

Thompson Sampling for Stochastic Bandits with Graph Feedback

Aristide C. Y. Tossou

Computer Science and Engineering
Chalmers University of Technology
Gothenburg, Sweden
aristide@chalmers.se

Christos Dimitrakakis

University of Lille, France
Chalmers University of Technology
Harvard University, USA
christos.dimitrakakis@gmail.com

Devdatt Dubhashi

Computer Science and Engineering
Chalmers University of Technology
Gothenburg, Sweden
dubhashi@chalmers.se

Abstract

We present a novel extension of Thompson Sampling for stochastic sequential decision problems with graph feedback, even when the graph structure itself is unknown and/or changing. We provide theoretical guarantees on the Bayesian regret of the algorithm, linking its performance to the underlying properties of the graph. Thompson Sampling has the advantage of being applicable without the need to construct complicated upper confidence bounds for different problems. We illustrate its performance through extensive experimental results on real and simulated networks with graph feedback. More specifically, we tested our algorithms on power law, planted partitions and Erdős–Rényi graphs, as well as on graphs derived from Facebook and Flixster data. These all show that our algorithms clearly outperform related methods that employ upper confidence bounds, even if the latter use more information about the graph.

1 Introduction

Sequential decision making problems under uncertainty appear in most modern applications, such as automated experimental design, recommendation systems and optimisation. The common structure of these applications is that, at each time step t , the decision-making agent is faced with a choice. After each decision, it obtains some problem-dependent feedback (Cesa-Bianchi and Lugosi 2006). For the so-called *bandit* problem, the choices are between different *arms*, and the feedback consists of a single scalar reward obtained by the arm at time t . For the *prediction* (or full-information) problem, it obtains the reward of the chosen arm, but also observes the rewards of all other choices at time t . In both cases, the problem is to maximise the total reward obtained over time. However, dealing with specific types of feedback may require specialised algorithms. In this paper, we show that the *Thompson sampling* algorithm can be applied successfully to a range of sequential decision problems, whose feedback structure is characterised by a graph.

Our algorithm is an extension of Thompson sampling, introduced in (Thompson 1933). Although easy to implement and effective in practice, it remained unpopular until relatively recently. Interest grew after empirical studies

(Scott 2010; Chapelle and Li 2011) demonstrated performance exceeding state of the art. This has prompted a surge of interest in Thompson sampling, with the first theoretical results (Agrawal and Goyal 2012) and industrial adoption (Scott 2015) appearing only recently. However, there are still only a few theoretical results and many of these are in the simplest settings. However, it is easy to implement and effective under very many different settings with complex feedback structures, and there is thus great need to extend the theoretical results to these wider settings.

Russo and Roy argue that Thompson sampling is a very effective and versatile strategy for different *information structures*. Their paper focuses on specific examples: the two extreme cases of no and full information mentioned above and the case of linear bandits and combinatorial feedback.

Here we consider the case where the feedback is defined through a graph (Caron et al. 2012; Alon et al. 2015). More specifically, the arms (choices) are vertices of a (potentially changing) graph and when an arm is chosen, we see the reward of that arm as well as its neighbours. On one hand, it is a clean model for theoretical and experimental analysis and on the other hand, it also corresponds to realistic settings in social networks, for example in advertisement settings (c.f. (Caron et al. 2012)).

We provide a problem-independent¹ regret bound that is parametrized by the clique cover number of the graph and naturally generalizes the two extreme cases of zero and full information. We present two variants of Thompson sampling, that are both very easy to implement and computationally efficient. The first is straightforward Thompson sampling, and so draws an arm according to its probability of being the best, but also uses the graph feedback to update the posterior distribution. The second one can be seen as sampling cliques in the graph according to their probability of containing the best arm, and then choosing the empirically best arm in that clique. Neither algorithm requires knowledge of the complete graph.

