
Adaptive Sampling for Discovery

Ziping Xu

Department of Statistics
University of Michigan
zipingxu@umich.edu

Eunjae Shim

Department of Chemistry
University of Michigan
eunjae@umich.edu

Ambuj Tewari

Department of Statistics
University of Michigan
tewaria@umich.edu

Paul Zimmerman

Department of Chemistry
University of Michigan
paulzim@umich.edu

Abstract

In this paper, we study a sequential decision-making problem, called Adaptive Sampling for Discovery (ASD). Starting with a large unlabeled dataset, algorithms for ASD adaptively label the points with the goal to maximize the sum of responses. This problem has wide applications to real-world discovery problems, for example drug discovery with the help of machine learning models. ASD algorithms face the well-known exploration-exploitation dilemma. The algorithm needs to choose points that yield information to improve model estimates but it also needs to exploit the model. We rigorously formulate the problem and propose a general information-directed sampling (IDS) algorithm. We provide theoretical guarantees for the performance of IDS in linear, graph and low-rank models. The benefits of IDS are shown in both simulation experiments and real-data experiments for discovering chemical reaction conditions.

1 Introduction

Machine Learning (ML) models are becoming increasingly popular for discovery problems in various areas like drug discovery [40, 5, 38] and scientific discoveries [13, 23, 19]. Despite its success in real-data applications, the theoretical aspects of the problem are not fully understood in the machine learning literature. In this paper, we consider a particular case of discovery and formalize it as a problem that we call Adaptive Sampling for Discovery (ASD). In ASD, an algorithm sequentially picks T samples X_1, \dots, X_T out of a size- n unlabeled dataset S^n ($T \ll n$) and obtains their labels Y_1, \dots, Y_T with the goal to maximize the sum $\sum_{t=1}^T Y_t$. The labels correspond to the quality of discoveries, which can be continuous or ordered categorical variables depending on the types of discoveries. As a sequential decision-making problem, we find it substantially connected to the existing literature on Multi-arm Bandit (MAB) and other related decision-making problems, which we briefly review in the rest of this section.

Multi-arm bandit. MAB is a decision-making problem that maximizes the total rewards by sequentially pulling arms and each arm corresponds to a reward distribution [27]. MAB faces the well-known exploration-exploitation trade-off dilemma. ASD has to deal with the same exploration-exploitation dilemma if one treats each arm as an unlabeled point. The major difference between ASD and MAB is that each arm after being pulled once can still be pulled for infinite number of times, while we can only label a sample once in ASD. This is because a discovery can only be made once: one receives no credit making the same discovery twice.

In fact, as we will show later, ASD can be an MAB when each unique value in S^n has larger than T repeated samples and $n \ll T^2$. Moreover, a variant of MAB called Sleeping Expert [22, 24] can be seen as a more general problem setup of ASD. One can intuitively interpret Sleeping Expert as the bandit problem where the set of available arms varies. More generally, bandit algorithms that allows adversarial choice of the action set can be used to solve the ASD problem. However, due to its generality, the regret bounds developed in [24] are loose or vacuous in our setup. Detailed discussions are deferred to Section 2.

Due to the similarities between bandits and ASD, we adopt the same performance measurement, regret, from bandit literature. Regret is the highest sum of labels that could have been achieved minus the actual sum of labels.

There is a long line of work in the bandit literature on solving the exploration-exploitation tradeoff. Popular methods includes UCB (upper confidence bound), TS (Thompson sampling). UCB keeps an optimistic view on the uncertain model [4, 26], which encourages the algorithm to select arms with higher uncertainty. TS [1] adapts a Bayesian framework, which samples a reward vector from its posterior distribution and select arms accordingly. In this paper, we adopt another family of methods called Information-directed Sampling (IDS). IDS [33] balances between the expected information gain on the unknown model and instant regret with respect to the posterior distribution under a Bayesian framework. IDS has been shown to outperform UCB in problem with rich structural information, e.g. sparse linear bandit [14]. As we show below (Proposition 1), the ASD problem is not interesting in the unstructured case which means we have to use structural assumptions. That makes IDS a natural choice.

Traditional MAB algorithms designed for finite-armed bandits are not applicable to ASD since the number of unlabeled samples n can be much larger than horizon T . There are works considering infinite-armed bandits by assuming structural information across arms. For instance, linear bandit assumes a linear model for the reward generation [27], [41] assumes Lipschitz-continuity and [42] considers neural network model for reward generation. However, there are few works studying sleeping expert with structure across arms beyond linear structure.

Other related literature. With a similar adaptive sampling procedure, Active Learning (AL) aims at achieving a better *model accuracy* with a limited number of labels. As the reader may have noticed, the goal of AL aligns with ASD in the early stage when the model has high uncertainty and the predictions are high unreliable, while in the later stage ASD aims at better discovery performance. Indeed, there are active learning algorithms based on the idea of maximum information gain [3]. [29] applies the idea to the matrix completion problem, which significantly improves the prediction accuracy from random selection. Other works with a similar goal go by the name *sequential exploration* [6, 7, 17]. Their works lack a theoretical justification and the algorithms are not generally applicable.

A more relevant setup in the active learning literature is Active Search (AS). Active search is an active learning setting with the goal of identifying as many members of a given class as possible under a labeling budget. AS concerns the Bayesian optimal methods named Efficient Nonmyopic Search (ENS) [21, 20, 32], which is not computationally efficient. Their approximate algorithms for the Bayesian optimal method do not provide strong theoretical guarantee. A more detailed comparison between our proposed approach and ENS is given in Appendix I.

Main contributions. In this paper, we formulate the adaptive sampling for discovery problem and propose an generic algorithm using information-direct sampling (IDS) strategy. We theoretically analyze the regret of IDS under generalized linear model, low-rank matrix model and graph model assumptions. Indeed, the analysis for linear model are directly from IDS for linear bandit. The regret analysis for low-rank model and graph model are new even in the bandit literature. The results are complemented by simulation studies. We apply the algorithms to the real-world problems in reaction condition discovery, where our algorithm can discover about 20% more plausible reaction conditions than random policies and other simple baseline policies adapted from the bandit literature.

2 Formulation

In this section, we formally formulate ASD and rigorously discuss its connections to MAB.

Notations. We first introduce some notations that are repeatedly used in this paper. For any finite set S , we let $\mathcal{D}(S)$ be the set of all the distributions over the set. For any positive integer n , we let $[n] = \{1, \dots, n\}$. We use the standard $O(\cdot)$ and $\Omega(\cdot)$ notation to hide universal constant factors. We also use $a \lesssim b$ and $a \gtrsim b$ to indicate $a = O(b)$ and $a = \Omega(b)$. ASD is a sequential decision-making problem over discrete time steps. In general, we let \mathcal{F}_t be the observations up to decision step t . We adapt a Bayesian framework and let $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_t]$ be the conditional expectation given the history up to t . Let \mathbb{P}_t be the probability measure of posterior distribution.

Adaptive sampling for discovery. We formulate ASD problem as follows. Consider a problem with the covariate space $\mathcal{X} \subset \mathbb{R}^d$ and label space $\mathcal{Y} \subset [0, 1]$. For categorical discoveries, the labels are ordered such that higher labels reflects better discoveries. Given a input $x \in \mathcal{X}$, its label is generated from a distribution $\mathcal{D}_{\theta|x}$ over \mathcal{Y} parametrized by unknown parameters θ . We adopt a Bayesian framework and denote the prior distribution of θ by ϕ . Given a sequence of $t - 1$ observations $\mathcal{F}_t = \{(X_i, Y_i)\}_{i=1}^{t-1}$, the posterior distribution of ϕ is denoted by $\phi(\cdot | \mathcal{F}_t)$. We also let $f_{\theta}(x) = \mathbb{E}_{Y \sim \mathcal{D}_{\theta|x}}[Y]$ be the expected label of input x under model θ .

ASD starts with an unlabeled dataset $S^n \subset \mathcal{X}$ with n samples. Our goal is to adaptively label T samples $\{X_1, \dots, X_T\} \subset S^n$ to maximize the sum of labels $\sum_{t=1}^T Y_t$. The discovered set is removed from available unlabeled set, i.e. each arm can only be pulled once. At each step t , an algorithm chooses an input X_t from $S_t^n := S^n \setminus \{X_1, \dots, X_{t-1}\}$. The environment returns the label $Y_t \sim \mathcal{D}_{\theta|X_t}$.

