of θ for which that action is optimal.

Lemma 25 *Let \mathcal{X} be finite and $\mathcal{C} \subset \mathbb{R}^d$ be any convex set. Then*

- i) *The Pareto optimal actions within $\mathcal{P}(\mathcal{C})$ are connected on the neighborhood graph.*
- ii) *For two Pareto optimal actions $x, y \in \mathcal{P}(\mathcal{C})$ it holds that $\mathcal{N}_{xy} \subset \mathcal{P}(\mathcal{C})$.*
- iii) *Any $x \in \mathcal{P}(\mathcal{C})$ can be written as convex combination of Pareto optimal actions in $\mathcal{P}(\mathcal{C})$.*

Proof

- i) The proof is intuitively simple. Take any Pareto optimal actions $x, y \in \mathcal{P}(\mathcal{C})$ and let $\theta_x \in C_x$ and $\theta_y \in C_y$. Then take the chord $[\theta_x, \theta_y] \subset \mathcal{C}$ and consider the path $(x_i)_{i=1}^n$ defined by the cells that intersect $[\theta_x, \theta_y] \cup C_{x_i} \neq \emptyset$. There is a technicality that this chord may pass through intersections of cells that have dimension $d - 2$. A perturbation and dimension argument fixes the proof (see Lemma 28 below). For a similar result see (Lattimore and Szepesvári, 2019c, Lemma 23).
- ii) Let $x, y \in \mathcal{P}(\mathcal{C})$ be Pareto optimal actions. Pick any $\theta \in C_x \cap C_y \cap \mathcal{C}$. If $z \in \mathcal{N}_{xy}$, we have $C_x \cap C_y \subset C_z$, hence $\theta \in C_z$ and z is optimal for θ . Therefore $z \in \mathcal{P}(\mathcal{C})$.
- iii) Let F be the lowest dimensional face of $\text{conv}(\mathcal{X})$ containing $x \in \mathcal{P}(\mathcal{C})$. Assume that x is in the interior of F (otherwise it would be an extreme point and so Pareto optimal). Then let θ be a parameter such that x is optimal. $H = \{z : (x - z)^\top \theta = 0\}$ is a supporting hyperplane of $\text{conv}(F)$. Hence F is a subset of H . Note that $H \cap \mathcal{X} \subset \mathcal{P}(\theta)$ contains actions that are optimal for θ . Therefore all extreme points of F are in $\mathcal{P}(\mathcal{C})$ and since x is in the convex hull of the extreme points of F the result follows. ■

The next lemma shows that observability can be characterized in terms of the neighborhood relation.

Lemma 26 *The following conditions equivalently characterize locally observable games:*

- i) *For any convex \mathcal{C} and all $x, y \in \mathcal{P}(\mathcal{C})$, $x - y \in \text{span}\{A_z : z \in \mathcal{P}(\mathcal{C})\}$.*
- ii) *For any two neighboring Pareto optimal actions x, y , $x - y \in \text{span}\{A_z : z \in \mathcal{N}_{xy}\}$.*

Proof “i) \Rightarrow ii)”. Let $x, y \in \mathcal{X}$ be neighboring Pareto optimal actions. Pick $\theta \in \text{relint}(C_x \cap C_y)$. Then $\mathcal{N}_{xy} = \mathcal{P}(\theta)$ by Lemma 25 and therefore $x - y \in \text{span}\{A_x : x \in \mathcal{P}(\mathcal{C})\}$ by i).

“ii) \Rightarrow i)”. Let $x, y \in \mathcal{P}(\mathcal{C})$. First note that by Lemma 25, iii), $x - y$ can be written as linear combination of Pareto optimal actions in $\mathcal{P}(\mathcal{C})$. Therefore we can assume that x, y are Pareto optimal. By Lemma 25, i), there exists a sequence $(x_i)_{i=1}^m$ of Pareto optimal actions with $x_1 = x$ and $x_m = y$, such that x_i, x_{i+1} are neighbors and $\{x_i : i = 1, \dots, m\} \subset \mathcal{P}(\mathcal{C})$. By assumption, $x_i - x_{i-1} \in \text{span}\{A_z : z \in \mathcal{N}_{x_i x_{i-1}}\}$. Since $x - y = \sum_{i=1}^{m-1} x_i - x_{i+1}$ the claim follows. \blacksquare

Lastly, the key lemma for proving the lower bound for globally observable games shows that in games that are not locally observable, there exists a pair of neighbouring Pareto optimal actions x, y and a parameter θ such that both actions x, y are optimal, but $\langle x - y, \theta \rangle$ can *not* be estimated by playing only actions from the neighborhood \mathcal{N}_{xy} .