Almost all previous algorithms require the full structure of the feedback graph in order to operate. Some require the entire graph for performing their updates only at the end of round (e.g. (Alon et al. 2015)). Others actually need the description of the graph at the beginning of the round to make

¹In the sense that it does not depend on the reward structure.

their decision and almost none of the algorithms previously proposed in the literature is able to provide non-trivial regret guarantees without the feedback graphs being disclosed. However, Cohen, Hazan, and Koren (2016) argue that the assumption that the entire observation system is revealed to the learner on each round, even if only after making her prediction, is rather unnatural. In principle, the learner need not be even aware of the fact that there is a graph underlying the feedback model; the feedback graph is merely a technical notion for us to specify a set of observations for each of the possible arms. Ideally, the only signal we would like the agent to receive following each round is the set of observations that corresponds to the arm she has taken on that round (in addition to the obtained reward). Our algorithms work in this setup - they do not need the whole graph to be disclosed either when selecting the arm or when updating beliefs - only the local neighborhood is needed. Furthermore, the underlying graph is allowed to change arbitrarily at each step. The detailed proofs of all our main results are available in the full version of this paper.

2 Setting

2.1 The stochastic bandit model

The stochastic K-armed bandit problem is a well known sequential decision problem involving an agent sequentially choosing among a set of K arms $\mathcal{V} = \{1 \dots K\}$. At each round t , the agent plays an arm $A_t \in \mathcal{V}$ and receives a reward $r_t = R(Y_{t,A_t})$, where $Y_{t,A_t} : \Omega \rightarrow \mathcal{Y}$ is a random variable defined on some probability space (P, Ω, Σ) and $R : \mathcal{Y} \rightarrow \mathbb{R}$ is a reward function.

Each arm i has mean reward $\mu_i(P) = \mathbb{E}_P R(Y_{t,i})$. Our goal is to maximize its expected cumulative reward after T rounds. An equivalent notion is to minimize the expected regret against an oracle which knows P . More formally, the expected regret $\mathbb{E}_P^\pi \mathcal{L}$ of an agent policy π for a bandit problem P is defined as:

$$\mathbb{E}_P^\pi \mathcal{L} = T\mu_*(P) - \mathbb{E}_P^\pi \sum_{t=1}^T r_{A_t}, \quad (2.1)$$

where $\mu_*(P) = \max_{i \in \mathcal{V}} \mu_i(P)$ is the mean of the optimal arm and $\pi(A_t|h_t)$ is the policy of the agent, defining a probability distribution on the next arm A_t given the history $h_t = \langle A_{1:t-1}, r_{1:t-1} \rangle$ of previous arms and rewards.

The main challenge in this model is that the agent does not know P , and it only observes the reward of the arm it played. As a consequence, the agent must trade-off exploitation (taking the apparently best arm) with exploration (trying out other arms to assess their quality).

The Bayesian setting offers a natural way to model this uncertainty, by assuming that the underlying probability law P is in some set $\mathcal{P} = \{P_\theta \mid \theta \in \Theta\}$ parametrised by θ , over which we define a prior probability distribution \mathbb{P} . In that case, we can define the Bayesian regret:

$$\mathbb{E}^\pi \mathcal{L} = \int_{\Theta} \mathbb{E}_{P_\theta}^\pi (\mathcal{L}) d\mathbb{P}(\theta). \quad (2.2)$$

A policy with small Bayesian regret may not be uniformly good in all P . Since in the Bayesian setting we frequently

need to discuss posterior probabilities and expectations, we also introduce the notation $\mathbb{E}_t f \triangleq \mathbb{E}(f \mid h_t)$ and $\mathbb{P}_t(\cdot) \triangleq \mathbb{P}(\cdot \mid h_t)$ for expectations and probabilities conditioned on the current history.

2.2 The graph feedback model

In this model, we assume the existence of an undirected graph $G = (\mathcal{V}, \mathcal{E})$ with vertices corresponding to arms. By taking an arm $a \in \mathcal{V}$, we not only receive the reward of the arm we played, but we also observe the rewards of all neighbouring arms $\mathcal{N}_a = \{a' \in \mathcal{V} \mid (a, a') \in \mathcal{E}\}$. More precisely, at each time-step t we observe $Y_{t,a'}$ for all $a' \in \mathcal{N}_{A_t}$, while our reward is still $r_t = R(Y_{t,A_t})$.

If the graph is empty, then the setting is equivalent to the bandit problem. If the graph is fully connected, then it is equivalent to the prediction (i.e. full information) problem. However, many practical graphs, such as those derived from social networks, have an intermediate connectivity. In such cases, the amount of information that we can obtain by picking an arm can be characterised by graph properties, such as the clique cover number:

Definition 2.1 (Clique cover number). A clique covering \mathcal{C} of a graph G is a partition of all its vertices into sets $S \in \mathcal{C}$ such that the sub-graph formed by each S is a clique i.e. all vertices in S are connected to each other in G . The smallest number of cliques into which the nodes of G can be partitioned is called the clique cover number. We denote by $\mathcal{C}(G)$ the minimum clique cover and $\chi(G)$ its size, omitting G when clear from the context.