To measure the performance of any algorithm, we borrow the definition of Bayesian regret from bandit literature [27]. Note that the regret definition is also used in Sleeping Expert [22], which is a more general framework than ASD. With respect to any chosen random sequence (X_1, \dots, X_T) , we define regret by

$$R_T = \max_{x_1, \dots, x_T \in S^n} \sum_{t=1}^T f_{\theta}(x_t) - \sum_{t=1}^T f_{\theta}(X_t). \quad (1)$$

Let $S_t^n = S^n \setminus \{X_1, \dots, X_{t-1}\}$ be the remaining unlabeled set at the start of step t . An algorithm \mathcal{A} is a sequence of policies (π_1, \dots, π_T) . Let $\mathcal{D}(S_t^n) \subset \mathbb{R}^{n-t-1}$ be the set of all the distributions over S_t^n that are denoted as $n - t - 1$ vectors. Each π_t is a mapping from history \mathcal{F}_t to a distribution $\mathcal{D}(S_t^n)$ over S_t^n . The Bayesian Regret w.r.t. an algorithm is

$$\mathcal{BR}(T, \mathcal{A}) := \mathbb{E}R_T,$$

where the expectation is taken over the randomness in the choice of input X_t , the labeling Y_t and over the prior distribution on θ . Let X_1^*, \dots, X_T^* be the sequence that realizes the maximum in (1). Note that X_1^*, \dots, X_T^* are random variables depending on θ . In general, the goal the algorithm is to achieve a sublinear regret upper bound.

Connections to bandit problem. This problem degenerates to a classical bandit problem when \mathcal{X} is discrete and S^n has infinite number of unlabeled samples for each unique value in \mathcal{X} . In the bandit literature some frequentist results are available. The problem has the minimax regret rate $\sqrt{|\mathcal{X}|T}$. However, in our problem, we generally consider large unlabeled set ($n \gg T$). In fact, when $|\mathcal{X}|$ is larger than T^2 , Proposition 1 shows that a linear regret is unavoidable without a structural assumption. Therefore, we discuss different types of structural information that allows information sharing across different inputs.

Proposition 1 (Lower bound for nonstructural case). *Consider the following set of ASD problems. Let $n = T^2$, $S^n = [n]$ and $\mathcal{D}_{\theta|x} = \mathcal{N}(\theta_x, 1)$ for $x \in [n]$ and $\theta \in [0, 1]^n$. For any algorithm \mathcal{A} , there exists some prior distribution over the set of ASDs described above such that $\mathcal{BR}(T, \mathcal{A}) \gtrsim T$.*

Compared with sleeping expert. Sleeping expert considers a generalized bandit framework where the available (awake) action set \mathcal{A}_t are changing. It is the same as ASD except for that the changes of the unlabeled dataset (or awake action set) follow certain dynamic. That is whenever an input is labeled, it is taken out of the unlabeled dataset. However, sleep expert normally assumes a stochastic distribution on \mathcal{A}_t or adversarial changes [24, 9], which makes their regret bound loose. To be specific, [22] derives regret bounds for unstructured bandits, which leads to vacuous regret bound of \sqrt{nT} . For structured case, [24] derives the same regret bound in this paper, due to the strong

structure assumption for linear model. There are few literature for sleeping bandits with complicated structures beyond linear model. The regret bounds in the bandit literature for the other two models, graph and low-rank matrix models, are not applicable both resulting a $\Omega(T)$ regret bound, which assures the necessity of studying ASD problem by itself.

3 Information-directed Sampling and Generic Regret Bound

In this section, we introduce a strategy called Information-Directed Sampling and develop a generic regret bound. IDS runs by evaluating the information gain on the indices of the top T inputs with respect to the labeling of a certain input x . It balances between obtaining low expected regret in the current step and acquiring high information about the inputs that are worth labeling. We first introduce *entropy* and *information gain*.

Definition 1 (Entropy and information gain). *For any random variable $X \in \mathcal{X}$ with a density function \mathbb{P} , its entropy is defined as $H(X) = -\int_{x \in \mathcal{X}} P(x) \log(P(x))$. Let $H(X | Z)$ be the conditional entropy of X given Z . The mutual information between any two random variables Y and X can be defined as $I(X; Y) = H(X) - H(X | Y)$. Conditional mutual information is $I(X; Y | Z) = H(X | Z) - H(X | Y, Z)$. Furthermore, we use H_t and I_t for the entropy and mutual information under the posterior distribution up to step t .*

For IDS, we let $\Delta_t(x) = \mathbb{E}_t[\max_{x' \in S_t^n} f_\theta(x') - f_\theta(x)]$ be the expected instant regret. Let the information gain of selecting x at the step t be $g_t(x) = I_t(X_{t,1}^*, \dots, X_{t,T-t+1}^*; Y_t | \mathcal{F}_{t-1}, X_t = x)$, where $(X_{t,1}^*, \dots, X_{t,T-t+1}^*)$

$$= \left\{ (x_1, \dots, x_{T-t+1}) : x_i \in S_t^n, f_\theta(x_1) \geq \dots \geq f_\theta(x_{T-t+1}) \geq \max_{x' \neq x_1, \dots, x_{T-t+1}} f_\theta(x') \right\},$$

are the random variable for the top $T - t + 1$ unlabeled points. Note that $\{X_{t,1}^*, \dots, X_{t,T-t+1}^*\} \subset \{X_1^*, \dots, X_T^*\}$. Intuitively, $g_t(x)$ captures the amount of information on the indices of the top $T - t + 1$ points gained by labeling x .

IDS minimizes the following ratio

$$\pi_t^{\text{IDS}} \in \arg \min_{\pi \in \mathcal{D}(S_t^n)} \Psi_{t,\lambda}(\pi) := \frac{(\Delta_t^T \pi)^\lambda}{g_t^T \pi}, \text{ for some constant } \lambda > 0,$$

where $\Delta_t \in \mathbb{R}^{n-t+1}$ and $g_t \in \mathbb{R}^{n-t+1}$ are the corresponding vectors for the expected single-step regret and information gain. Constant ratio here weighs between instant regret and information gain. A higher λ prefers points with lower instant regrets. An abstract algorithm is given Algorithm 1.

Algorithm 1 IDS (Information-Directed Sampling) for Discovery

Input: Unlabeled dataset S^n , prior distribution ϕ , total number of steps T , constant λ .
Initialize history $\mathcal{F}_0 = \{\}$, $S_1^n = S^n$.
for $t = 1$ **to** T **do**
 Calculate $\Delta_t(x)$ and $g_t(x)$ for each $x \in S_t^n$ using posterior $\phi(\cdot | \mathcal{F}_{t-1})$.
 Sample $X_t \sim \arg \max_{\pi} \Psi_{t,\lambda}(\pi)$ and receive label Y_t from $\mathcal{D}_{\theta|X_t}$.
 Update history $\mathcal{F}_t = \mathcal{F}_{t-1} \cup \{(X_t, Y_t)\}$ and unlabeled dataset $S_{t+1}^n = S_t^n \setminus \{X_t\}$.
end for

We can show that the following generic Bayesian regret upper bound holds.

Lemma 1. *Suppose $\pi^{\text{IDS}} = (\pi_t)^{t \in [T]}$ and $\Psi_{t,\lambda}(\pi_t) \leq \Psi_{*,\lambda}$. Then we have IDS using constant λ has the following regret bound*

$$\text{BR}(T, \text{IDS}) \lesssim (\Psi_{*,\lambda} H(X_1^*, \dots, X_T^*) T^{\lambda-1})^{1/\lambda}.$$

As the case in the bandit problem, the regret bound depends on the worst-case information ratio bound and the entropy of the random variables for the top T unlabeled points. Now we will discuss how to bound information ratios under various structural information assumptions.

4 Bounding Information Ratio under Structural Assumptions

In this section, we study three cases with different structural information.

4.1 Generalized linear model.

In this subsection, we discuss a generalized regression model, which includes logistic regression, corresponding to a natural binary discovery problem. Consider d -dimensional generalized linear regression problem. The label Y given an input X is generated by

$$Y = \mu(X^\top \theta^*) + \epsilon,$$

where μ is usually referred as link function. We assume that Y is uniformly bounded in $[0, 1]$. In addition, we assume an upper bound on the derivatives of μ .