Lemma 27 *Suppose a game is not locally observable. Then there exists a pair x, y of neighbouring Pareto optimal actions and $\theta \in \text{relint}(C_x \cap C_y)$ such that $x - y \notin \text{span}\{A_z : z \in \mathcal{N}_{xy}\}$.*

Proof The lemma follows from the definition of local observability its equivalent characterization provided in Lemma 26. \blacksquare

Lemma 28 *Let \mathcal{U} be a collection of disjoint open sets of \mathbb{R}^d with the usual metric $d(x, y) = \|x - y\|$ such that:*

1. *The union $\mathcal{K} = \bigcup_{i=1}^{\infty} \text{cl}(U_i)$ is convex.*
2. *$\dim(\mathcal{K} \cap \{x : d(x, y) \leq \epsilon\}) = \dim(\mathcal{K})$ for all $\epsilon > 0$ and $y \in U \in \mathcal{U}$.*
3. *$A = \bigcup_{U, V, W \in \mathcal{U} \text{ are distinct}} (\text{cl}(U) \cap \text{cl}(V) \cap \text{cl}(W))$ has $\dim(A) < \dim(\mathcal{K}) - 1$.*
4. *For any compact set $W \subset \mathbb{R}^d$, at most finitely many elements of \mathcal{U} have non-empty intersection with W .*

We say that $U_1, U_2 \in \mathcal{U}$ are connected if $(\text{cl}(U_1) \cap \text{cl}(U_2)) \setminus A \neq \emptyset$. Then, for any $x \in \text{cl}(U)$ and $y \in \text{cl}(V)$, there exists a sequence $(U_i)_{i=1}^m$ of connected sets with $U_1 = U$ and $U_m = V$ and $U_i \cap [x, y] \neq \emptyset$.

Proof Suppose that $x \in \text{cl}(U)$ and $y \in V$. Let $B = \{z \in \mathcal{K} : d(z, y) < \epsilon\}$ with ϵ sufficiently small that $B \subset V$ and $S = \{z \in B : [z, x] \cap A \neq \emptyset\}$. A straightforward calculation shows that $\dim(S) \leq \dim(A) + 1 < \dim(\mathcal{K}) = \dim(B)$. Hence, there exists a $z \in B \setminus S$ and $[x, z] \cap A = \emptyset$. By the last assumption, $[x, z]$ intersects with at most finitely many elements of \mathcal{U} , which form the path between U and V . Next, suppose that $y \in \text{cl}(V)$ and let (y_n) be a sequence in V with $\lim_{n \rightarrow \infty} d(y_n, y) = 0$. By the previous argument, for each n , there exists a sequence $(U_i^n)_{i=1}^{m_n}$ of connected sets with $U_1^n = U$ and $U_{m_n}^n = V$. By the fourth assumption, for suitably large n , there

are only finitely many sets in all the $(U_i^n)_{i,n}$ and hence, by re-labelling if necessary, the sequence of connected sets can be chosen so that $(U_i^n)_{i=1}^{m_n}$ converges (in the sense that the identity/order of the sequences converges – the discrete topology on finite sequences of \mathcal{U}) to some sequence $(U_i)_{i=1}^m$. We need to show that $\text{cl}(U_i) \cap [x, y] \neq \emptyset$ for each i . By the definition of convergence we have $[x, y_n] \cap \text{cl}(U_i) \neq \emptyset$ for all suitably large n . Taking a sequence (z_n) with $z_n \in U_i$ for all suitably large n . Compactness again allows us to assume that (z_n) converges to some z , which is easily seen to lie on $[x, z]$ and by closure of U_i also lies in U_i , as required. ■

Appendix G. Lower Bounds

The lower bounds complete the classification theorem. These results are almost implied by existing theorems from finite partial monitoring. The only difference is that here the outcome space is infinite, which does not change the structure of the proofs. We include here the key details and intuition. As expected, the key tool is Le Cam’s method in combination with the Bretagnolle–Huber inequality (Bretagnolle and Huber, 1979) and an elementary calculation of the relative entropy between measures on interaction sequences induced by a fixed policy and for different environments. For the remainder of this section, we fix an arbitrary policy and finite game with actions \mathcal{X} and feedback functions $(A_x)_{x \in \mathcal{X}}$. For simplicity, we assume the noise is Gaussian and \mathcal{X} spans \mathbb{R}^d . Given a $\theta \in \mathbb{R}^d$ let \mathbb{P}_θ^n be the measure on action/observation sequences of length n when the learner interacts with the game for parameter θ . Before the theorems and proofs we need a little more notation. Let

$$V_n(\theta) = \mathbb{E}_\theta[V_n] = \mathbb{E}_\theta \left[\sum_{t=1}^n A_{x_t} A_{x_t}^\top \right].$$