The domination number is another useful similar notion for the amount of information that we can obtain.

Definition 2.2 (Domination number). A dominating set in a graph $G = (\mathcal{V}, \mathcal{E})$ is a subset $U \subseteq \mathcal{V}$ such that for every vertex $u \in \mathcal{V}$, either $u \in U$ or $(u, v) \in \mathcal{E}$ for some $v \in U$. The smallest size of a dominating set in G is called the domination number of G and denoted $\gamma(G)$.

3 Related work and our contribution

Optimal policies for the stochastic multi-armed bandit problem were first characterised by (Lai and Robbins 1985), while index-based optimal policies for general non-parametric problems were given by (Burnetas and Katehakis 1997). Later (Auer, Cesa-Bianchi, and Fischer 2002a) proved finite-time regret bounds for a number of UCB (Upper Confidence Bound) index policies, while (Garivier and Cappé 2011) proved finite-time bounds for index policies similar to those of (Burnetas and Katehakis 1997), with problem-dependent bounds $O(K \ln T)$. Recently, a number of policies based on sampling from the posterior distribution (i.e. Thompson sampling (Thompson 1933)) were analysed in both the frequentist (Agrawal and Goyal 2012) and Bayesian setting (Russo and Roy 2016) and shown to obtain the same order of regret bound for the stochastic case. For the *adversarial* bandit problem the bounds are of order $O(\sqrt{KT})$. The analysis for the full information case generally results in $O(\ln(K)\sqrt{T})$ bounds on the regret (Cesa-Bianchi and Lugosi 2006), i.e. with a much lower dependence on the number of arms.

Intermediate cases between full information and bandit feedback can be obtained through graph feedback, introduced in (Mannor and Shamir 2011), which is the focus of this paper. In particular, (Caron et al. 2012) and (Alon et al. 2015) analysed graph feedback problems with stochastic and adversarial reward sequences respectively. Specifically, Caron et al. analysed variants of Upper Confidence Bound policies, for which they obtained $O(\chi(\bar{G}) \ln T)$ problem-dependent bounds. In more recent work, (Cohen, Hazan, and Koren 2016) also introduced algorithms for graphs where the structure is never fully revealed showing that (unlike the bandit setting) there is a large gap in the regret between the adversarial and stochastic cases. In particular, they show that in the adversarial setting one cannot do any better than ignore all additional feedback, while they provide an action-elimination algorithm for the stochastic setting. Finally, (Buccapatnam, Eryilmaz, and Shroff 2014) obtain a problem-dependent bound of the form $O(\gamma^*(G) \log T + K\delta)$ where γ^* is the linear programming relaxation to γ and δ is the minimum degree of G .

Contributions. In this paper, we provide much simpler strategies based on Thompson sampling, with a matching regret bound. Unlike previous work, these are also applicable to graphs whose structure is unknown or changing over time. More specifically:

1. We extend (Russo and Roy 2016) to graph-structured feedback, and obtain a problem-independent bound of $O(\sqrt{\frac{1}{2}\chi(\bar{G})T})$.
2. Using planted partition models, we verify the bound's dependence on the clique cover.
3. We provide experiments on data drawn from two types of random graphs: Erdős–Rényi graphs and power law graphs, showing that our algorithms clearly outperform UCB and its variations (Caron et al. 2012).
4. Finally, we measured the performance on graphs estimated from the data used in (Caron et al. 2012). Once again, Thompson sampling clearly outperforms UCB and its variants.

4 Algorithms and analysis

We consider two algorithms based on Thompson sampling. The first uses standard Thompson sampling to select arms. As this also reveals the rewards of neighbouring arms, the posterior is conditioned on those as well. The second algorithm uses Thompson sampling to select an arm, and then chooses the empirically best arm within that arm's clique.

4.1 The TS-N policy

The TS-N policy is an adaptation of Thompson Sampling for graph-structured feedback. Thompson Sampling maintains a distribution over the problem parameters. At each step, it selects an arm according to the probability of its mean being the largest. It then observes a set of rewards which it uses to update its probability distribution over the parameters.