Assumption 1. *The first order derivatives of μ are upper-bounded by L_μ .*

Theorem 1. *Under Assumption 1, we have $\Psi_{*,2} \leq L_\mu d/2$, which gives the Bayesian regret bound*

$$\mathcal{BR}(T, IDS) \lesssim \sqrt{dH(X_1^*, \dots, X_T^*)L_\mu T} \quad (2)$$

A naive bound for $H(X_1^*, \dots, X_T^*)$ in [33] does not apply here because it introduces a $\log(n)$ dependency. Instead, we notice that X_1^*, \dots, X_T^* is a deterministic function of θ and S^n . We can bound the entropy of θ instead. Assuming a Gaussian prior on θ , $H(\theta) \lesssim d \log(d)$.

Indeed, the regret in (2) is analogous to the regret bound in linear bandit [33, 27]. Theoretically, IDS does not show an advantage. This is partly because carefully designed algorithms for linear model are sufficiently utilizing the structural information efficiently. [14] has shown that IDS achieves a much tighter lower bound for sparse linear bandit compared to UCB-type algorithms and Thompson sampling algorithms. We show in the Appendix G that IDS enjoys a regret bound of $\tilde{O}(sT^{2/3})$ in the discovery setting with a sparse linear structure, with s being the sparsity.

4.2 Low-rank matrix

In this subsection, we discuss a novel problem setup: low-rank matrix discovery. Low-rank matrix completion problem has been extensively studied in the literature [25, 18, 31]. However, most of the algorithms select entries with a uniform random distribution. In many applications, appropriately identifying a subset of missing entries in an incomplete matrix is of paramount practical importance. For example, we can encourage users to rate certain movies for movie rating matrix recovery. While selecting the users that reveal more information is important, it is also important to push the movies aligning with users tastes. Active matrix completion has long been studied in the literature [8, 29, 15]. As far as we know, no work has considered a discovery setup, where the goal of adaptive sampling is not only to get a better recovery accuracy but also to maximize the sum of observed entries.

Formulation. Consider an unknown matrix $Y \in \mathbb{R}^{m_1 \times m_2}$, whose entries $(i, j) \in [m_1] \times [m_2]$ are sampled in the following manner:

$$Y_{i,j} = e_i^T M_{i,j} e_j + \epsilon_{i,j} \text{ for } \epsilon_{i,j} \text{ being the noise for entry } i, j,$$

where $M \in \mathbb{R}^{m_1 \times m_2}$ is the unknown rank- r matrix and $r \ll \min\{m_1, m_2\}$. The matrix M admits SVD, i.e. $M = UDV^T$, where $U \in \mathbb{R}^{m_1 \times r}$ and $V \in \mathbb{R}^{r \times m_2}$ are orthogonal matrix and $D \in \mathbb{R}^{r \times r}$. For simplicity, we let $m_1 = m_2 = m$. Our analysis can be easily extended to the case $m_1 \neq m_2$.

Note that this setup is closely connected to the low-rank bandit problem except for that the actions are standard basis and can only be selected once. The best results in the literature [28] achieve a regret bound of $\tilde{O}(m^{3/2}\sqrt{rT})$. Since in our case, $T \leq m^2$, the regret bound becomes vacuous. This finding implies that we need some further assumptions on the structure of the matrix. A common assumption in the matrix completion is *incoherence*, an important concept that measures how the subspace aligns with the standard basis.

Definition 2 (Coherence). *The coherence of subspace U with respect to the standard i -th basis vector e_i is defined as $\mu_i(U) = \|P_U e_i\|_2^2$, where P_U is the orthogonal projection onto the subset defined by U . Note that $\mu_i(U) = \|U^T e_i\|_2^2$ as well.*

Assumption 2 (Incoherence). *There exists some constant $\gamma > 0$, such that the unknown matrix M is incoherent, i.e. $\max_{i \in [m]} \mu_i(U) \leq \sqrt{\gamma r/m}$ and $\max_{i \in m} \mu_i(V) \leq \sqrt{\gamma r/m}$.*

To introduce our results, we further define some constants.

Remark 1. *Let \bar{d} be the maximum singular value of the unknown matrix M . We have*

$$\max_{i,j} |M_{i,j}| \leq \max_{i,j} \mu_i(U) \mu_j(V) \|D\|_F \leq \gamma r^{3/2} \bar{d}/m := B.$$

Theorem 2. *Under the low-rank matrix model with Assumption 2, the worst-case information ratio can be bounded by $\Psi_{*,3} \leq 4B(B^2 + 1)r^3\gamma^2$, which gives us a regret bound of*

$$B\mathcal{R}(T, \text{IDS}) \lesssim (4B(B^2 + 1)r^3\gamma^2 H(M)T^2)^{1/3}.$$

The entropy term $H(M)$ depends on the distribution of M . To give an example, let $M = UV^T$, where $U, V^T \in \mathbb{R}^{m \times r}$ and each element in U and V are sampled independently from standard Gaussian. Then $H(M) \leq H((U, V)) = H(U) + H(V) \lesssim mr \log(mr)$. In general, the regret bound is still sublinear when $mr^4 \ll T \ll m^2$.

4.3 Graph

We consider the discovery problem with graph feedback. Specifically, we are given a graph $G = (S^n, E)$, where the unlabeled set S^n is the node set. By labeling a node $x \in S^n$, we receive noisy outcomes for all the nodes connecting to x . Let the side information at step t be $O_t(X_t) = \{\tilde{Y}_{x'}\}_{x':(x', X_t) \in E}$, where $\tilde{Y}_{x'} = x' + \epsilon_{x'}$ for $\epsilon_{x'}$ being the zero-mean noise given x being selected. In addition to the side information, we also receive the label for X_t , which is sampled in the same manner. This setup is also studied in the sleeping expert literature [9]. [9] gives a regret bound depending on $\mathbb{E}[\sum_x T_x/Q_x]$, where T_x is the total number of visits on action x and Q_x is the total number of observations for action x . The ratio can be low when the algorithms tend to select nodes with high degrees and their algorithms are not designed for that purpose. Therefore, the structure are not sufficiently exploited and the regret bound can be loose.

To measure the complexity of a graph, we introduce the following two definitions.

Definition 3 (Maximal independent set). *An independent set is a set of vertices in a graph such that no two of which are adjacent. A maximal independent set (MIS) is an independent set that is not a subset of any other independent set. We denote the cardinality of the smallest maximum independent set of a graph G by $\mathcal{C}(G)$.*

Definition 4 (Clique cover number). *A Clique of a graph $G = (\mathcal{K}, \mathcal{E})$ is a subset $S \subseteq \mathcal{K}$ such that the sub-graph formed by S and \mathcal{E} is a complete graph. A Clique cover of a graph $G = (\mathcal{K}, \mathcal{E})$ is a partition of \mathcal{K} , denoted by \mathcal{Q} , such that Q is a clique for each $Q \in \mathcal{Q}$. The cardinality of the smallest clique cover is called the clique cover number, which is denoted by $\chi(G)$.*

Remark 2. *Clique cover number is guaranteed to be larger than the cardinality of the smallest maximal independent set, i.e. $\mathcal{C}(G) \leq \chi(G)$.*

Since the nodes in the graph are diminishing, we further define a stable version of MIS.

Definition 5. *Let $\mathcal{G}_t(G)$ be the set of all subgraphs of G with $N - t$ nodes and define $\mathcal{C}_t(G) = \max_{G_t \in \mathcal{G}_t(G)} \mathcal{C}(G_t)$.*

Assumption 3. *We assume that $x + \epsilon_x \leq B$ almost surely for all $x \in S^n$ and some constant $B > 0$.*

Theorem 3. *Under Assumption 3, the Bayesian regret of the IDS with $\lambda = 2$ can be upper bounded by*

$$B\mathcal{R}(T, \text{IDS}) \lesssim \min\{(BC_T(G)T)^{2/3}, (\mathcal{X}(G)T)^{1/2}\}.$$

Let us discuss certain graphs where the regret bound can be low. First, if the graph can be decomposed into $K \ll T$ complete subgraphs, then $\mathcal{X}(G) \leq K$ and we have a regret bound of \sqrt{KT} . An example, in which $\mathcal{C}(G)$ is low while $\mathcal{X}(G)$ is high, is a star-shaped graph, with a set of nodes in the center of the graph connecting to all the other nodes. See Figure 2 in the Appendix for an illustration.