Then define $E_n(\theta)$ as the binary random variable that the algorithm plays a suboptimal action at least $n/2$ times.

$$E_n(\theta) = \mathbb{1} \left(\sum_{t=1}^n \mathbb{1}(x_t \notin \mathcal{P}(\theta)) \geq n/2 \right).$$

Notice that if θ, θ' are such that $\mathcal{P}(\theta) \cap \mathcal{P}(\theta') = \emptyset$, then $E_n(\theta') \geq 1 - E_n(\theta)$. For simplicity we focus on proving lower bounds on the expected regret. The extension to high probability bounds is possible using the techniques of Gerchinovitz and Lattimore (2016). Let

$$R_n(\theta) = \mathbb{E}_\theta[R_n]$$

be the expected regret when the learner interacts with the environment determined by θ .

Lemma 29 *The relative entropy between \mathbb{P}_θ^n and $\mathbb{P}_{\theta'}^n$ satisfies $\text{KL}(\mathbb{P}_\theta^n, \mathbb{P}_{\theta'}^n) = \frac{1}{2} \|\theta - \theta'\|_{V_n(\theta)}^2$.*

For a proof refer to (Lattimore and Szepesvári, 2018, Theorem 24.1).

Lemma 30 (Bretagnolle-Huber inequality) *Let P and Q be probability measures on the same measurable space (F, Ω) and let $A \in F$ be an arbitrary event. Then*

$$P(A) + Q(A^c) \geq \frac{1}{2} \exp(-\text{KL}(P, Q)). \quad (16)$$

Theorem 31 *Suppose that $\text{span}(A_x : x \in \mathcal{X}) \neq \mathbb{R}^d$, then there exists a game-dependent constant $c > 0$ such that for all $n \geq 1$ there exists a θ for which $R_n(\theta) \geq cn$.*

Proof Let $\theta \in \mathbb{R}^d$ be a non-zero vector such that $A_x \theta = 0$ for all $x \in \mathcal{X}$, which exists by the assumption that $\text{span}(A_x : x \in \mathcal{X}) \neq \mathbb{R}^d$. Next, let $\theta' = -\theta$ and notice that by Lemma 29,

$$\text{KL}(\mathbb{P}_\theta^n, \mathbb{P}_{\theta'}^n) = 0.$$

By our choice, the optimal action for the environment determined by θ and θ' are different: $\mathcal{P}(\theta) \cap \mathcal{P}(\theta') = \emptyset$. The Bretagnolle-Huber inequality (Lemma 30) implies that

$$\mathbb{P}_\theta^n(E_n(\theta)) + \mathbb{P}_{\theta'}^n(E_n(\theta')) \geq \mathbb{P}_\theta^n(E_n(\theta)) + \mathbb{P}_{\theta'}^n(1 - E_n(\theta)) \geq \frac{1}{2} \exp(-\text{KL}(\mathbb{P}_\theta^n, \mathbb{P}_{\theta'}^n)) \geq \frac{1}{2}. \quad (17)$$

Furthermore, there exists an $\epsilon > 0$ such that $R_n(\theta) \geq \epsilon n \mathbb{P}_\theta^n(E_n(\theta))/2$. Hence, by Eq. (17), the regret is linear for either environment θ or θ' . \blacksquare

Theorem 32 *Suppose the game is globally observable, but not locally observable. Then there exists a game-dependent constant $c > 0$ and $\theta \in \mathbb{R}^d$ such that the regret is $R_n(\theta) \geq cn^{2/3}$.*

Proof By Lemma 27, there exists a pair of neighboring Pareto optimal actions $x, y \in \text{ext}(\text{conv}(\mathcal{X}))$ and $\theta \in \text{relint}(C_x \cap C_y)$ such that $x - y \notin \text{span}(\{A_z : z \in \mathcal{P}(\theta)\}) = L$. Let $x - y = u + v$, where $u \in L$ and $v \in L^\perp$. Since $x - y \notin L$ it follows that

$$\langle x - y, v \rangle = \langle u + v, v \rangle = \|v\|^2 > 0.$$

In particular, for suitably small $\epsilon > 0$ it holds that $\theta + \epsilon v \in C_x$ and $\theta - \epsilon v \in C_y$. Define