For the case where each arm has an independent parameter defining its reward distribution, we can update the distribution of all arms observed separately. A particularly simple

case is when all the reward are generated from Bernoulli distributions. Then we can simply use a Beta prior for each arm, illustrated by the TS-N policy in Algorithm 1. We note that the algorithm trivially extends to other priors and families.

Algorithm 1 TS-N (Bernoulli case)

```

For each arm  $i$ , set  $S_i = 1$  and  $F_i = 1$ 
for all round  $t = 1, \dots, T$  do
  For each arm  $i$ , sample  $\theta_i$  from the Beta distribution  $\text{Beta}(S_i, F_i)$ 
  Play arm  $A_t = \arg \max_{i \in \mathcal{V}} \theta_i$ 
  for all  $k \in \mathcal{N}_{A_t}$  do
     $\hat{r}_k = \text{Bernoulli}(r_k)$ 
    If  $\hat{r}_k = 1$  the  $S_k = S_k + 1$ , else  $F_k = F_k + 1$ 
  end for
end for

```

4.2 The TS-MaxN policy

The TS-N policy does not fully exploit the graphical structure. For example, as noted by (Caron et al. 2012), instead of doing exploration on arm i we could explore an apparently better neighbour, which would give us the same information. More precisely, instead of picking arm i , we pick the arm $j \in \mathcal{N}_i$ with the best empirical mean. The intuition behind it is that, if we take any arm in \mathcal{N}_i , we are going to observe anyway the reward of i . So it is always better to exploit the best arm in \mathcal{N}_i . The resulting policy, TS-MaxN is summarized in Algorithm 2. Although our theoretical results do not apply to this policy, it can have better performance as it uses more information.

Algorithm 2 TS-MaxN

```

For each arm  $i$ , set  $S_i = 1$  and  $F_i = 1$ 
Let  $\bar{x}_i$  be the empirical mean of arm  $i$ 
for all round  $t = 1, \dots, T$  do
  For each arm  $i$ , sample  $\theta_i$  from the Beta distribution  $\text{Beta}(S_i, F_i)$ 
  Let  $j = \arg \max_{i \in \mathcal{V}} \theta_i$ 
  Play arm  $A_t = \arg \max_{k \in \mathcal{N}_j} \bar{x}_k$ 
  for all  $k \in \mathcal{N}_{A_t}$  do
     $\hat{r}_k = \text{Bernoulli}(r_k)$ 
    If  $\hat{r}_k = 1$  the  $S_k = S_k + 1$ , else  $F_k = F_k + 1$ 
  end for
end for

```

4.3 Analysis of TS-N policy

Russo and van Roy introduced an elegant approach to the analysis of Thompson sampling. They define the *information ratio* as a key quantity for analysing information structures:

$$\Gamma_t := \frac{\mathbb{E}_t [R(Y_{t,A^*}) - R(Y_{t,A_t})]^2}{\mathbb{I}_t(A^*, (A_t, Y_{t,A_t}))}, \quad (4.1)$$

where \mathbb{E}_t and \mathbb{I}_t denote expectation and mutual information respectively, conditioned on the history of arms and observations until time t . They show that it follows very generally that

Proposition 4.1. If $\Gamma_t \leq \Gamma$ almost surely for all $1 \leq t \leq T$, then, $\mathbb{E} \mathcal{L}(T, \pi^{TS}) \leq \sqrt{\Gamma \mathbb{H}(\alpha_1) T}$.

Here \mathbb{H} denotes entropy. Thus to analyse the performance of Thompson sampling on a specific problem, one may focus on bounding the information ratio (4.1). For the (independent) K -armed bandit case, they show that $\Gamma_t \leq K/2$, while for full-information (K experts) case, they show that $\Gamma_t \leq 1/2$. We now give a simple but useful extension of their results which is intermediate between these cases.

Proposition 4.2. Let \equiv be an equivalence relation defined on the arms with \bar{a} denoting the equivalence class of a . Let $Y_{t,a} = (a, Z_{t,\bar{a}})$ for sequence of random variables $Z_{t,\bar{a}} : \Omega \rightarrow \mathcal{Z}$. Then $\Gamma_t \leq \frac{1}{2} |K/\equiv|$, half the number of equivalence classes.

This is a direct generalisation of propositions 3 and 4 in (Russo and Roy 2016), to which it reduces when the equivalence relation is trivial (bandit case) or full (expert case).