4.4 A generic results for models with structural information

As shown in sections 4.2 and 4.3, a $T^{2/3}$ regret upper bound with a small coefficient can be derived when the model has certain structural information. We give the following generic result, stating that as long as there exists a policy μ , whose information gain can be lower bounded by the instant regret of Thompson sampling policy an upper bound for $\Psi_{*,3}$ can be derived.

Proposition 2. *Assume that there exists a policy μ such that $g_t^\top \mu \geq \phi(\Delta_t^\top \pi_t^{ts})^2$, where π_t^{ts} is the Thompson sampling policy at the step t for some constant ϕ and assume that the instant regrets are uniformly bounded by B . Then $\Psi_{*,3} \leq 2B/\phi$, which gives a regret bound of $\mathcal{O}((BTH(\theta)/\phi)^{2/3})$.*

5 Experiments

We start with simulation studies on the three problems: (generalized) linear model, low-rank matrix and graph model. Since information gain can be hard to calculate in some problems, we introduce sampling-based approximate algorithms, which generate random samples from the posterior distribution and evaluate a variance term instead of the original information gain. We will also introduce general Thompson Sampling algorithms for all the three problem setups.

5.1 Approximate algorithms

It may not be efficient to evaluate the information gain in certain problem setups. For example, one way to generate random low-rank matrix is to generate its row and column spaces from Gaussian distribution. We do not have a closed form for the posterior distribution of the resulting matrix. Neither can we evaluate the information gain. To that end, we follow an approximate algorithmic design in [33], which replaces the information gain with a conditional variance $v_t(x) = \text{Var}_t(\mathbb{E}[Y_{t,x} | X_1^*, X_t = x])$, which will be evaluated under random samples from posterior distribution using MCMC. The conditional variance is a lower bound of the information gain. An approximate algorithm is given by Algorithm 2 in Appendix H.

Proposition 3. *The following lower bound for information gain holds, $g_t(x) \geq v_t(x)$.*

Compared algorithms. For the rest of section, we compare IDS (the approximate algorithm), TS, UCB (if the bandit version is available) and random policy. As a method that is closely connected to IDS, Thompson sampling (also referred as posterior sampling) is also compared in our experiments. Thompson sampling [36] at each round samples a θ_t from its posterior distribution and then selects the input with lowest instant regret w.r.t to θ_t . TS does not account for the amount of information gain, thus can be insufficient in utilizing the structural information. UCB is an important algorithm design in bandit literature. Though it is not specifically designed for ASD, we apply UCB to (generalized) linear ASD by manually removing arms that have been selected.

5.2 Simulation studies

Linear and logistic regression model. The data generating distribution are specified below. Unlabeled dataset is sampled from $\mathcal{N}(0, \sigma_x^2)$ and the prior distribution of θ^* is $\mathcal{N}(0, \sigma_\theta^2)$. The noise ϵ is sampled from $\mathcal{N}(0, \sigma_\epsilon^2)$. We study the effects of $d = 20, 50, 100$. The same setup is used for logistic regression model. The posterior distribution for θ in simple linear regression is also Gaussian. For logistic regression, we used a Laplace approximation for the posterior distribution [16]. We use the GLM-UCB [12] for logistic regression.

Graph model. We consider different types of graphs. For each graph, the node values are sampled from standard Gaussian distribution and the noise are sampled from Gaussian with zero mean and σ_ϵ^2 variance. We tested $\sigma_\epsilon^2 = 0.1, 1.0, 10$. We experiment on graphs with $N = 900$. For a better demonstration of the early stage performance, we only show the first 50 steps. We experiment the following types of graphs. 1) Random connections: any node has a probability of $p (= 0.01)$ to be connected; 2) Complete graph: every pair of nodes has an edge; 3) Star graph: $2/3$ of nodes are randomly connected with a probability of $p (= 0.01)$. And the rest of $1/3$ nodes are connected to all the nodes. Note that an example of star graph is given in Figure 2. Since, no node in complete graph provides more information than others, we expect TS and IDS perform the same. We expect IDS

outperforms TS more significantly in star graphs than it does in random graphs as there are more nodes gaining global information.

Low-rank matrix We sample each entry of $U \in \mathbb{R}^{d \times r}$ and $V \in \mathbb{R}^{d \times r}$ from $\mathcal{N}(0, \sigma_0^2)$ and let the unknown matrix $Y = UV^T + \epsilon$, with each entry of ϵ from $\mathcal{N}(0, \sigma_\epsilon^2)$. Since we do not have a closed form for the posterior distribution, posterior samples are generated from a MCMC algorithm. In the experiment, we let $m = 30$, which gives us 900 steps at maximum. For a better demonstration of the early stage, we show the cumulative regrets in the first 100 steps.

Results. Figure 1 shows the simulations results for the four models above. IDS has lower cumulative regrets over all the steps in linear model, logistic model, graph random and graph star model. TS and IDS always outperform random policy and UCB (if available), showing a benefit of complete algorithms for ASD. In the graph experiments, IDS and TS performs almost the same for developing graph because any input gains information of the full graph. IDS has slightly lower regrets than TS in random graph because of the weak structural information. It has significantly lower regrets than TS and random because of the stronger structural information. As shown in (b, d), the cumulative regrets at step 100 grows linearly with d . A higher noise also leads to a higher regret as shown in (f, h, j, l). IDS performs better in different levels of d for linear models and different levels of σ for graph random, graph star, matrix.




Figure 1: Cumulative regret curves for linear regression (a, b), logistic regression (c, d), random graph, complete graph, star graph (e-j) and low-rank matrix (k, l). Columns 1 and 3 are the cumulative regret curves for different models. Columns 2 and 4 are cumulative regret at an early step for different hyperparameters. For linear and logistic regression model, d ranges in $\{20, 50, 100\}$. For graph model, σ_ϵ^2 ranges in $\{0.1, 1.0, 10\}$. For low-rank matrix, σ_ϵ^2 ranges in $\{0.1, 0.5, 1.0\}$.

5.3 Reaction condition discovery

To further test the usefulness of IDS and TS for ASD in real-world problems, we consider a *reaction condition discovery* problem in organic chemistry.

Background and dataset. In the chemical literature, a single reaction condition that gives the highest yield for one set of reacting molecules is often demonstrated on a couple dozen analogous

molecules. However, for molecules with structures that significantly differ from this baseline, one needs to *rediscover* working reaction conditions. This provides a significant challenge as the number of plausible reaction conditions can grow rapidly, due to factors such as the many combinations of individual reagents, varying concentrations of the reagents, and temperature. Therefore, strategies that can identify as many high-yielding reaction conditions as early as possible would greatly facilitate the preparation of functional molecules. Similar reaction condition discovery problem is studied under a transfer learning framework [30, 37].

The IDS algorithms may therefore be highly useful in examining experimental data compared to UCB, TS and random policies. Two experimental datasets on reactions that form carbon-nitrogen bonds, which are of high importance for synthesizing drug molecules, were therefore selected as specific test cases. The first dataset, named Photoredox Nickel Dual-Catalysis (PNDC) [11], comes from a reaction that involves two types of catalysts (a photocatalyst and a nickel catalyst) that work together to form carbon-nitrogen bonds. It has $n = 80$ reaction conditions for discovery. The variables in the dataset are the identity of photocatalysts, relative catalyst concentration, and relative substrate concentration, totaling 80 reaction conditions. The second C-N Cross-Coupling with Isoxazoles (CNCCI) dataset [2] comes from a collection of a cross-coupling reactions in the presence of an additive. This dataset was designed to test the reaction’s tolerance to the additive (isoxazole) and includes 330 candidate reactions. Both datasets have continuous response variables (reaction yield).

Results. Table 1 demonstrated the results of linear IDS, TS, UCB for PNDC and CNCCI. We tuned hyperparameters for all three algorithms. Details are given in Appendix H. Over 17 steps, IDS showed a dominant performance in all 11 molecules tested in PNDC with an average of 1.58 more discoveries than random policy. Either IDS or TS performs the best in the 9 combinations of catalyst and base tested in CNCCI dataset. Over 66 steps, IDS discovers 14.1 more plausible reaction conditions on average than random.

Table 1: Summary of cumulative regrets at 20% of total number of steps. Each column corresponds to a target molecule. The regrets are averages of 10 independent runs, whose standard deviation are given in the brackets. Best algorithm for each column is marked in red. (-) denote that the algorithms are deterministic. The columns are the indices for target molecules, whose details are provided in Appendix H.

