

# Thompson Sampling for Loosely-Coupled Multi-Agent Systems: An Application to Wind Farm Control

Timothy Verstraeten  
Artificial Intelligence Lab Brussels,  
Vrije Universiteit Brussel  
tiverstr@vub.be

Eugenio Bargiacchi  
Artificial Intelligence Lab Brussels,  
Vrije Universiteit Brussel  
ebargiac@vub.be

Pieter J.K. Libin  
Artificial Intelligence Lab Brussels,  
Vrije Universiteit Brussel  
plibin@vub.be

Jan Helsen  
Acoustics & Vibrations Research  
Group, Vrije Universiteit Brussel  
jahelsen@vub.be

Diederik M. Roijers  
Microsystems Technology, HU  
University of Applied Sciences  
Utrecht  
diederik.yamamoto-roijers@hu.nl

Ann Nowé  
Artificial Intelligence Lab Brussels,  
Vrije Universiteit Brussel  
anowe@vub.be

## ABSTRACT

Multi-agent coordination is prevalent in many real-world applications. However, such coordination is challenging due to its combinatorial nature. An important observation in this regard is that agents in the real world often only directly affect a limited set of neighbouring agents. Leveraging such loose couplings among agents is key to making coordination in multi-agent systems feasible. In this work, we focus on *learning* to coordinate. Specifically, we consider the multi-agent multi-armed bandit framework, in which fully cooperative loosely-coupled agents must learn to coordinate their decisions to optimize a common objective. We propose multi-agent Thompson sampling (MATS), a new Bayesian exploration-exploitation algorithm that leverages loose couplings. We provide a regret bound that is sublinear in time and low-order polynomial in the highest number of actions of a single agent for sparse coordination graphs. Additionally, we empirically show that MATS outperforms the state-of-the-art algorithm, MAUCE, on two synthetic benchmarks, and a novel benchmark with Poisson distributions. An example of a loosely-coupled multi-agent system is a wind farm. Coordination within the wind farm is necessary to maximize power production. As upstream wind turbines only affect nearby downstream turbines, we can use MATS to efficiently learn the optimal control mechanism for the farm. To demonstrate the benefits of our method toward applications we apply MATS to a realistic wind farm control task. In this task, wind turbines must coordinate their alignments with respect to the incoming wind vector in order to optimize power production. Our results show that MATS improves significantly upon state-of-the-art coordination methods in terms of performance, demonstrating the value of using MATS in practical applications with sparse neighbourhood structures.

## KEYWORDS

multi-agent; multi-armed bandits; Thompson sampling; coordination graphs

## 1 INTRODUCTION

Multi-agent decision coordination is prevalent in many real-world applications, such as traffic light control [39], warehouse commissioning [11] and wind farm control [14, 37]. Often, such settings

can be formulated as coordination problems in which agents have to cooperate in order to optimize a shared team reward [5].

Handling multi-agent settings is challenging, as the size of the joint action space scales exponentially with the number of agents in the system. Therefore, an approach that directly considers all agents' actions jointly is computationally intractable. This has made such coordination problems the central focus in the planning literature [15–17, 22]. Fortunately, in real-world settings agents often only directly affect a limited set of neighbouring agents. This means that the global reward received by all agents can be decomposed into local components that only depend on small subsets of agents. Exploiting such loose couplings is key in order to keep multi-agent decision problems tractable [9].

In this work, we consider learning to coordinate in multi-agent systems. For example, consider a wind farm control task, which is comprised of a set of wind turbines, and we aim to maximize the farm's total productivity. When upstream turbines directly face the incoming wind stream, energy is extracted from wind. This reduces the productivity of downstream turbines, potentially damaging the overall power production. However, turbines have the option to rotate, in order to deflect the turbulent flow away from turbines downwind [35]. Due to the complex nature of the aerodynamic interactions between the turbines, constructing a model of the environment and deriving a control policy using planning techniques is extremely challenging [27]. Instead, a joint control policy among the turbines can be *learned* to effectively maximize the productivity of the wind farm. The system is loosely coupled, as redirection only directly affects adjacent turbines.

While most of the literature only considers approximate reinforcement learning methods for learning in multi-agent systems, it has recently been shown [4] that it is possible to achieve theoretical bounds on the regret (i.e., how much reward is lost due to learning). In this work, we use the multi-agent multi-armed bandit problem definition, and improve upon the state of the art. Specifically, we propose the multi-agent Thompson sampling (MATS) algorithm, which exploits loosely-coupled interactions in multi-agent systems. The loose couplings are formalized as a *coordination graph*, which defines for each pair of agents whether their actions depend on each other. We assume the graph structure is known beforehand, which is the case in many real-world applications with sparse agent interactions (e.g., wind farm control).

Our method leverages the exploration-exploitation mechanism of Thompson sampling (TS). TS has been shown to be highly competitive to other popular methods, e.g., UCB [8]. Recently, theoretical guarantees on its regret have been established [1], which renders the method increasingly popular in the literature. Additionally, due to its Bayesian nature, problem-specific priors can be specified. We argue that this has strong relevance in many practical fields, such as advertisement selection [8] and influenza mitigation [24, 25].

We provide a finite-time Bayesian regret analysis and prove that the upper regret bound of MATS is low-order polynomial in the number of actions of a single agent for sparse coordination graphs (Corollary 4.6). This is a significant improvement over the exponential bound of classic TS, which is obtained when the coordination graph is ignored [1]. We show that MATS improves upon the state of the art in various synthetic settings. Finally, we demonstrate that MATS achieves high performance on a realistic wind farm control task, in which multiple wind turbines have to be jointly aligned to maximize the total power production.

We define the problem setting, the multi-armed multi-agent bandit, in Section 2 and describe our method, MATS, in Section 3. We provide a theoretical and empirical analysis of MATS in Sections 4 and 5, respectively. Then, we apply our method to a realistic wind farm control task in Section 6. Next, we discuss the results in Section 7. Finally, we compare with related work in Section 8 and conclude that MATS achieves state-of-the-art performance, both empirically and theoretically, in Section 9.

## 2 PROBLEM STATEMENT

In this work, we adopt the multi-agent multi-armed bandit (MAMAB) setting [4, 32]. A MAMAB is similar to the multi-armed bandit formalism [34], but considers multiple agents factored into groups. When the agents have pulled a joint arm, each group receives a reward. The goal shared by all agents is to maximize the total sum of rewards. Formally,

*Definition 2.1.* A multi-agent multi-armed bandit (MAMAB) is a tuple  $\langle \mathcal{D}, \mathcal{A}, f \rangle$  where

- $\mathcal{D}$  is the set of  $m$  enumerated agents. This set is factorized into  $\rho$ , possibly overlapping, subsets of agents  $\mathcal{D}^e$ .
- $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_m$  is the set of joint actions, or joint arms, which is the Cartesian product of the sets of actions  $\mathcal{A}_i$  for each of the  $m$  agents in  $\mathcal{D}$ . We denote  $\mathcal{A}^e$  as the set of local joint actions, or local arms, for the group  $\mathcal{D}^e$ .
- $f(\mathbf{a})$  is a stochastic function providing a global reward when a joint arm,  $\mathbf{a} \in \mathcal{A}$ , is pulled. The global reward function is decomposed into  $\rho$  noisy, observable and independent local reward functions, i.e.,  $f(\mathbf{a}) = \sum_{e=1}^{\rho} f^e(\mathbf{a}^e)$ . A local function  $f^e$  only depends on the local arm  $\mathbf{a}^e$  of the subset of agents in  $\mathcal{D}^e$ .