We can now use Proposition 4.2 to analyse graph structured arms:

Lemma 4.1. Let $G = (\mathcal{V}, \mathcal{E})$ be a graph with V corresponding to the arms and suppose that when an arm a is played, we observe the rewards $R(Y_{t,a'})$ for all $a' \in N(a)$ i.e we observe the rewards corresponding to both a and all its neighbours. Let \mathcal{C} be a clique cover of G i.e. a partition of V into cliques. Then $\Gamma_t \leq \frac{1}{2} |\mathcal{C}|$.

Applying Proposition 4.1 and Lemma 4.1, we get a performance guarantee for Thompson sampling with graph-structured feedback:

Theorem 4.3. For Thompson sampling with feedback from the graph G , we have $\mathbb{E}^{\pi^{TS}} \mathcal{L} \leq \sqrt{\frac{1}{2} \chi(\bar{G}) \mathbb{H}(\alpha_1) T}$, where $\chi(\bar{G})$ is the clique cover number of G .

Remark 4.1. The bandit and expert cases are special cases corresponding to the empty graph and the complete graph respectively since $\chi(\bar{G}) = K$ for the empty graph and $\chi(\bar{G}) = 1$ for the complete graph.

Remark 4.2 (Planted Partition Models). The *planted partition models* or *stochastic block models* graphs $G(n, k, p, q)$ are defined as follows (McSherry 2001; Condon and Karp 2001): first a fixed partition of the n vertices into k parts is chosen, then an edge between two vertices within the same class exists with probability p and that between vertices in different classes exists with probability q , independently with $p > q$. If $p = 1$, then with high probability, the clique cover number of the resulting graph is k (corresponding to the planted k cliques). Thus for this class of graphs, the regret grows as $O(\sqrt{k})$ as per Theorem 4.3. This is explored in Section 5. When $p \neq 1$ but large, the planted partition graph is considered a good model of the structure of network communities.

If the underlying graph changes at each time step, then we also have the bound for the same algorithm:

Corollary 4.1. Suppose the underlying graph at time $t \geq 1$ is G_t , then:

$$\mathbb{E}^{\pi^{TS}} \mathcal{L} \leq \sqrt{\frac{1}{2} \max_t \chi(\bar{G}_t) \mathbb{H}(\alpha_1) T}$$

Proof. The information ratio at time t is bounded by $\chi(\bar{G}_t) \leq \max_t \chi(\bar{G}_t)$. \square

5 Experiments

We compared our proposed algorithms in terms of the actual expected regret against a number of other algorithms that can take advantage of the graph structure. Our comparison was performed over both synthetic graphs and networks derived from real-world data.²

5.1 Algorithms and hyperparameters.

In all our experiments, we tested against the UCB-MaxN and UCB-N algorithms, introduced in (Caron et al. 2012). These are the analogues of our algorithms, using upper confidence bounds instead of Thompson sampling.

ϵ -greedy- \mathcal{D} and ϵ -greedy-LP. For the real-world networks, we also evaluated our algorithms against a variant of ϵ -greedy-LP from (Buccapatnam, Eryilmaz, and Shroff 2014). This is based on a linear program formulation for finding a lower bound $\gamma(G)$ on the size of the minimum dominating set. We observe first that their analysis holds for any fixed dominating set D and the bound so obtained is $O(|D| \ln T)$. In particular, we may use a simple greedy algorithm to compute a near-optimal dominating set D' such that $|D'| \leq \gamma(G) \log \Delta$, where Δ is the maximum degree of the graph.³ Using such a near optimal dominating set in place of the LP relaxation and choosing arms from it uniformly at random, we obtain a variant of the original algorithm, which we call ϵ -greedy- \mathcal{D} , which is much more computationally efficient, and which enjoys a similar regret bound:

Theorem 5.1. The regret of ϵ -greedy- \mathcal{D} is at most $O(\gamma(G) \ln \Delta \ln T)$, where Δ is the maximum degree of the graph.

ϵ -greedy- \mathcal{D} and ϵ -greedy-LP have the hyper-parameters c and d , which control the amount of exploration. We found that its performance is highly sensitive to their choice. In our experiments, we find the optimal values for these parameters by performing a separate grid search for each problem, and only reporting the best results. Since there is no obvious way to tune these parameters online, this leads to a favourable bias in our results for this algorithm.⁴

As Thompson sampling is a Bayesian algorithm, we can view the prior distribution as a hyper-parameter. In our experiments, we always set that to a Beta(1,1) prior for all rewards.