Molecules \ Algorithms	X2	X3	X4	X5	X6	X8
Random	8.77 (1.44)	7.74 (0.75)	2.12 (0.20)	1.70 (0.15)	4.96 (0.72)	9.90 (1.47)
IDS	6.53 (0.20)	5.70 (0.52)	1.56 (0.06)	0.81 (0.01)	3.44 (0.81)	9.35 (0.96)
TS	7.64 (1.11)	6.67 (1.09)	2.15 (0.48)	1.22 (0.38)	4.07 (1.30)	9.78 (1.69)
UCB	9.63 (-)	7.66 (-)	2.58 (-)	1.42 (-)	5.26 (-)	10.43 (-)
	X11	X12	X13	X14	X15	
Random	7.15 (1.21)	5.27 (0.92)	4.44 (0.43)	5.72 (0.35)	2.65 (0.47)	
IDS	2.40 (0.12)	3.95 (0.19)	3.22 (0.21)	4.94 (0.07)	1.09 (0.03)	
TS	2.90 (1.76)	4.42 (0.90)	4.79 (1.17)	5.39 (0.64)	1.35 (0.43)	
UCB	3.58 (-)	5.28 (-)	4.53 (-)	5.86 (-)	2.03 (-)	

(a) Results for PNDC

Catalyst and Base \ Algorithms	L1+B1	L2+B1	L3+B1	L4+B1	L1+B2
Random	25.65 (2.11)	23.08 (1.98)	19.90 (1.19)	15.96 (0.83)	29.59 (1.26)
IDS	15.55 (1.52)	7.90 (0.28)	10.03 (0.88)	10.81 (0.79)	10.76 (0.51)
TS	11.84 (1.04)	9.82 (2.23)	11.58 (4.31)	10.39 (0.81)	12.49 (1.10)
UCB	13.63 (-)	10.08 (-)	12.00 (-)	11.98 (-)	17.41 (-)
	L4+B2	L1+B3	L3+B3	L4+B3	
Random	28.78 (2.59)	19.84 (0.81)	29.03 (2.22)	27.08 (2.51)	
IDS	10.51 (1.32)	8.63 (0.97)	8.76 (0.32)	9.44 (0.52)	
TS	12.68 (2.33)	10.61 (0.99)	12.44 (7.39)	9.80 (2.39)	
UCB	15.60 (-)	12.42 (-)	11.49 (-)	12.32 (-)	

(b) Results for CNCCI

6 Discussion

In this paper, we posed and comprehensively studied the Adaptive Sampling for Discovery problem and proposed the Information-directed sampling algorithm along with its regret analysis. The results are complemented with both simulation and real-data experiments. There are a few open directions following the results in the paper. First, although our paper shows bandit-like regret bounds, ASD problem by nature is different from bandit in the way that the unlabeled set is diminishing, which may further reduce the complexity of the problem. For example, in the linear model, all the unlabeled points may fall on a lower dimensional hyperplane, which would reduce the dependence on d . Second, frequentist regret may be studied instead of adopting a Bayesian framework. More practical algorithms e.g. random forests and neural network may be considered in the future work. Sampling from posterior can be slow for certain models. A fast and generalizable algorithm using IDS for more complicated and larger scale real-data applications may be considered.

7 Acknowledgement

We acknowledge the support of National Science Foundation via grant IIS-2007055 and CHE-1551994.

References

- [1] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*, pages 39–1. JMLR Workshop and Conference Proceedings, 2012.
- [2] Derek T Ahneman, Jesús G Estrada, Shishi Lin, Spencer D Dreher, and Abigail G Doyle. Predicting reaction performance in c–n cross-coupling using machine learning. *Science*, 360(6385):186–190, 2018.
- [3] Kasra Arnavaz, Aasa Feragen, Oswin Krause, and Marco Loog. Bayesian active learning for maximal information gain on model parameters. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 10524–10531. IEEE, 2021.
- [4] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.
- [5] Jürgen Bajorath, Steven Kearnes, W Patrick Walters, Nicholas A Meanwell, Gunda I Georg, and Shaomeng Wang. Artificial intelligence in drug discovery: into the great wide open, 2020.
- [6] J Eric Bickel and James E Smith. Optimal sequential exploration: A binary learning model. *Decision Analysis*, 3(1):16–32, 2006.
- [7] David B Brown and James E Smith. Optimal sequential exploration: Bandits, clairvoyants, and wildcats. *Operations research*, 61(3):644–665, 2013.
- [8] Shayok Chakraborty, Jiayu Zhou, Vineeth Balasubramanian, Sethuraman Panchanathan, Ian Davidson, and Jieping Ye. Active matrix completion. In *2013 IEEE 13th international conference on data mining*, pages 81–90. IEEE, 2013.
- [9] Corinna Cortes, Giulia DeSalvo, Claudio Gentile, Mehryar Mohri, and Scott Yang. Online learning with sleeping experts and feedback graphs. In *International Conference on Machine Learning*, pages 1370–1378. PMLR, 2019.
- [10] Shi Dong and Benjamin Van Roy. An information-theoretic analysis for thompson sampling with many actions. *Advances in Neural Information Processing Systems*, 31, 2018.
- [11] Spencer D Dreher and Shane W Krska. Chemistry informer libraries: Conception, early experience, and role in the future of cheminformatics. *Accounts of Chemical Research*, 54(7):1586–1596, 2021.
- [12] Sarah Filippi, Olivier Cappé, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. *Advances in Neural Information Processing Systems*, 23, 2010.

- [13] Yolanda Gil, Mark Greaves, James Hendler, and Haym Hirsh. Amplify scientific discovery with artificial intelligence. *Science*, 346(6206):171–172, 2014.
- [14] Botao Hao, Tor Lattimore, and Wei Deng. Information directed sampling for sparse linear bandits. *Advances in Neural Information Processing Systems*, 34, 2021.
- [15] Zichang He, Bo Zhao, and Zheng Zhang. Active sampling for accelerated mri with low-rank tensors. *arXiv preprint arXiv:2012.12496*, 2020.
- [16] Tommi S Jaakkola and Michael I Jordan. A variational approach to bayesian logistic regression models and their extensions. In *Sixth International Workshop on Artificial Intelligence and Statistics*, pages 283–294. PMLR, 1997.
- [17] Babak Jafarizadeh and Reidar Bratvold. The two-factor price process in optimal sequential exploration. *Journal of the Operational Research Society*, 72(7):1637–1647, 2021.
- [18] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 665–674, 2013.
- [19] Xiaowei Jia, Jared Willard, Anuj Karpatne, Jordan S Read, Jacob A Zwart, Michael Steinbach, and Vipin Kumar. Physics-guided machine learning for scientific discovery: An application in simulating lake temperature profiles. *ACM/IMS Transactions on Data Science*, 2(3):1–26, 2021.
- [20] Shali Jiang, Gustavo Malkomes, Matthew Abbott, Benjamin Moseley, and Roman Garnett. Efficient nonmyopic batch active search. *Advances in Neural Information Processing Systems*, 31, 2018.
- [21] Shali Jiang, Gustavo Malkomes, Geoff Converse, Alyssa Shofner, Benjamin Moseley, and Roman Garnett. Efficient nonmyopic active search. In *International Conference on Machine Learning*, pages 1714–1723. PMLR, 2017.
- [22] Varun Kanade, H Brendan McMahan, and Brent Bryan. Sleeping experts and bandits with stochastic action availability and adversarial rewards. In *Artificial Intelligence and Statistics*, pages 272–279. PMLR, 2009.
- [23] Anuj Karpatne, Gowtham Atluri, James H Faghmous, Michael Steinbach, Arindam Banerjee, Auroop Ganguly, Shashi Shekhar, Nagiza Samatova, and Vipin Kumar. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on knowledge and data engineering*, 29(10):2318–2331, 2017.
- [24] Robert Kleinberg, Alexandru Niculescu-Mizil, and Yogeshwer Sharma. Regret bounds for sleeping experts and bandits. *Machine learning*, 80(2):245–272, 2010.
- [25] Vladimir Koltchinskii, Karim Lounici, and Alexandre B Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302–2329, 2011.
- [26] Tor Lattimore. Optimally confident ucb: Improved regret for finite-armed bandits. *arXiv preprint arXiv:1507.07880*, 2015.
- [27] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [28] Yangyi Lu, Amirhossein Meisami, and Ambuj Tewari. Low-rank generalized linear bandit problems. In *International Conference on Artificial Intelligence and Statistics*, pages 460–468. PMLR, 2021.
- [29] Simon Mak, Henry Shaowu Yushi, and Yao Xie. Information-guided sampling for low-rank matrix completion. *arXiv preprint arXiv:1706.08037*, 2017.
- [30] Gregory McInnes, Rachel Dalton, Katrin Sangkuhl, Michelle Whirl-Carrillo, Seung-been Lee, Philip S Tsao, Andrea Gaedigk, Russ B Altman, and Erica L Woodahl. Transfer learning enables prediction of cyp2d6 haplotype function. *PLoS Computational Biology*, 16(11):e1008399, 2020.