We denote the mean reward of a joint arm as  $\mu(\mathbf{a}) = \sum_{e=1}^{\rho} \mu^e(\mathbf{a}^e)$ . For simplicity, we refer to the  $i^{\text{th}}$  agent by its index  $i$ .

The dependencies between the local reward functions and the agents are described as a coordination graph [16].

*Definition 2.2.* A coordination graph is a bipartite graph  $G = \langle \mathcal{D}, \{f^e\}_{e=1}^{\rho}, E \rangle$ , whose nodes  $\mathcal{D}$  are agents and components of a

factored reward function  $f = \sum_{e=1}^{\rho} f^e$ , and an edge  $(i, f^e) \in E$  exists if and only if agent  $i$  influences component  $f^e$ .

The dependencies in a MAMAB can be described by setting  $E = \{(i, f^e) \mid i \in \mathcal{D}^e\}$ .

In this setting, the objective is to minimize the expected cumulative regret [2], which is the cost incurred when pulling a particular joint arm instead of the optimal one.

*Definition 2.3.* The expected cumulative regret of pulling a sequence of joint arms until time step  $T$  according to policy  $\pi$  is

$$\mathbb{E}[R(T, \pi)] \triangleq \mathbb{E}\left[\sum_{t=1}^T \Delta(\mathbf{a}_t) \mid \pi\right] \quad (1)$$

with

$$\begin{aligned} \Delta(\mathbf{a}_t) &\triangleq \mu(\mathbf{a}_*) - \mu(\mathbf{a}_t) \\ &= \sum_{e=1}^{\rho} \mu^e(\mathbf{a}_*^e) - \mu^e(\mathbf{a}_t^e), \end{aligned} \quad (2)$$

where  $\mathbf{a}_*$  is the optimal joint arm and  $\mathbf{a}_t$  is the joint arm pulled at time  $t$ . For the sake of brevity, we will omit  $\pi$  when the context is clear.

Cumulative regret can be minimized by using a policy that considers the full joint arm space, thereby ignoring loose couplings between agents. This leads to a combinatorial problem, as the joint arm space scales exponentially with the number of agents. Therefore, loose couplings need to be taken into account whenever possible.

## 3 MULTI-AGENT THOMPSON SAMPLING

We propose the multi-agent Thompson sampling (MATS) algorithm for decision making in loosely-coupled multi-agent multi-armed bandit problems. Consider a MAMAB with groups  $\mathcal{D}^e$  (Definition 2.1). The local means  $\mu^e(\mathbf{a}^e)$  are treated as unknown. According to the Bayesian formalism, we exert our beliefs over the local means  $\mu^e(\mathbf{a}^e)$  in the form of a prior,  $Q_{\mathbf{a}^e}^e(\cdot)$ . At each time step  $t$ , MATS draws a sample  $\mu_t^e(\mathbf{a}^e)$  from the posterior for each group and local arm given the history,  $\mathcal{H}_{t-1}$ , consisting of local actions and rewards associated with past pulls:

$$\begin{aligned} \mu_t^e(\mathbf{a}^e) &\sim Q_{\mathbf{a}^e}^e(\cdot \mid \mathcal{H}_{t-1}^e) \\ \mathcal{H}_{t-1}^e &\triangleq \{(\mathbf{a}_i^e, f_i^e(\mathbf{a}_i^e))\}_{i=1}^{t-1}. \end{aligned} \quad (3)$$

Note that during this step, MATS samples directly the posterior over the unknown local means, which implies that the sample  $\mu_t^e(\mathbf{a}^e)$  and the unknown mean  $\mu^e(\mathbf{a}^e)$  are independent and identically distributed at time step  $t$ .

Thompson sampling (TS) chooses the arm with the highest sample, i.e.,

$$\mathbf{a}_t = \arg \max_{\mathbf{a}} \mu_t(\mathbf{a}). \quad (4)$$

However, in our case, the expected reward is decomposed into several local means. As conflicts between overlapping groups will arise, the optimal local arms for an agent in two groups may differ. Therefore, we must define the argmax-operator to deal with the

factored representation of a MAMAB, while still returning the full joint arm that maximizes the sum of samples, i.e.,

$$\mathbf{a}_t = \arg \max_{\mathbf{a}} \sum_{e=1}^{\rho} \mu_t^e(\mathbf{a}^e). \quad (5)$$

To this end, we use variable elimination (VE), which computes the joint arm that maximizes the global reward without explicitly enumerating over the full joint arm space [16]. Specifically, VE consecutively eliminates an agent from the coordination graph, while computing its best response with respect to its neighbours. VE is guaranteed to return the optimal joint arm and has a computational complexity that is combinatorial in terms of the induced width of the graph, i.e., the number of neighbours of an agent at the time of its elimination. However, as the method is typically applied to a loosely-coupled coordination graph, the induced width is generally much smaller than the size of the full joint action space, which renders the maximization problem tractable [16, 17]. Approximate efficient alternatives exist, such as max-plus [38], but using them will invalidate the proof for the Bayesian regret bound (Theorem 4.5).

Finally, the joint arm that maximizes Equation 5,  $\mathbf{a}_t$ , is pulled and a reward  $f_t^e(\mathbf{a}_t^e)$  will be obtained for each group. MATS is formally described in Algorithm 1.

**Data:** Prior  $Q_{\mathbf{a}^e}^e$  per group  $\mathcal{D}^e$  and local action  $\mathbf{a}^e$   
 $\mathcal{H}_0 \leftarrow \{\}$   
**for**  $t \in [1..T]$  **do**  
     $\forall e \in [1..\rho], \mathbf{a}^e \in \mathcal{A}^e :$   
         $\mu_t^e(\mathbf{a}^e) \sim Q_{\mathbf{a}^e}^e(\cdot \mid \mathcal{H}_{t-1})$   
         $\mathbf{a}_t \leftarrow \arg \max_{\mathbf{a}} \sum_{e=1}^{\rho} \mu_t^e(\mathbf{a}^e)$  using VE  
         $\langle f_t^e(\mathbf{a}_t^e) \rangle_{e=1}^{\rho} \leftarrow$  Pull joint arm  $\mathbf{a}_t$   
         $\mathcal{H}_t \leftarrow \mathcal{H}_{t-1} \cup \{ \langle \mathbf{a}_t^e, f_t^e(\mathbf{a}_t^e) \rangle_{e=1}^{\rho} \}$   
**end**

**Algorithm 1: MATS**

MATS belongs to the class of probability matching methods [23].