5.2 General experimental setup.

For all of our experiments, we performed 210 independent trials and reported the *median-of-means* estimator⁵ of the

²Our source code and data sets will be made available on an open hosting website.

³No polynomial time algorithm can guarantee a better approximation unless P=NP, (Ruan et al. 2004)

⁴A similar observation was made in (Auer, Cesa-Bianchi, and Fischer 2002b), which noted that an optimally tuned ϵ -greedy performs almost always best, but its performance can degrade significantly when the parameters are changed. Although (Buccapatnam, Eryilmaz, and Shroff 2014) suggests a method for selecting these parameters, we find that using it leads to a near-linear regret.

⁵Used heavily in the streaming literature (Alon, Matias, and Szegedy 1996)

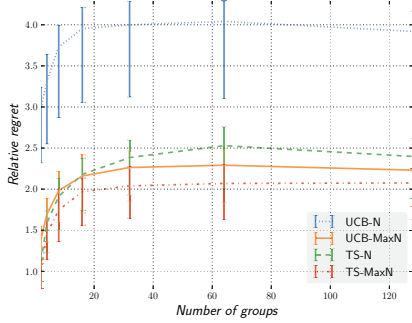


Figure 1: Relative regret in the planted partition setting as the number of groups increases.

cumulative regret. It partitions the trials into a_0 equal groups and return the median of the sample means of each group. We set the number of groups to $a_0 = 14$, so that the confidence interval holds with probability at least 0.955.

We also reported the deviation of each algorithm using the Gini’s Mean Difference (GMD hereafter) (Gini and Pearson 1912). GMD computes the deviation as $\sum_{j=1}^N (2j - N - 1)x_{(j)}$ with $x_{(j)}$ the j -th order statistics of the sample (that is $x_{(1)} \leq x_{(2)} \leq \dots \leq x_{(N)}$). As shown in (Yitzhaki and others 2003; David 1968) the GMD provides a superior approximation of the true deviation than the standard one. To account for the fact that the cumulative regret of our algorithms might not follow a symmetric distribution, we computed the GMD separately for the values above and below the *median-of-means*.

5.3 Simulated graphs

In our synthetic problems, unless otherwise stated, the rewards are drawn from a Bernoulli distribution whose mean is generated uniformly randomly in $[0.45, 0.55]$ except for the optimal arm whose mean is generated randomly in $[0.55, 0.6]$. The number of nodes in the graph is 500. We tested with a sparse graph of 2500 edges and also with a dense graph of 62625 edges.

Erdős–Rényi graphs In our first experiment, we generate the graph randomly using the Erdős–Rényi model. Figure 2e and 2f respectively show the result in the sparse and dense graph.

Our first observation here is that all policies take advantage of a large number of edges as their cumulative regret is better by using the dense graph (Figure 2f) rather than the sparse one (Figure 2e). This confirms the theoretical result as a dense graph will have a smaller clique cover number than a sparse one.

The policy TS-MaxN outperforms all other in both the sparse and dense graph model. However, the performance of TS-N is very close to that of TS-MaxN in the near complete graph. This is explained by the fact that in a near complete graph we have many cliques. It is revealing to see that TS-N outperforms both the UCB-N and UCB-MaxN policies.

Power Law graph Such graphs are commonly used to generate static scale-free networks (Goh, Kahng, and Kim 2001). In this experiment, we generated a non-growing random graph with expected power-law degree distribution.

show the results respectively for the dense and sparse graph Figure 2d and 2c show the results respectively for the dense and sparse graph. Again, the policy TS-MaxN clearly outperforms all other. In the sparse graph model, TS-N is beaten by UCB-MaxN at the beginning of the rounds ($t \leq 100000$), but catches and ended up beating UCB-MaxN.

Planted Partition Model The aim of the experiment on this model is to check the dependency on the number of cliques for each policy. Figure 1 shows the results where on the x-axis we have the parameter k of the planted partition graph (which is almost equal to the number of cliques) on a graph with 1024 nodes. On the y-axis we have the relative regret of each policy, i.e. the ratio between the regret of each policy with the regret of the best policy when there are two groups, for ease of comparison. As we can see, all methods’ regret scales similarly. Thus, the theoretical bounds appear to hold in practice, and to be somewhat pessimistic. For a larger number of nodes, we would expect the plots to flatten later.