- [31] Luong Trung Nguyen, Junhan Kim, and Byonghyo Shim. Low-rank matrix completion: A contemporary survey. *IEEE Access*, 7:94215–94237, 2019.
- [32] Quan Nguyen and Roman Garnett. Nonmyopic multiclass active search for diverse discovery. *arXiv preprint arXiv:2202.03593*, 2022.
- [33] Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. *Advances in Neural Information Processing Systems*, 27, 2014.
- [34] Daniel Russo and Benjamin Van Roy. An information-theoretic analysis of thompson sampling. *The Journal of Machine Learning Research*, 17(1):2442–2471, 2016.
- [35] Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. *Operations Research*, 66(1):230–252, 2018.
- [36] Daniel J Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, Zheng Wen, et al. A tutorial on thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018.
- [37] Eunjae Shim, Joshua Kammeraad, Ziping Xu, Ambuj Tewari, Tim Cernak, and Paul Zimmerman. Predicting reaction conditions from limited data through active transfer learning. *Chem. Sci.*, pages –, 2022.
- [38] Thomas J Struble, Juan C Alvarez, Scott P Brown, Milan Chytil, Justin Cisar, Renee L DesJarlais, Ola Engkvist, Scott A Frank, Daniel R Greve, Daniel J Griffin, et al. Current and future roles of artificial intelligence in medicinal chemistry synthesis. *Journal of medicinal chemistry*, 63(16):8667–8682, 2020.
- [39] Kiran K Thekumparampil, Prateek Jain, Praneeth Netrapalli, and Sewoong Oh. Statistically and computationally efficient linear meta-representation learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- [40] Jessica Vamathevan, Dominic Clark, Paul Czodrowski, Ian Dunham, Edgardo Ferran, George Lee, Bin Li, Anant Madabhushi, Parantu Shah, Michaela Spitzer, et al. Applications of machine learning in drug discovery and development. *Nature reviews Drug discovery*, 18(6):463–477, 2019.
- [41] Yizao Wang, Jean-Yves Audibert, and Rémi Munos. Algorithms for infinitely many-armed bandits. *Advances in Neural Information Processing Systems*, 21, 2008.
- [42] Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural contextual bandits with ucb-based exploration. In *International Conference on Machine Learning*, pages 11492–11502. PMLR, 2020.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? **[Yes]** See Supplementary material.
- Did you include the license to the code and datasets? **[No]** The code and the data are proprietary.
- Did you include the license to the code and datasets? **[N/A]**

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...

- (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] The limitations are discussed in the Section 6
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] Depending on the area of application, the impact can be different. The potential impact of reaction discovery is in 6.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
- (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes] The proofs are given in Appendix
3. If you ran experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The codes and dataset are in the supplemental material.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] They are provided along with code
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Figure 1 and Table 1.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] They are given in Appendix H
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We include our code in the supplementary material.
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [Yes] The data are realised with their publications.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] The data does not have personally identifiable information or offensive content.
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Proof of Proposition 1

Our proof follows the standard proof of MAB lower bound. Let Δ be some constant. Define some $\theta_1 = (\underbrace{\Delta, \dots, \Delta}_{T \text{ times}}, \underbrace{0, \dots, 0}_{T(T-1) \text{ times}})^\top$. For any algorithm running on this ASD, let $T_1(n)$ be the total number of selections for the first T samples. If $T_1(n) \leq T/2$, the regret for θ_1 is greater than $\Delta T/2$. Let

$$k = \arg \min_{t>1} \mathbb{E}_{\theta_1} T_t(n).$$

Since $\sum_t \mathbb{E}_{\theta_1} T_t(n) = T$, we have $\mathbb{E}_{\theta_1} T_k(n) \leq T/(T-1) \leq 2$. Without a loss of generality, we let $k = 2$. Define $\tilde{\theta}_2 = (\underbrace{0, \dots, 0}_{T \text{ times}}, \underbrace{2\Delta, \dots, 2\Delta}_{T \text{ times}}, \underbrace{0, \dots, 0}_{T(T-2) \text{ times}})^\top$. Then consider a uniform prior over $\{\theta_1, \tilde{\theta}_2\}$. The Bayesian regret is at least

$$\frac{T\Delta}{4} \mathbb{P}_{\theta_1}(T_1(n) \leq T/2) + \frac{T\Delta}{4} \mathbb{P}_{\tilde{\theta}_2}(T_1(n) > T/2)$$

where \mathbb{P}_θ is the probability measure under the ASD problem with parameter θ . Using Lemma 15.1 [27], we have

$$\begin{aligned} \mathbb{P}_{\theta_1}(T_1(n) \leq T/2) + \mathbb{P}_{\tilde{\theta}_2}(T_1(n) > T/2) &\geq \exp(-D(\mathbb{P}_{\theta_1}, \mathbb{P}_{\tilde{\theta}_2}))/2 \\ &= \exp(-\mathbb{E}_{\theta_1}[T_2(n)] D(\mathcal{N}(0, 1), \mathcal{N}(2\Delta, 1)))/2 \\ &\geq \exp(-4\Delta^2)/2, \end{aligned}$$

where $D(P_1, P_2)$ is the KL-divergence of two probability measure P_1 and P_2 . Now choosing $\Delta = 1/2$, we have for any algorithm \mathcal{A} , $\mathcal{BR}(T, \mathcal{A}) \gtrsim T$.

B Proof of Lemma 1

We decompose the Bayesian regret in terms of the instant regret

$$\begin{aligned} \mathcal{BR}(T; IDS) &= \mathbb{E}[\sum_{t=1}^T f_\theta(X_t^*) - \sum_{t=1}^T f_\theta(X_t)] \\ &= \mathbb{E}[\sum_{t=1}^T \langle \pi_t, \Delta_t \rangle] \\ &\leq \mathbb{E}[\sum_{t=1}^T (\Psi_{*,\lambda} g_t^\top \pi_t)^{1/\lambda}] \\ &= \Psi_{*,\lambda}^{1/\lambda} \mathbb{E}[\sum_{t=1}^T (g_t^\top \pi_t)^{1/\lambda}] \\ &\leq \Psi_{*,\lambda}^{1/\lambda} T^{1-1/\lambda} \mathbb{E}[\sum_{t=1}^T (g_t^\top \pi_t)]^{1/\lambda} \\ &\leq \Psi_{*,\lambda}^{1/\lambda} T^{1-1/\lambda} \mathbb{E}[\sum_{t=1}^T I_t(X_{t,1}^*, \dots, X_{t,T-t+1}^*; (X_t, Y_t))]^{1/\lambda} \\ &\quad \text{(using the fact that } \{X_{t,1}^*, \dots, X_{t,T-t+1}^*\} \subset \{X_1^*, \dots, X_T^*\}) \\ &\leq \Psi_{*,\lambda}^{1/\lambda} T^{1-1/\lambda} \mathbb{E}[\sum_{t=1}^T I_t(X_1^*, \dots, X_T^*; (X_t, Y_t))]^{1/\lambda} \\ &\leq \Psi_{*,\lambda}^{1/\lambda} T^{1-1/\lambda} H(X_1^*, \dots, X_T^*)^{1/\lambda} \end{aligned}$$

C Generalized Linear Model

We start from proving for the simple linear regression i.e. $\mu(x) = x$.