*Definition 3.1.* Given history  $\mathcal{H}_{t-1}$ , the probability distribution of the pulled arm  $\mathbf{a}_t$  is equal to the probability distribution of the optimal arm  $\mathbf{a}_*$ . Formally,

$$P(\mathbf{a}_t = \cdot \mid \mathcal{H}_{t-1}) = P(\mathbf{a}_* = \cdot \mid \mathcal{H}_{t-1}). \quad (6)$$

Intuitively, MATS samples the local mean rewards according to the beliefs of the user at each time step, and maximizes over those means to find the optimal joint arm according to Definition 2.1. This process is conceptually similar to traditional TS [34].

## 4 BAYESIAN REGRET ANALYSIS

Many multi-agent systems are composed of locally connected agents. When formalized as a MAMAB (Definition 2.1), our method is able to exploit these local structures during the decision process. We provide a regret bound for MATS that scales sublinearly with a factor  $\tilde{A}T$ , where  $\tilde{A}$  is the total number of local arms.

Consider a MAMAB  $(\mathcal{D}, \mathcal{A}, f)$  with  $\rho$  groups and the following assumption on the rewards:

**ASSUMPTION 4.1.** The global rewards have a mean between 0 and 1, i.e.,

$$\mu(\mathbf{a}) \in [0, 1], \forall \mathbf{a} \in \mathcal{A}.$$

**ASSUMPTION 4.2.** The local rewards shifted by their mean are  $\sigma$ -subgaussian distributed, i.e.,  $\forall e \in [1..\rho], \mathbf{a}^e \in \mathcal{A}^e$ ,

$$\mathbb{E} [\exp(t(f^e(\mathbf{a}^e) - \mu^e(\mathbf{a}^e)))] \leq \exp(0.5\sigma^2 t^2).$$

We maintain the pull counters  $n_{t-1}^e(\mathbf{a}^e)$  and estimated means  $\hat{\mu}_{t-1}^e(\mathbf{a}^e)$  for local arms  $\mathbf{a}^e$ .

Consider the event  $\mathcal{E}_T$ , which states that, until time step  $T$ , the differences between the local sample means and true means are bounded by a time-dependent threshold, i.e.,

$$\mathcal{E}_T \triangleq (\forall e, \mathbf{a}^e, t : |\hat{\mu}_{t-1}^e(\mathbf{a}^e) - \mu^e(\mathbf{a}^e)| \leq c_t^e(\mathbf{a}^e)) \quad (7)$$

with

$$c_t^e(\mathbf{a}^e) \triangleq \sqrt{\frac{2\sigma^2 \log(\delta^{-1})}{n_{t-1}^e(\mathbf{a}^e)}}. \quad (8)$$

where  $\delta$  is a free parameter that will be chosen later. We denote the complement of the event by  $\bar{\mathcal{E}}_T$ .

**LEMMA 4.3.** (Concentration inequality) The probability of exceeding the error bound on the local sample means is linearly bounded by  $\tilde{A}T\delta$ . Specifically,

$$P(\bar{\mathcal{E}}_T) \leq 2\tilde{A}T\delta. \quad (9)$$

**PROOF.** Using the union bound (U), we can bound the probability of observing event  $\bar{\mathcal{E}}_T$  as

$$\begin{aligned} P(\bar{\mathcal{E}}_T) &\stackrel{(7)}{=} P(\exists t, e, \mathbf{a}^e : |\hat{\mu}_{t-1}^e(\mathbf{a}^e) - \mu^e(\mathbf{a}^e)| > c_t^e(\mathbf{a}^e)) \\ &\stackrel{(U)}{\leq} \sum_{t=1}^T \sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} P(|\hat{\mu}_{t-1}^e(\mathbf{a}^e) - \mu^e(\mathbf{a}^e)| > c_t^e(\mathbf{a}^e)). \end{aligned} \quad (10)$$

The estimated mean  $\hat{\mu}_{t-1}^e(\mathbf{a}^e)$  is a weighted sum of  $n_{t-1}^e(\mathbf{a}^e)$  random variables distributed according to a  $\sigma$ -subgaussian with mean  $\mu^e(\mathbf{a}^e)$ . Hence, Hoeffding's inequality (H) is applicable [36].

$$\begin{aligned} &P(|\hat{\mu}_{t-1}^e(\mathbf{a}^e) - \mu^e(\mathbf{a}^e)| > c_t^e(\mathbf{a}^e) \mid \mu^e(\mathbf{a}^e)) \\ &\stackrel{(H)}{\leq} 2 \exp\left(-\frac{n_{t-1}^e(\mathbf{a}^e)}{2\sigma^2} (c_t^e(\mathbf{a}^e))^2\right) \\ &\stackrel{(8)}{=} 2 \exp\left(-\frac{n_{t-1}^e(\mathbf{a}^e)}{2\sigma^2} \frac{2\sigma^2 \log(\delta^{-1})}{n_{t-1}^e(\mathbf{a}^e)}\right) \\ &= 2 \exp(-\log(\delta^{-1})) \\ &= 2\delta. \end{aligned} \quad (11)$$

Therefore, the following concentration inequality on  $\bar{\mathcal{E}}_T$  holds:

$$P(\bar{\mathcal{E}}_T) \leq \sum_{t=1}^T \sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} 2\delta = 2\tilde{A}T\delta. \quad (12)$$

□

LEMMA 4.4. (Bayesian regret bound under  $\mathcal{E}_T$ ) *Provided that the error bound on the local sample means is never exceeded until time  $T$ , the Bayesian regret bound, when using the MATS policy  $\pi$ , is of the order*

$$\mathbb{E}[R(T, \pi) \mid \mathcal{E}_T] \leq \sqrt{32\sigma^2 \tilde{A} \rho T \log(\delta^{-1})}. \quad (13)$$

PROOF. Consider this upper bound on the sample means:

$$u_t(\mathbf{a}) \triangleq \sum_{e=1}^{\rho} \hat{\mu}_{t-1}^e(\mathbf{a}^e) + c_t^e(\mathbf{a}^e). \quad (14)$$

Given history  $\mathcal{H}_{t-1}$ , the statistics  $\hat{\mu}_{t-1}^e(\mathbf{a}^e)$  and  $n_{t-1}^e(\mathbf{a}^e)$  are known, rendering  $u_t(\cdot)$  a deterministic function. Therefore, the probability matching property of MATS (Equation 6) can be applied as follows:

$$\mathbb{E}[u_t(\mathbf{a}_t) \mid \mathcal{H}_{t-1}] = \mathbb{E}[u_t(\mathbf{a}_*) \mid \mathcal{H}_{t-1}]. \quad (15)$$