5.4 Social networks datasets

Our experiments on real world datasets follow the methodology described in (Caron et al. 2012). We first infer a graph from data, and then define a reward function for movie recommendation from user ratings. Missing ratings are predicted using matrix factorization. This enables us to generate rewards from the graph. We explain the datasets, reward function and graph inference in the full version.

Results Figure 2a shows the results for the Facebook graph and Figure 2b for the Flixster graph. Once again, the Thompson sampling strategies dominate all other strategies for the Facebook and they are matched by the optimised ϵ -greedy- \mathcal{D} policy in the Flixster graph. We notice that in this setting the gap between the UCB policies and the rest is much larger, as is the overall regret of all policies. This can be attributed to the larger size of these graphs.

6 Conclusion

We have presented the first Thompson sampling algorithms for sequential decision problems with graph feedback, where we not only observe the reward of the arm we select, but also those of the neighbouring arms in the graph. Thus, the graph feedback allows us the flexibility to model different types of feedback information, from bandit feedback to expert feedback. Since the structure of the graph need not be known in advance, our algorithms are directly applicable to problems with changing and/or unknown topology. Our analysis leverages the information-theoretic construction of (Russo and Roy 2016), by bounding the expected information gain in terms of fundamental graph properties. Although our problem-independent bound of is not directly comparable to (Caron et al. 2012), we believe that a problem-independent

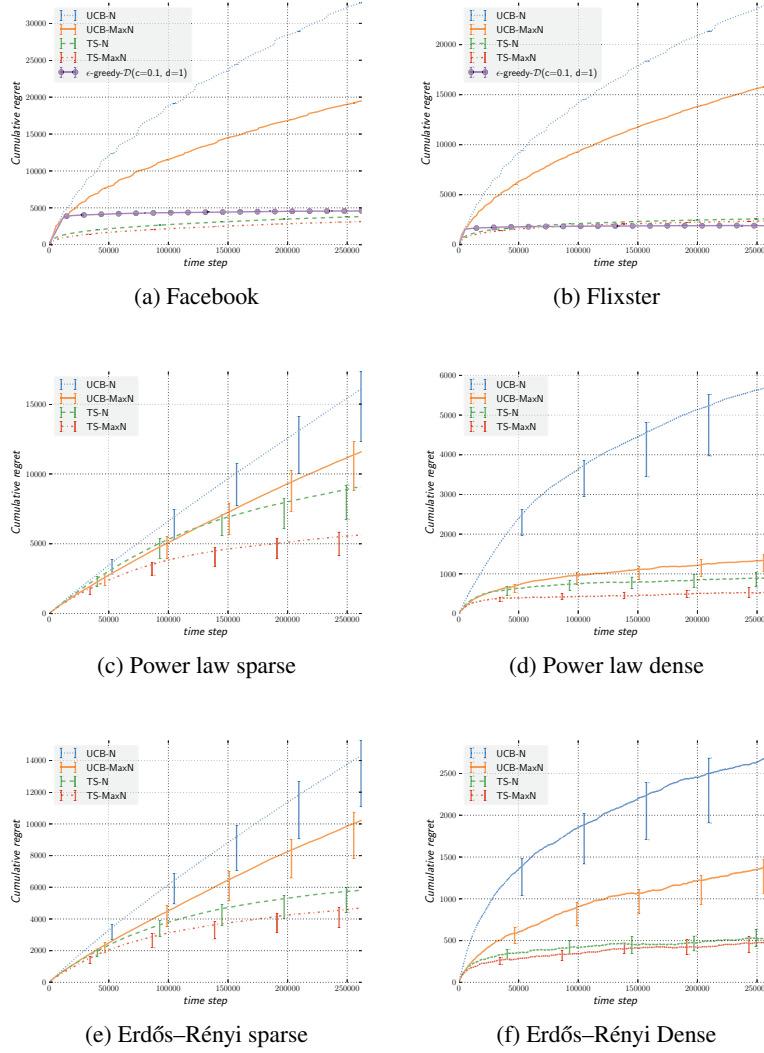


Figure 2: Regret and error bar on simulated (sparse and dense) and real social network graphs

version of the latter should be $O(\sqrt{\chi \ln T})$, in which case our results would represent an improvement of $O(\sqrt{\chi})$.

In practice, our two variants always outperform UCB-N, UCB-MaxN, which also use graph feedback but rely on upper confidence bounds. We are also favourably compared against ϵ -greedy- \mathcal{D} , even when we tune the parameters of the latter *post hoc*.