Proof. Let π_t^{ts} be the Thompson sample policy at the step t , i.e. $\pi_t^{ts}(x) = \mathbb{P}_t(X_{t,1}^* = x)$. We have

$$\Psi_t = \frac{(\Delta_t^\top \pi_t)^2}{g_t^\top \pi_t} \leq \frac{(\Delta_t^\top \pi_t^{ts})^2}{g_t^\top \pi_t^{ts}}.$$

We first write the instant regret in the following form:

$$\begin{aligned} \Delta_t^\top \pi_t^{ts} &= \mathbb{E}_t[\theta^\top X_{t,1}^*] - \sum_{x \in S_t^n} \pi_t^{ts}(x) \mathbb{E}_t[\theta^\top x] \\ &\leq \sum_{x \in S_t^n} \mathbb{P}_t(X_{t,1}^* = x) (\mathbb{E}_t[\theta^\top x \mid X_{t,1}^* = x] - \mathbb{E}_t[\theta^\top x]). \end{aligned}$$

Now we deal with the information gain term. Let $Y_{t,x}$ be the observation at step t when selecting input x and $\epsilon_{t,x}$ be the noise generated. We have $Y_{t,x} = \mu(\theta^\top x) + \epsilon_{t,x}$. The mutual information can be represented by the Kullback-Leibler divergence between the joint distribution of the two variables and the product of their marginal distribution, i.e.

$$I_t(X_{t,1}^*; Y_{t,x}) = \sum_{x' \in S_t^n} D(\mathbb{P}_t(Y_{t,x} \mid X_{t,1}^* = x') \parallel \mathbb{P}_t(Y_{t,x})). \quad (3)$$

Lemma 2 (Fact 9 [34]). *For any distribution P and Q such that P is absolutely continuous with respect to Q , any random variable $X : \Omega \mapsto \mathcal{X}$ and any $g : \mathcal{X} \mapsto \mathbb{R}$ such that $\sup g - \inf g \leq 1$,*

$$\mathbb{E}_P[g(X)] - \mathbb{E}_Q[g(X)] \leq \sqrt{\frac{1}{2} D(P \parallel Q)},$$

where \mathbb{E}_P and \mathbb{E}_Q denote the expectation operators under P and Q .

By Lemma 2 with $g(x) = x$ and $X = Y_{t,x}$, the information gain can be lower bounded by:

$$g_t^\top \pi_t^{ts} \geq 2 \sum_{x, x' \in S_t^n} \mathbb{P}_t(X_{t,1}^* = x) \mathbb{P}_t(X_{t,1}^* = x') (\mathbb{E}_t[Y_{t,x} \mid X_{t,1}^* = x'] - \mathbb{E}_t[Y_{t,x}])^2$$

Let

$$M_{x,x'} = \sqrt{P(X_{t,1}^* = x)P(X_{t,1}^* = x')} (\mathbb{E}[Y_{t,x} \mid X_{t,1}^* = x] - \mathbb{E}[Y_{t,x}]).$$

Then $\Delta_t^\top \pi_t^{ts} = \text{trace}(M)$ and $v_t^\top \pi_t^{ts} = \|M\|_F^2$.

By Fact 2 of [35], we have

$$\Delta_t^\top \pi_t^{ts} / v_t^\top \pi_t^{ts} = \text{trace}(M) / (2\|M\|_F^2) \leq \text{rank}(M) \leq d/2.$$

□

C.1 Proof of Theorem 1

Proof. Using the similar strategy for linear model, we let π_t^{ts} be the Thompson sample policy at the step t .

Using the fact that μ is L_μ -Lipschitz

$$\begin{aligned} \Delta_t^\top \pi_t^{ts} &= \mathbb{E}_t[\mu(\theta^\top X_{t,1}^*)] - \sum_{x \in S_t^n} \pi_t^{ts}(x) \mathbb{E}_t[\mu(\theta^\top x)] \\ &= \sum_{x \in S_t^n} \mathbb{P}_t(X_{t,1}^* = x) (\mathbb{E}_t[\mu(\theta^\top X_{t,1}^*)] - \mathbb{E}_t[\mu(\theta^\top x)]) \\ &\leq L_\mu \sum_{x \in S_t^n} \mathbb{P}_t(X_{t,1}^* = x) (\mathbb{E}_t[(\theta^\top X_{t,1}^*)] - \mathbb{E}_t[(\theta^\top x)]) \\ &\leq L_\mu \sum_{x \in S_t^n} \mathbb{P}_t(X_{t,1}^* = x) (\mathbb{E}_t[(\theta^\top x) \mid X_{t,1}^* = x] - \mathbb{E}_t[(\theta^\top x)]). \end{aligned}$$

The major difference is on the lower bound of the information gain term. We follow a similar strategy in [10]. Let $f(x) = \mathbb{E}[\mu^{-1}(x - \tilde{\epsilon})]$, where $\tilde{\epsilon}$ follows exactly the same distribution of ϵ . Let $\tilde{Y}_{t,x} = f(Y_{t,x}) = \mathbb{E}[\mu^{-1}(Y_{t,x} - \tilde{\epsilon}) | Y_{t,x}]$. Then we have

$$\begin{aligned}\mathbb{E}[\tilde{Y}_{t,x} | \theta] &= \mathbb{E}[\mathbb{E}[\mu^{-1}(Y_{t,x} - \tilde{\epsilon}) | Y_{t,x}] | \theta] \\ &= \mathbb{E}[\mathbb{E}[\mu^{-1}(Y_{t,x} - \epsilon) | Y_{t,x}] | \theta] \\ &= \mathbb{E}[\mathbb{E}[\mu^{-1}(\theta^\top x) | Y_{t,x}] | \theta] \\ &= \theta^\top x.\end{aligned}$$

Since $\tilde{Y}_{t,x}$ is a linear regression outcome, the information gain

$$\begin{aligned}I(X_1^*; \tilde{Y}_{t,x}) &\geq 2 \sum_{x' \in S_t^n} \mathbb{P}_t(X_{t,1}^* = x') (\mathbb{E}_t[\tilde{Y}_{t,x} | X_{t,1}^* = x'] - \mathbb{E}_t[\tilde{Y}_{t,x}])^2 \\ &= 2 \sum_{x' \in S_t^n} \mathbb{P}_t(X_{t,1}^* = x') (\mathbb{E}_t[\theta^\top x | X_{t,1}^* = x'] - \mathbb{E}_t[\theta^\top x])^2.\end{aligned}$$

To proceed, we notice that $\tilde{Y}_{t,x}$ is a deterministic function of $Y_{t,x}$. Hence, $I(X_1^*; \tilde{Y}_{t,x}) \leq I(X_1^*; Y_{t,x})$. Then we have

$$g_t^\top \pi_t^{ts} \geq 2 \sum_{x, x' \in S_t^n} \mathbb{P}_t(X_{t,1}^* = x) \mathbb{P}_t(X_{t,1}^* = x') (\mathbb{E}_t[\theta^\top x | X_{t,1}^* = x'] - \mathbb{E}_t[\theta^\top x])^2$$

From here the same analysis for linear model can be applied. \square

D Low-rank Matrix

We denote a policy at time t by $\pi_t \in [0, 1]^{|\Omega_t^c|}$ each dimension corresponding to a unlabeled entry. Now we derive the instant regret and information gain. We denote an index by $a = (a_1, a_2)$, where a_1, a_2 are row and column indices.

Let μ be uniform distribution over Ω_t^c . Using Lemma 3 [34], we have

$$\begin{aligned}\langle \pi_t, I_t \rangle &\geq \frac{2}{B^2 + 1} \sum_{a \in \Omega_t^c} \mu(a) \sum_{a^* \in \Omega_t^c} \mathbb{P}_t(X_{t,1}^* = a^*) (\mathbb{E}_t[Y_{t,a} | X_{t,1}^* = a^*] - \mathbb{E}_t[Y_{t,a}])^2 \\ &= \frac{2}{B^2 + 1} \sum_{a^* \in \Omega_t^c} \mathbb{P}_t(X_{t,1}^* = a^*) \sum_{a \in \Omega_t^c} \mu(a) (\mathbb{E}_t[M_a | X_{t,1}^* = a^*] - \mathbb{E}_t[M_a])^2 \\ &= \frac{2}{B^2 + 1} \sum_{a^* \in \Omega_t^c} \mathbb{P}_t(X_{t,1}^* = a^*) \sum_{a \in \Omega_t^c} \mu(a) (\mathbb{E}_t[M_a | X_{t,1}^* = a^*] - \mathbb{E}_t[M_a])^2 \\ &= \frac{2}{B^2 + 1} \sum_{a^* \in \Omega_t^c} \mathbb{P}_t(X_{t,1}^* = a^*) \frac{1}{|\Omega_t^c|} \|M^{t,a^*}\|_F^2 \\ &\geq \frac{2}{B^2 + 1} \sum_{a^* \in \Omega_t^c} \mathbb{P}_t(X_{t,1}^* = a^*) \frac{1}{m^2} \|M^{t,a^*}\|_F^2 \\ &\geq \frac{2}{B^2 + 1} \sum_{a^* \in \Omega_t^c} \mathbb{P}_t(X_{t,1}^* = a^*) \frac{1}{Rm^2} \text{trace}^2(M^{t,a^*}).\end{aligned}$$

where $M^{t,a^*} := \mathbb{E}_t[M | X_{t,1}^* = a^*] - \mathbb{E}_t[M]$.