Hence, using the tower-rule (T), the regret can be bounded as

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^T \Delta(\mathbf{a}_t) \mid \mathcal{E}_T\right] &\stackrel{(T)}{=} \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}[\mu(\mathbf{a}_*) - \mu(\mathbf{a}_t) \mid \mathcal{H}_{t-1}, \mathcal{E}_T]\right] \\ &= \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}[\mu(\mathbf{a}_*) - u_t(\mathbf{a}_t) \mid \mathcal{H}_{t-1}, \mathcal{E}_T] \right. \\ &\quad \left. + \sum_{t=1}^T \mathbb{E}[u_t(\mathbf{a}_t) - \mu(\mathbf{a}_t) \mid \mathcal{H}_{t-1}, \mathcal{E}_T]\right] \\ &\stackrel{(15)}{=} \mathbb{E}\left[\sum_{t=1}^T \mathbb{E}[\mu(\mathbf{a}_*) - u_t(\mathbf{a}_*) \mid \mathcal{H}_{t-1}, \mathcal{E}_T] \right. \\ &\quad \left. + \sum_{t=1}^T \mathbb{E}[u_t(\mathbf{a}_t) - \mu(\mathbf{a}_t) \mid \mathcal{H}_{t-1}, \mathcal{E}_T]\right]. \end{aligned} \quad (16)$$

Note that the expression  $\mu(\mathbf{a}_*) - u_t(\mathbf{a}_*)$  is always negative under  $\mathcal{E}_T$ , i.e.,

$$\begin{aligned} \mu(\mathbf{a}_*) - u_t(\mathbf{a}_*) &\stackrel{(14)}{=} \sum_{e=1}^{\rho} \mu^e(\mathbf{a}_*) - \hat{\mu}_{t-1}^e(\mathbf{a}_*) - c_t^e(\mathbf{a}_*) \\ &\stackrel{(7)}{\leq} \sum_{e=1}^{\rho} c_t^e(\mathbf{a}_*) - c_t^e(\mathbf{a}_*) = 0, \end{aligned} \quad (17)$$

while  $u_t(\mathbf{a}_t) - \mu(\mathbf{a}_t)$  is bounded by twice the threshold  $c_t^e(\mathbf{a}^e)$ , i.e.,

$$\begin{aligned} u_t(\mathbf{a}_t) - \mu(\mathbf{a}_t) &\stackrel{(14)}{=} \sum_{e=1}^{\rho} \hat{\mu}_{t-1}^e(\mathbf{a}_t^e) + c_t^e(\mathbf{a}_t^e) - \mu^e(\mathbf{a}_t^e) \\ &\stackrel{(7)}{\leq} \sum_{e=1}^{\rho} c_t^e(\mathbf{a}_t^e) + c_t^e(\mathbf{a}_t^e) = 2 \sum_{e=1}^{\rho} c_t^e(\mathbf{a}_t^e). \end{aligned} \quad (18)$$

Thus, Equation 16 can be bounded as

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^T \Delta(\mathbf{a}_t) \mid \mathcal{E}_T\right] &\leq 2 \sum_{t=1}^T \sum_{e=1}^{\rho} c_t^e(\mathbf{a}_t^e) \\ &\leq 2 \sum_{t=1}^T \sum_{e=1}^{\rho} \sqrt{\frac{2\sigma^2 \log(\delta^{-1})}{n_{t-1}^e(\mathbf{a}_t^e)}} \\ &= 2 \sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} \sum_{t=1}^T \mathcal{I}\{\mathbf{a}_t^e = \mathbf{a}^e\} \sqrt{\frac{2\sigma^2 \log(\delta^{-1})}{n_{t-1}^e(\mathbf{a}^e)}}, \end{aligned} \quad (19)$$

where  $\mathcal{I}\{\cdot\}$  is the indicator function. The terms in the summation are only non-zero at the time steps when the local action  $\mathbf{a}^e$  is pulled, i.e., when  $\mathcal{I}\{\mathbf{a}_t^e = \mathbf{a}^e\} = 1$ . Additionally, note that only at these time steps, the counter  $n_t^e(\mathbf{a}^e)$  increases by exactly 1. Therefore, the following equality holds:

$$\sum_{t=1}^T \mathcal{I}\{\mathbf{a}_t^e = \mathbf{a}^e\} \sqrt{(n_{t-1}^e(\mathbf{a}^e))^{-1}} = \sum_{k=1}^{n_T^e(\mathbf{a}^e)} \sqrt{k^{-1}}. \quad (20)$$

The function  $\sqrt{k^{-1}}$  is decreasing and integrable. Hence, using the right Riemann sum,

$$\sqrt{k^{-1}} \leq \int_{k-1}^k \sqrt{x^{-1}} dx. \quad (21)$$

Combining Equations 19-21 leads to a bound

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^T \Delta(\mathbf{a}_t) \mid \mathcal{E}_T\right] &\stackrel{(19)}{\leq} 2 \sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} \sum_{t=1}^T \mathcal{I}\{\mathbf{a}_t^e = \mathbf{a}^e\} \sqrt{\frac{2\sigma^2 \log(\delta^{-1})}{n_{t-1}^e(\mathbf{a}^e)}} \\ &\stackrel{(20)}{=} \sqrt{8\sigma^2 \log(\delta^{-1})} \sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} \sum_{k=1}^{n_T^e(\mathbf{a}^e)} \sqrt{k^{-1}} \\ &\stackrel{(21)}{\leq} \sqrt{8\sigma^2 \log(\delta^{-1})} \sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} \int_0^{n_T^e(\mathbf{a}^e)} \sqrt{x^{-1}} dx \\ &= \sqrt{8\sigma^2 \log(\delta^{-1})} \sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} \sqrt{4n_T^e(\mathbf{a}^e)}. \end{aligned} \quad (22)$$

We use the relationship  $\|\mathbf{x}\|_1 \leq \sqrt{n} \|\mathbf{x}\|_2$  between the 1- and 2-norm of a vector  $\mathbf{x}$ , where  $n$  is the number of elements in the vector, as follows:

$$\sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} \left| \sqrt{n_T^e(\mathbf{a}^e)} \right| \leq \sqrt{\tilde{A}} \sqrt{\sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} \left( \sqrt{n_T^e(\mathbf{a}^e)} \right)^2}. \quad (23)$$

Finally, note that the sum of all counts  $n_T^e(\mathbf{a}^e)$  is equal to the total number of local pulls done by MATS until time  $T$ , i.e.,

$$\sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} n_T^e(\mathbf{a}^e) = \rho T. \quad (24)$$

Using the Equations 22-24, the complete regret bound under  $\mathcal{E}_T$  is given by

$$\begin{aligned}
& \mathbb{E} \left[ \sum_{t=1}^T \Delta(\mathbf{a}_t) \mid \mathcal{E}_T \right] \\
& \stackrel{(22)}{\leq} \sqrt{8\sigma^2 \log(\delta^{-1})} \sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} \sqrt{4n_T^e(\mathbf{a}^e)} \\
& \stackrel{(23)}{\leq} \sqrt{32\sigma^2 \log(\delta^{-1})} \sqrt{\tilde{A}} \sqrt{\sum_{e=1}^{\rho} \sum_{\mathbf{a}^e \in \mathcal{A}^e} \left( \sqrt{n_T^e(\mathbf{a}^e)} \right)^2} \\
& \stackrel{(24)}{=} \sqrt{32\sigma^2 \log(\delta^{-1})} \sqrt{\tilde{A}} \sqrt{\rho T}.
\end{aligned} \tag{25}$$