$$\theta_n = \theta + n^{-1/3}v \quad \text{and} \quad \theta'_n = \theta - n^{-1/3}v$$

and let assume n is sufficiently large that $\theta_n \in C_x$ and $\theta'_n \in C_y$. Next, decompose $V_n(\theta)$ as $V_n(\theta) = U_n(\theta) + W_n(\theta)$, where

$$U_n(\theta) = \mathbb{E}_\theta \left[\sum_{t=1}^n \mathbb{1}(x_t \in \mathcal{P}(\theta)) A_{x_t} A_{x_t}^\top \right] \quad \text{and} \quad W_n(\theta) = \mathbb{E}_\theta \left[\sum_{t=1}^n \mathbb{1}(x_t \notin \mathcal{P}(\theta)) A_{x_t} A_{x_t}^\top \right].$$

Let $T_n(\mathcal{Y}) = \sum_{t=1}^n \mathbb{1}(x_t \in \mathcal{Y})$ be the number of times an action in $\mathcal{Y} \subset \mathcal{X}$ is played. Notice, since $v \in L^\perp$, that

$$\frac{1}{2} \|\theta_n - \theta'_n\|_{V_n(\theta_n)}^2 = 2n^{-2/3} \|v\|_{V_n(\theta_n)}^2 = 2n^{-2/3} \|v\|_{W_n(\theta_n)}^2 \leq 2n^{-2/3} \mathbb{E}_{\theta_n} [T_n(\mathcal{P}(\theta)^c)] \|v\|_G^2,$$

where $G = \sum_{x \in \mathcal{X}} A_x A_x^\top$. Now, there exists a game-dependent constant $\epsilon > 0$ such that

$$R_n(\theta_n) \geq \epsilon \mathbb{E}_{\theta_n} [T_n(\mathcal{P}(\theta)^c)].$$

Hence if $\mathbb{E}_{\theta_n} [T_n(\mathcal{P}(\theta)^c)] \geq n^{2/3}$, then $R_n(\theta_n) \geq \epsilon n^{2/3}$. Assume that $\mathbb{E}_{\theta_n} [T_n(\mathcal{P}(\theta)^c)] \leq n^{2/3}$. By the Bretagnolle-Huber inequality (Lemma 30), there exists another game-dependent constant $\epsilon' > 0$ such that

$$R_n(\theta_n) + R_n(\theta'_n) \geq n^{2/3} \epsilon' \exp\left(-2n^{-2/3} \mathbb{E}_{\theta_n} [T_n(\mathcal{P}(\theta)^c)] \|v\|_G^2\right) \geq n^{2/3} \epsilon' \exp(-2 \|v\|_G^2).$$

Combining the last two displays completes the proof. \blacksquare

Theorem 33 *Suppose the game is locally observable, then there exists a constant $c > 0$ such that for all n there is a θ for which $R_n(\theta) \geq cn^{1/2}$.*

Proof Let $\theta \in \mathbb{R}^d$ be arbitrary and $\theta_n = n^{-1/2}\theta$ and $\theta'_n = -\theta_n$. By the assumption that \mathcal{X} spans \mathbb{R}^d , it follows that $\mathcal{P}(\theta_n) \cap \mathcal{P}(\theta'_n) = \emptyset$. By Lemma 29,

$$\text{KL}(\mathbb{P}_{\theta_n}^n, \mathbb{P}_{\theta'_n}^n) = \frac{1}{2} \|\theta'_n - \theta_n\|_{V_n(\theta_n)}^2 = \frac{1}{2} \|\theta\|_{V_n(\theta_n)/n}^2.$$

Clearly, $G = \sum_{x \in \mathcal{X}} A_x A_x^\top \succ V_n(\theta_n)/n$. Hence, there exists a constant $c > 0$ such that for all $n \geq 1$,

$$\text{KL}(\mathbb{P}_{\theta_n}^n, \mathbb{P}_{\theta'_n}^n) \leq c.$$

Then, using the same argument as in the proof of Theorem 31, we have

$$\mathbb{P}_{\theta_n}^n(E_n(\theta_n)) + \mathbb{P}_{\theta'_n}^n(E_n(\theta'_n)) \geq \mathbb{P}_{\theta_n}^n(E_n(\theta_n)) + \mathbb{P}_{\theta'_n}^n(1 - E_n(\theta_n)) \geq \frac{1}{2} \exp(-c).$$

The result follows because there exists an $\epsilon > 0$ such that $R_n(\theta) \geq \mathbb{P}_{\theta}^n(E_n(\theta))\epsilon\sqrt{n}/2$. ■