It would be interesting to extend our techniques to other types of feedback. For example, the Bayesian foundations of Thompson sampling render our algorithms applicable to arbitrary dependencies between arms. In future work, we will analytically and experimentally consider such problems and related applications. Finally, an open question is the existence of information-theoretic lower bounds in settings with partial feedback.

Acknowledgments. This research was partially supported by the SNSF grants “Adaptive control with approximate Bayesian computation and differential privacy” (IZK0Z2_167522), “Swiss Sense Synergy” (CRSII2_154458), by the the People Programme (Marie Curie Actions) of the European Union’s Seventh Framework Programme (FP7/2007-2013) under REA grant agreement number 608743, and the Future of Life Institute.

References

- Agrawal, S., and Goyal, N. 2012. Analysis of thompson sampling for the multi-armed bandit problem. In *COLT 2012*.
- Alon, N.; Cesa-Bianchi, N.; Dekel, O.; and Koren, T. 2015. Online learning with feedback graphs: beyond bandits. In *Proceedings of the 28th Annual Conference on Learning Theory*, 23–35.

- Alon, N.; Matias, Y.; and Szegedy, M. 1996. The space complexity of approximating the frequency moments. In *28th STOC*, 20–29. ACM.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002a. Finite time analysis of the multiarmed bandit problem. *Machine Learning* 47(2/3):235–256.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002b. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47(2-3):235–256.
- Buccapatnam, S.; Eryilmaz, A.; and Shroff, N. B. 2014. Stochastic bandits with side observations on networks. *ACM SIGMETRICS Performance Evaluation Review* 42(1):289–300.
- Burnetas, A. N., and Katehakis, M. N. 1997. Optimal adaptive policies for markov decision processes. *Mathematics of Operations Research* 22(1):222–255.
- Caron, S.; Kveton, B.; Lelarge, M.; and Bhagat, S. 2012. Leveraging side observations in stochastic bandits. *UAI*.
- Cesa-Bianchi, N., and Lugosi, G. 2006. *Prediction, Learning and Games*. Cambridge Press.
- Chapelle, O., and Li, L. 2011. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, 2249–2257.
- Cohen, A.; Hazan, T.; and Koren, T. 2016. Online learning with feedback graphs without the graphs. In *ICML 2016*.
- Condon, A., and Karp, R. M. 2001. Algorithms for graph partitioning on the planted partition model. *Random Structures and Algorithms* 18(2):116–140.
- David, H. 1968. Miscellanea: Gini’s mean difference rediscovered. *Biometrika* 55(3):573–575.
- Garivier, A., and Cappé, O. 2011. The kl-ucb algorithm for bounded stochastic bandits and beyond. *arXiv preprint arXiv:1102.2490*.
- Gini, C., and Pearson, K. 1912. *Variabilità e mutabilità: contributo allo studio delle distribuzioni e delle relazioni statistiche. Fascicolo I*. tipografia di Paolo Cuppini.
- Goh, K.-I.; Kahng, B.; and Kim, D. 2001. Universal behavior of load distribution in scale-free networks. *Physical Review Letters* 87(27):278701.
- Lai, T. L., and Robbins, H. 1985. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics* 6(1):4–22.
- Mannor, S., and Shamir, O. 2011. From bandits to experts: On the value of side-observations. In *Advances in Neural Information Processing Systems*, 684–692.
- McSherry, F. 2001. Spectral partitioning of random graphs. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, 529–537. IEEE.
- Ruan, L.; Du, H.; Jia, X.; Wu, W.; Li, Y.; and Ko, K.-I. 2004. A greedy approximation for minimum connected dominating sets. *Theoretical Computer Science* 329(1):325–330.
- Russo, D., and Roy, B. V. 2016. An information-theoretic analysis of thompson sampling. *Journal of Machine Learning Research* 17(68):1–30.
- Scott, S. L. 2010. A modern bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry* 26(6):639–658.
- Scott, S. L. 2015. Multi-armed bandit experiments in the on-line service economy. *Applied Stochastic Models in Business and Industry* 31(1):37–45.
- Thompson, W. 1933. On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of two Samples. *Biometrika* 25(3-4):285–294.
- Yitzhaki, S., et al. 2003. Gini’s mean difference: A superior measure of variability for non-normal distributions. *Metron* 61(2):285–316.