□

**THEOREM 4.5.** *Let  $\langle \mathcal{D}, \mathcal{A}, f \rangle$  be a MAMAB. If Assumptions 4.1 and 4.2 hold, then the MATS policy  $\pi$  satisfies a Bayesian regret bound of*

$$\begin{aligned}
\mathbb{E}[R(T, \pi)] & \leq \sqrt{64\sigma^2 \tilde{A} \rho T \log(\tilde{A}T)} + \frac{2}{\tilde{A}} \\
& \in O \left( \sqrt{\sigma^2 \tilde{A} \rho T \log(\tilde{A}T)} \right).
\end{aligned} \tag{26}$$

**PROOF.** Using the law of excluded middle (M) and the fact that  $\Delta(\mathbf{a}_t)$  and  $P(\mathcal{E}_T \mid \mathcal{H}_{t-1})$  are between 0 and 1 (B), the regret can be decomposed as

$$\begin{aligned}
& \mathbb{E} \left[ \sum_{t=1}^T \Delta(\mathbf{a}_t) \right] \\
& \stackrel{(M)}{=} \mathbb{E} \left[ \sum_{t=1}^T \Delta(\mathbf{a}_t) \mid \mathcal{E}_T \right] P(\mathcal{E}_T) + \mathbb{E} \left[ \sum_{t=1}^T \Delta(\mathbf{a}_t) \mid \bar{\mathcal{E}}_T \right] P(\bar{\mathcal{E}}_T) \\
& \stackrel{(B)}{\leq} \mathbb{E} \left[ \sum_{t=1}^T \Delta(\mathbf{a}_t) \mid \mathcal{E}_T \right] + TP(\bar{\mathcal{E}}_T).
\end{aligned} \tag{27}$$

Then, according to Lemmas 4.3 and 4.4 (L), we have

$$\begin{aligned}
\mathbb{E} \left[ \sum_{t=1}^T \Delta(\mathbf{a}_t) \right] & \stackrel{(27)}{\leq} \mathbb{E} \left[ \sum_{t=1}^T \Delta(\mathbf{a}_t) \mid \mathcal{E}_T \right] + TP(\bar{\mathcal{E}}_T) \\
& \stackrel{(L)}{\leq} \sqrt{32\sigma^2 \tilde{A} \rho T \log(\delta^{-1})} + 2\tilde{A}T^2\delta.
\end{aligned} \tag{28}$$

Finally, choosing  $\delta = (\tilde{A}T)^{-2}$ , we conclude that

$$\begin{aligned}
\mathbb{E}[R(T, \pi)] & \stackrel{(28)}{\leq} \sqrt{32\sigma^2 \tilde{A} \rho T \log(\delta^{-1})} + 2\tilde{A}T^2\delta \\
& \leq \sqrt{64\sigma^2 \tilde{A} \rho T \log(\tilde{A}T)} + \frac{2}{\tilde{A}} \\
& \in O \left( \sqrt{\sigma^2 \tilde{A} \rho T \log(\tilde{A}T)} \right).
\end{aligned} \tag{29}$$

□

**COROLLARY 4.6.** *If  $|\mathcal{A}_i| \leq k$  for all agents  $i$ , and if  $|\mathcal{D}^e| \leq d$  for all groups  $\mathcal{D}^e$ , then*

$$\mathbb{E}[R(T, \pi)] \in O \left( \rho \sqrt{\sigma^2 k^d T \log(\rho k^d T)} \right). \tag{30}$$

**PROOF.**  $\tilde{A} = \sum_{e=1}^{\rho} |\mathcal{A}^e| = \sum_{e=1}^{\rho} \prod_{i \in \mathcal{D}^e} |\mathcal{A}_i| \leq \rho k^d$ . □

Corollary 4.6 tells us that the regret is sub-linear in terms of time  $T$  and low-order polynomial in terms of the largest action space of a single agent when the number of groups and agents per group are small. This reflects the main contribution of this work. When agents are loosely coupled, the *effective* joint arm space is significantly reduced, and MATS provides a mechanism that efficiently deals with such settings. This is a significant improvement over the established classic regret bounds of vanilla TS when the MAMAB is ‘flattened’ and the factored structure is neglected [23, 30]. The classic bounds scale exponentially with the number of agents, which renders the use of vanilla TS unfeasible in many multi-agent environments.

## 5 EXPERIMENTS

We evaluate the performance of MATS on the benchmark problems proposed in the paper that introduced Multi-Agent Upper-Confidence Exploration (MAUCE) [4], which is the current state-of-the-art algorithm for multi-agent bandit problems, and one novel setting that falls outside the domain of the theoretical guarantees for both MAUCE and MATS. First, we evaluate the performance of MATS on two benchmarks that were introduced in the MAUCE paper, i.e., Bernoulli 0101-Chain and Gem Mining. We compare against a random policy (rnd), Sparse Cooperative Q-Learning (SCQL) [21] and the state-of-the-art algorithm, MAUCE [4]. For SCQL and MAUCE, we use the same exploration parameters as in previous work [4]. For MATS, we always use non-informative Jeffreys priors, which are invariant toward reparametrization of the experimental settings [29]. Although including additional prior domain knowledge could be useful in practice, we use well-known non-informative priors in our experiments to compare fairly with other state-of-the-art techniques. Then, we introduce a novel variant of the 0101-Chain with Poisson-distributed local rewards. A Poisson distribution is supergaussian, meaning that its tails tend slower towards zero than the tails of any Gaussian. Therefore, both the assumptions made in Theorem 4.5 and in the established regret bound of MAUCE are violated. Additionally, as the rewards are highly skewed, we expect that the use of symmetric exploration bounds in MAUCE will often lead to either over- or underexploration of the local arms. We assess the performance of both methods on this benchmark.

### 5.1 Bernoulli 0101-Chain

The Bernoulli 0101-Chain consists of  $n$  agents and  $n-1$  local reward distributions. Each agent can choose between two actions: 0 and 1. In the coordination graph, agents  $i$  and  $i+1$  are connected to a local reward  $f^i(a_i, a_{i+1})$ . Thus, each pair of agents should locally coordinate in order to find the best joint arm. The local rewards are drawn from a Bernoulli distribution with a different success probability per group. These success probabilities are given in Table 1. The optimal joint action is an alternating sequence of zeros and ones, starting with 0.

To ensure that the assumptions made in the regret analyses of MAUCE and MATS hold, we divide the local rewards by the number of groups, such that the global rewards are between 0 and 1.

We provide non-informative Jeffreys priors on the unknown means to MATS, which for the Bernoulli likelihood is a Beta prior,




Figure 1: Cumulative normalized regret averaged over 100 runs for the (a) Bernoulli 0101-Chain, (b) Gem Mining and (c) Poisson 0101-Chain. Both the mean (line) and standard deviation (shaded area) are plotted.

$f^i \sim \mathcal{B}$	$a_{i+1} = 0$	$a_{i+1} = 1$
$a_i = 0$	0.75	1
$a_i = 1$	0.25	0.9

Table 1: Bernoulli 0101 Chain – The unscaled local reward distributions of agents  $i$  and  $i + 1$ , where  $i$  is even. Each entry shows the success probability for each local arm of agents  $i$  and  $i + 1$ , where  $i$  is even. The table is transposed for the case where  $i$  is odd.

$\mathcal{B}(\alpha = 0.5, \beta = 0.5)$  [26]. The results for the Bernoulli 0101-chains are shown in Figure 1(a).

## 5.2 Gem Mining

In the Gem Mining problem, a mining company wants to excavate a set of mines for gems (i.e., local rewards). The goal is to maximize the total number of gems found over all mines. However, the company’s workers live in separate villages (i.e., agents), and only one van per village is available. Therefore, each village needs to decide to which mine it should send its workers (i.e., local action). Moreover, workers can only commute to nearby mines (i.e., coordination graph). Hence, a group can be constructed per mine, consisting of all agents that can travel toward the mine.

The reward is drawn from a Bernoulli distribution, where the probability of finding a gem at a mine is  $1.03^{w-1}p$  with  $w$  the number of workers at the mine and  $p$  a base probability that is sampled uniformly random from the interval  $[0, 0.5]$  for each mine. When more workers are excavating a mine, the probability of finding a gem increases. Each village is populated by a number sampled uniformly random from  $[1..5]$ . The coordination graph is generated by sampling for each village  $i$  a number of mines  $m_i$  in  $[2..4]$  to which it should be connected. Then, each village  $i$  is connected to the mines  $i$  to  $(i + m_i - 1)$ . The last village is always connected to 4 mines.

We provide non-informative Jeffreys priors on the unknown means to MATS, which for the Bernoulli likelihood is a Beta prior,

$\mathcal{B}(\alpha = 0.5, \beta = 0.5)$  [26]. The results for the Gem Mining problem are shown in Figure 1(b).

## 5.3 Poisson 0101-Chain

We introduce a novel benchmark with Poisson distributed local rewards, for which the established regret bounds of MATS and MAUCE do not hold. Similar to the Bernoulli 0101-Chain, agents need to coordinate their actions in order to obtain an alternating sequence of zeroes and ones. However, as the rewards are highly skewed and supergaussian, this setting is much more challenging. The means of the Poisson distributions are given in Table 2. We also divide the rewards by the number of groups, similar to the Bernoulli 0101-Chain.

$f^i \sim \mathcal{P}$	$a_{i+1} = 0$	$a_{i+1} = 1$
$a_i = 0$	0.1	0.3
$a_i = 1$	0.2	0.1

Table 2: Poisson 0101 Chain – The unscaled local reward distributions of agents  $i$  and  $i + 1$ . Each entry shows the mean for each local arm of agents  $i$  and  $i + 1$ .

For MAUCE, an exploration parameter must be chosen. This exploration parameter denotes the range of the observed rewards. As a Poisson distribution has unbounded support, we rely on percentiles of the reward distribution. Specifically, as 95% of the rewards when pulling the optimal arm falls below 1, we choose 1 as the exploration parameter of MAUCE. For MATS we use non-informative Jeffreys priors on the unknown means, which for the Poisson likelihood is a Gamma prior,  $\mathcal{G}(\alpha = 0.5, \beta = 0)$  [26]. The results are shown in Figure 1(c).


## 6 WIND FARM CONTROL APPLICATION

We demonstrate the benefits of MATS on a state-of-the-art wind farm simulator and compare its performance to MAUCE and SCQL. A wind farm consists of a group of wind turbines, instantiated to extract energy from wind. From the perspective of a single turbine,

aligning with the incoming wind vector usually ensures the highest productivity. However, translating this control policy directly towards an entire wind farm may be sub-optimal. As wind passes through the farm, downstream turbines observe a significantly lower wind speed. This is known as the *wake effect*, which is due to the turbulence generated behind operational turbines.

In recent work, the possibility of deflecting wake away from the farm through rotor misalignment is investigated [35]. While a misaligned turbine produces less energy on its own, the group’s total productivity is increased. Physically, the wake effect reduces over long distances, and thus, turbines tend to only influence their neighbours. We can use this domain knowledge to define groups of agents and organize them in a graph structure. Note that the graph structure depends on the incoming wind vector. Nevertheless, atmospheric conditions are typically discretized when analyzing operational regimes [19], thus, a graph structure can be made independently for each possible incoming discretized wind vector. We construct a graph structure for one possible wind vector.

We demonstrate our method on a virtual wind farm, consisting of 11 turbines, of which the layout is shown in Figure 2. We use the state-of-the-art WISDEM FLORIS simulator [28]. For MATS, we



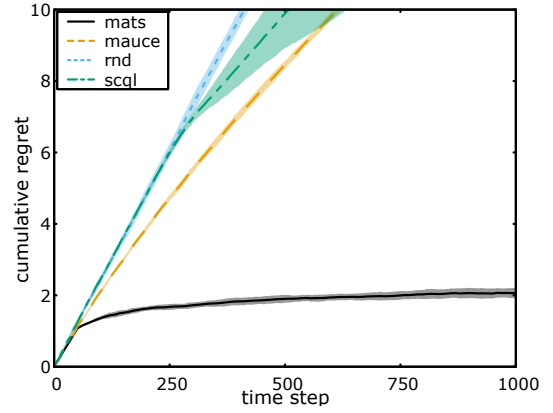
**Figure 2: Wind farm layout – Dependency graph where the nodes are the turbines and the edges describe the dependencies between the turbines. The incoming wind is denoted by an arrow.**

assume the local power productions are sampled from Gaussians with unknown mean and variance, which leads to a Student’s *t*-distribution on the mean when using a Jeffreys prior [18]. The results for the wind farm control setting are shown in Figure 3.

## 7 DISCUSSION

MATS is a Bayesian method, which means that it can leverage prior knowledge about the data distribution. This property is highly beneficial in many practical applications, e.g., influenza mitigation [24, 25] and wind farm control [37].

Both MAUCE and MATS achieve sub-linear regret in terms of time and low-order polynomial regret in terms of the number of local arms for sparse coordination graphs. However, empirically, MATS consistently outperforms MAUCE as well as SCQL. We can see that MATS solves the Bernoulli 0101-Chain problem in only a few time steps, while MAUCE still pulls many sub-optimal actions



**Figure 3: Cumulative normalized regret averaged over 10 runs for Wind Farm task. Both the mean (line) and standard deviation (shaded area) are plotted.**

after 10000 time steps (see Figure 1(a)). In the more challenging Gem Mining problem, the cumulative regret of MAUCE is three times as high as the cumulative regret of MATS around 40000 time steps (see Figure 1(b)). In the wind farm control task, we can see that MATS allowed for a five-fold increase of the normalized power productions with respect to the state of the art (see Figure 3). We argue that the high performance of MATS is due to the ability to seamlessly include domain knowledge about the shape of the reward distributions and treat the problem parameters as unknowns. To highlight the power of this property, we introduced the Poisson 0101-chain. In this setting, the reward distributions are highly skewed, for which the mean does not match the median. Therefore, in our case, since the mean falls well above 50% of all samples, it is expected that for the initially observed rewards, the true mean will be higher than the sample mean. Naturally, this bias averages out in the limit, but may have a large impact during the early exploration stage. The high standard deviations in Figure 1(c) support this impact. Although the established regret bounds of MATS and MAUCE do not apply for supergaussian reward distributions, we demonstrate that MATS exploits density information of the rewards to achieve more targeted exploration. In Figure 1(c), the cumulative regret of MATS stagnates around 7500 time steps, while the cumulative regret of MAUCE continues to increase significantly. As MAUCE only supports symmetric exploration bounds, it is challenging to correctly assess the amount of exploration needed to solve the task.

Throughout the experiments, exploration constants had to be specified for MAUCE, which were challenging to choose and interpret in terms of the density of the data. In contrast, MATS uses either statistics about the data (if available) or, potentially non-informative, beliefs defined by the user. For example, in the wind farm case, the spread of the data is unknown. MATS effectively maintains a posterior on the variance and uses it to balance exploration and exploitation, while still outperforming MAUCE with a manually calibrated exploration range (see Figure 3).

## 8 RELATED WORK

Multi-agent reinforcement learning and planning with loose couplings has been investigated in sequential decision problems [12, 17, 20, 31]. In sequential settings, the value function cannot be factorized exactly. Therefore, it is challenging to provide convergence and optimality guarantees. While for planning some theoretical guarantees can be provided [31], in the learning literature the focus has been on empirical validation [20]. In this work, we focus on MAMABs, which are single-shot stateless problems. In such settings, the reward function is factored exactly into components that only depend on a subset of agents.

The combinatorial bandit [6, 7, 10, 13] is a variant of the multi-armed bandit, in which, rather than one-dimensional arms, an arm vector has to be pulled. In our work, the arms’ dimensionality corresponds to the number of agents in our system, and similarly to combinatorial bandits, the number of arms exponentially increases with this quantity. We consider a variant of this framework, called the semi-bandit problem [3], in which local components of the global reward are observable. Chen et. al (2013) constructed an algorithm for this setting that assumes access to an  $(\alpha, \beta)$ -oracle, which provides a joint action that outputs a fraction  $\alpha$  of the optimal expected reward with probability  $\beta$ . Instead, we assume the availability of a coordination graph, which we argue is a reasonable assumption in many multi-agent settings.

Sparse cooperative Q-learning is an algorithm that also assumes the availability of a coordination graph [21]. However, although strong experimental results are given, no theoretical guarantees were provided. Later, the UCB-like algorithm, HEIST, for exploration and exploitation in MAMABs was introduced [32], which uses a message-passing scheme for resolving coordination graphs. They provide some theoretical guarantees on the regret for problems with acyclic coordination graphs. Multi-Agent Upper-Confidence Exploration (MAUCE) [4] is a more general method that uses variable elimination to resolve (potentially cyclic) coordination graphs. MAUCE demonstrates high performance on a variety of benchmarks and provides a tight theoretical upper bound on the regret. MATS provides a Bayesian alternative to MAUCE based on Thompson sampling (TS).

Our problem definition is related to distributed constraint optimization (DCOP) problems [40]. In DCOP problems, multiple agents control a set of variables in a distributed manner under a set of constraints. The objective is the same as for a MAMAB, i.e., optimize the sum over group rewards. However, in DCOPs, the rewards are assumed to be known beforehand. The Distributed Coordination of Exploration and Exploitation (DCEE) framework [33] extends this setting to unknown rewards, but considers the optimization of the cumulative reward achieved over a time span, rather than of a single-step reward. MAMABs, or MAB-DCOPs [32], consider the optimization of a single-step expected reward over time.

In recent research on wind farm control, the impact of optimized rotor alignments on power production is heavily investigated [35]. To search for the optimal alignments within the wind farm, data-driven methods are usually adopted, where the turbines’ alignments are perturbed iteratively until they locally converge [27]. When optimizing the alignment of a wind turbine, only considering its neighbours can significantly boost the learning speed [14]. MATS is

also able to leverage neighbourhood structures. In addition, rather than random perturbation of the alignments, MATS leverages an exploration-exploitation mechanism that is inspired by TS and variable elimination, which allows for a global exploration mechanism that targets the optimal alignment configuration, while retaining a small regret during the learning process itself.

## 9 CONCLUSIONS

We proposed multi-agent Thompson sampling (MATS), a novel Bayesian algorithm for multi-agent multi-armed bandits. The method exploits loose connections between agents to solve multi-agent coordination tasks efficiently. Specifically, we proved that, for  $\sigma$ -subgaussian rewards with bounded means, the expected cumulative regret decreases sub-linearly in time and low-order polynomially in the highest number of actions of a single agent when the coordination graph is sparse. Empirically, we showed a significant improvement over the state-of-the-art algorithm, MAUCE, on several synthetic benchmarks. Additionally, we showed that MATS can seamlessly be adapted to the available prior knowledge, and achieves state-of-the-art performance on the Poisson 0101-Chain, a new benchmark with supergaussian rewards. Finally, we demonstrated that MATS achieves high performance on a realistic wind farm control task, where the optimal rotor alignments of the wind turbines need to be jointly optimized to maximize the farm’s power production. In many practical applications, there exist sparse neighbourhood structures between agents, and we have shown that MATS is able to successfully exploit these structures, while leveraging prior knowledge about the data.

## 10 ACKNOWLEDGMENTS

The authors would like to acknowledge FWO (Fonds Wetenschappelijk Onderzoek) for their support through the SB grants of Timothy Verstraeten (#1S47617N), Eugenio Bargiacchi (#1SA2820N) and Pieter J.K. Libin (#1S31916N). This research was supported by funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme and under the VLAIO Supersized 4.0 ICON project.

## REFERENCES

- [1] Shipra Agrawal and Navin Goyal. 2012. Analysis of Thompson sampling for the multi-armed bandit problem. In *Proceedings of the 25th Annual Conference on Learning Theory (COLT)*, Vol. 23. 39.1–39.26.
- [2] Shipra Agrawal and Navin Goyal. 2013. Further optimal regret bounds for Thompson sampling. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics (AISTATS)*, Vol. 31. 99–107.
- [3] Jean-Yves Audibert, Sébastien Bubeck Bubeck, and Gábor Lugosi. 2011. Minimax Policies for Combinatorial Prediction Games. In *Proceedings of the 24th Annual Conference on Learning Theory (COLT)*, Vol. 19. 107–132.
- [4] Eugenio Bargiacchi, Timothy Verstraeten, Diederik M. Roijers, Ann Nowé, and Hado Hasselt. 2018. Learning to Coordinate with Coordination Graphs in Repeated Single-Stage Multi-Agent Decision Problems. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. 482–490.
- [5] Craig Boutilier. 1996. Planning, learning and coordination in multiagent decision processes. In *TARK 1996: Proceedings of the 6th conference on Theoretical aspects of rationality and knowledge*. 195–210.
- [6] Sébastien Bubeck and Nicoló Cesa-Bianchi. 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning* 5, 1 (2012), 1–122.
- [7] Nicoló Cesa-Bianchi and Gábor Lugosi. 2012. Combinatorial bandits. *J. Comput. System Sci.* 78, 5 (2012), 1404–1422.
- [8] Olivier Chapelle and Lihong Li. 2011. An empirical evaluation of Thompson sampling. In *Advances in Neural Information Processing Systems (NIPS)*, Vol. 24. 2249–2257.



- [9] Archie C. Chapman, David S. Leslie, Alex Rogers, and Nicholas R. Jennings. 2013. Convergent learning algorithms for unknown reward games. *SIAM Journal on Control and Optimization* 51, 4 (2013), 3154–3180.
- [10] Wei Chen, Yajun Wang, and Yuan Yuan. 2013. Combinatorial multi-armed bandit: General framework, results and applications. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, Vol. 28. 151–159.
- [11] Daniel Claes, Frans Oliehoek, Hendrik Baier, and Karl Tuyls. 2017. Decentralised Online Planning for Multi-Robot Warehouse Commissioning. In *Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. International Foundation for Autonomous Agents and Multiagent Systems, 492–500.
- [12] Yann-Michaël De Hauwere, Peter Vrancx, and Ann Nowé. 2010. Learning multi-agent state space representations. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. International Foundation for Autonomous Agents and Multiagent Systems, 715–722.
- [13] Yi Gai, Bhaskar Krishnamachari, and Rahul Jain. 2012. Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *IEEE/ACM Transactions on Networking (TON)* 20, 5 (2012), 1466–1478.
- [14] Pieter MO Gebraad and Jan-Willem van Wingerden. 2015. Maximum power-point tracking control for wind farms. *Wind Energy* 18, 3 (2015), 429–447.
- [15] Carlos Guestrin, Daphne Koller, and Ronald Parr. 2001. Max-norm projections for factored MDPs. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI)*. 673–682.
- [16] Carlos Guestrin, Daphne Koller, and Ronald Parr. 2001. Multiagent Planning with Factored MDPs. In *Advances in Neural Information Processing Systems (NIPS)*, Vol. 14. 1523–1530.
- [17] Carlos Guestrin, Shobha Venkataraman, and Daphne Koller. 2002. Context-specific multiagent coordination and planning with factored MDPs. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI)*. 253–259.
- [18] Junya Honda and Akimichi Takemura. 2014. Optimality of Thompson sampling for Gaussian bandits depends on priors. In *Artificial Intelligence and Statistics*. 375–383.
- [19] International Electrotechnical Commission. 2012. Wind Turbines – Part 4: Design requirements for wind turbine gearboxes (No. IEC 61400-4). (2012). <https://www.iso.org/standard/44298.html> accessed 6 March 2019.
- [20] Jelle Kok and Nikos Vlassis. 2006. Using the Max-Plus Algorithm for Multiagent Decision Making in Coordination Graphs. In *RoboCup 2005: Robot Soccer World Cup IX*. Ansgar Bredendfeld, Adam Jacoff, Itsuki Noda, and Yasutake Takahashi (Eds.). Lecture Notes in Computer Science, Vol. 4020. Springer, 1–12.
- [21] Jelle R. Kok and Nikos Vlassis. 2004. Sparse Cooperative Q-learning. In *Proceedings of the 21st International Conference on Machine Learning (ICML)*. New York, NY, USA.
- [22] Daphne Koller and Ronald Parr. 2000. Policy Iteration for Factored MDPs. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 326–334.
- [23] Tor Lattimore and Csaba Szepesvári. 2018. Bandit algorithms. *preprint* (2018).
- [24] Pieter J.K. Libin, Timothy Verstraeten, Diederik M. Roijers, Jelena Grujic, Kristof Theys, Philippe Lemey, and Ann Nowé. 2018. Bayesian Best-Arm Identification for Selecting Influenza Mitigation Strategies. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*. Springer, 456–471.
- [25] Pieter J.K. Libin, Timothy Verstraeten, Diederik M. Roijers, Wenjia Wang, Kristof Theys, and Ann Nowé. 2019. Thompson Sampling for m-top Exploration. In *Proceedings of the IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*. 1414–1420.
- [26] David Lunn, Chris Jackson, Nicky Best, David Spiegelhalter, and Andrew Thomas. 2012. *The BUGS book: A practical introduction to Bayesian analysis*. Chapman and Hall/CRC.
- [27] Jason R. Marden, Shalom D. Ruben, and Lucy Y. Pao. 2013. A model-free approach to wind farm control using game theoretic methods. *IEEE Transactions on Control Systems Technology* 21, 4 (2013), 1207–1214.
- [28] NREL. 2019. FLORIS. Version 1.0.0. (2019). <https://github.com/NREL/floris>
- [29] Christian Robert. 2007. *The Bayesian choice: from decision-theoretic foundations to computational implementation*. Springer Science & Business Media.
- [30] Daniel Russo and Benjamin Van Roy. 2014. Learning to optimize via posterior sampling. *Mathematics of Operations Research* 39, 4 (2014), 1221–1243.
- [31] Joris Scharpf, Diederik M. Roijers, Frans A. Oliehoek, Matthijs T.J. Spaan, and Matthijs M. de Weerd. 2016. Solving Transition-Independent Multi-agent MDPs with Sparse Interactions. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*.
- [32] Ruben Stranders, Long Tran-Thanh, Francesco M. Delle Fave, Alex Rogers, and Nicholas R. Jennings. 2012. DCOPs and bandits: Exploration and exploitation in decentralised coordination. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. International Foundation for Autonomous Agents and Multiagent Systems, 289–296.
- [33] Matthew E. Taylor, Manish Jain, Prateek Tandon, Makoto Yokoo, and Milind Tambe. 2011. Distributed on-line multi-agent optimization under uncertainty: Balancing exploration and exploitation. *Advances in Complex Systems* 14, 03 (2011), 471–528.
- [34] William R. Thompson. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25, 3/4 (1933), 285–294.
- [35] Mike T. van Dijk, Jan-Willem Wingerden, Turaj Ashuri, Yaoyu Li, and Mario A. Rotea. 2016. Yaw-Misalignment and its Impact on Wind Turbine Loads and Wind Farm Power Output. *Journal of Physics: Conference Series* 753, 6 (2016).
- [36] Roman Vershynin. 2018. *High-dimensional probability: An introduction with applications in data science*. Vol. 47. Cambridge University Press.
- [37] Timothy Verstraeten, Ann Nowé, Jonathan Keller, Yi Guo, Shuangwen Sheng, and Jan Helsen. 2019. Fleetwide data-enabled reliability improvement of wind turbines. *Renewable and Sustainable Energy Reviews* 109 (2019), 428–437.
- [38] Nikos Vlassis, Reinoud Elhorst, and Jelle R. Kok. 2004. Anytime algorithms for multiagent decision making using coordination graphs. In *IEEE International Conference on Systems, Man and Cybernetics*, Vol. 1. 953–957.
- [39] Marco Wiering. 2000. Multi-agent reinforcement learning for traffic light control. In *Proceedings of the 17th International Conference on Machine Learning (ICML)*. 1151–1158.
- [40] Makoto Yokoo, Edmund H. Durfee, Toru Ishida, and Kazuhiro Kuwabara. 1998. The distributed constraint satisfaction problem: Formalization and algorithms. *IEEE Transactions on knowledge and data engineering* 10, 5 (1998), 673–685.