To upper bound the instant regret, let π_t^{ts} be the Thompson Sampling policy. We have

$$\begin{aligned}
\langle \pi_t, \Delta_t \rangle &= \sum_a \mathbb{P}_t(X_{t,1}^* = a) (\mathbb{E}_t [M_{X_1^*}] - \mathbb{E}_t [M_a]) \\
&= \sum_a \mathbb{P}_t(X_{t,1}^* = a) (\mathbb{E}_t [M_a | X_{t,1}^* = a] - \mathbb{E}_t [M_a]) \\
&\leq \sqrt{\sum_a \mathbb{P}_t(X_{t,1}^* = a) (\mathbb{E}_t [M_a | X_1^* = a] - \mathbb{E}_t [M_a])^2} \\
&\leq \sqrt{\sum_a \mathbb{P}_t(X_{t,1}^* = a) \max_{a'} (M_{a'}^{t,a})^2} \\
&\quad \text{(Using Assumption 2)} \\
&\leq \sqrt{\sum_a \mathbb{P}_t(X_{t,1}^* = a) (4 \frac{\gamma^r}{m})^2 \text{trace}^2(M^{t,a})}.
\end{aligned}$$

Now we consider a mixed policy: $\pi_t = p\mu + (1-p)\pi_t^{PS}$ for some $p \in [0, 1]$.

$$\langle \pi_t, I_t \rangle \geq p \langle \mu, I_t \rangle \geq \frac{2p}{B^2 + 1} \sum_{a^* \in \Omega_i^*} \mathbb{P}_t(X_{t,1}^* = a^*) \frac{1}{Rm^2} \text{trace}^2(M^{t,a}).$$

$$\begin{aligned}
\langle \pi_t, \Delta_t \rangle &\leq pB + (1-p) \langle \pi_t^{PS}, \Delta_t \rangle \\
&\leq pB + (1-p) \sqrt{\sum_a \mathbb{P}_t(X_{t,1}^* = a) (4 \frac{\gamma^r}{m})^2 \text{trace}^2(M^{t,a})}
\end{aligned}$$

By optimizing p , we have

$$\frac{\langle \pi_t, \Delta_t \rangle^3}{\langle \pi_t, I_t \rangle} \leq 4B(B^2 + 1)r^3\gamma^2.$$

Combined with Lemma 1, we have the following Bayesian regret bound

$$\mathcal{BR}(T, \text{IDS}) \lesssim (4B(B^2 + 1)r^3\gamma^2 H(M)T^2)^{1/3}.$$

E Graph

Proof. The \sqrt{T} term in the maximization can be achieved by replacing π_{IS} with π_{TS} and use the fact that each complete subgraph can be treated as a single node.

Let C_t be the smallest maximum independent set at the step t . To prove the $T^{2/3}$ term, we consider a mixture policy $\pi_t^{mix} = \gamma\pi_t^C + (1-\gamma)\pi_t^{TS}$ for some $\gamma > 0$, where π_t^C is the uniform distribution over C_t .

We have

$$\Delta_t(\pi_t^{mix}) \leq \gamma B + (1-\gamma)\Delta_t(\pi_t^{TS}).$$

For information gain, we have

$$g_t(\pi_t^{mix}) \geq \gamma g_t(\pi_t^C).$$

Since each node in S_t^n has an edge to at least one node in the maximum independent set (otherwise they have to be added to the set), we have

$$g_t(\pi_t^C) \geq \frac{1}{|C_t|} \sum_{x \in C_t} I_t(X_{t,1}^*; O_t(x)) \geq \frac{1}{|C_t|} \sum_{x \in S_t^n} I_t(X_{t,1}^*; Y_{t,x}).$$




Figure 2: An illustration of graphs with rich structural information. Each dashed circle represent a complete subgraph. The graph on the left hand side has $\mathcal{X}(G) = 4$, while $N = 28$. The graph on the right hand side is an illustration of the star graph, in which each node outside the dashed circle has an edge to every node inside of the circle. This graph has $\mathcal{C}(G) = 1$, while $\mathcal{X}(G) = 8$ and $N = 15$.

Using (3) and Lemma 2 again, we have

$$\sum_{x \in S_t^n} I_t(X_{t,1}^*; Y_{t,x}) \geq \sum_{x \in S_t^n} \mathbb{P}_t(X_{t,1}^* = x) (\mathbb{E}_t [Y_{t,x} | X_{t,1}^* = x'] - \mathbb{E}_t [Y_{t,x}])^2.$$

We also have

$$(\Delta_t^\top \pi_t^{TS})^2 \leq \left(\sum_{x \in S_t^n} \mathbb{P}_t(X_{t,1}^* = x) (E_t [Y_{t,x} | X_{t,1}^* = x'] - \mathbb{E}_t [Y_{t,x}]) \right)^2.$$

Therefore, we have

$$g_t(\pi_C^t) \geq \frac{1}{|C_t|} (\Delta_t(\pi_t^{TS}))^2 \geq \frac{1}{\mathcal{C}_T(G)} (\Delta_t(\pi_t^{TS}))^2.$$

Henceforth,

$$\begin{aligned} \Delta_t(\pi_t^{mix})^3 / g(\pi_t^{mix}) &\leq \frac{(\gamma B + (1 - \gamma)(\mathcal{C}_T(G)^{1/2} g_t^{1/2}) / \gamma^{1/2})^3}{g_t} \\ &\leq \frac{(\gamma B + (\mathcal{C}_T(G)^{1/2} g_t^{1/2}) / \gamma^{1/2})^3}{g_t}. \end{aligned}$$

By optimizing γ , we have

$$\Delta_t(\pi_t^{mix})^3 / g(\pi_t^{mix}) \leq BC(G).$$

□

F Generic results

The proof is analogous to the above proof for matrix and graph models. Consider a mixed policy $\tilde{\pi}_t = p\mu + (1 - p)\pi_t^{ts}$.

We have

$$g_t^\top \tilde{\pi}_t \geq pg_t^\top \mu \geq p\phi(\Delta_t^\top \pi_t^{ts})^2$$

and

$$\Delta_t^\top \tilde{\pi}_t \leq (1 - p)\Delta_t^\top \pi_t^{ts} + pB$$

Thus

$$\frac{(\Delta_t^\top \tilde{\pi}_t)^3}{g_t^\top \tilde{\pi}_t} \leq \frac{((1 - p)\Delta_t^\top \pi_t^{ts} + pB)^3}{p\phi(\Delta_t^\top \pi_t^{ts})^2} \leq \frac{(\Delta_t^\top \pi_t^{ts} + pB)^3}{p\phi(\Delta_t^\top \pi_t^{ts})^2}.$$

The proof is finished by optimizing p .

G Sparse linear model

Consider a linear regression problem

$$Y = X^\top \theta + \epsilon,$$

where $\theta \in \mathbb{R}^d$ and $\|\theta\|_0 = s$ is the sparsity and ϵ is the zero-mean noise. In general, we expect $d \gg T$, in which case, any dependence on d would lead to a linear regret.

Similarly to [14], we need to assume an exploratory unlabeled set.

Assumption 4. Let $C_{\min}(S^n) = \max_{\mu \in \mathcal{D}(S^n)} \sigma_{\min}(\mathbb{E}_{x \sim \mu}[xx^\top])$. If $C_{\min}(S_T^n) \geq 1$ almost surely for any algorithm, then we say that S^n is exploratory.

Theorem 4 (Theorem 5.3 in [14]). *The following regret holds for IDS with $\lambda = 3$, if S^n is exploratory*

$$\mathcal{BR}(T, \text{IDS}) \lesssim (s^2 T^2 \Delta)^{1/3},$$

where

$$\Delta = \min\left(\log(n), 2s \log\left(CdT^{1/2}/s\right)\right)$$

for some constant $C > 0$.

The proof in [14] on sparse linear bandit is also applicable here. The only difference is that we make a stronger assumption on the exploratory set stating that the unlabeled dataset is still exploratory after eliminating any T elements. This is to guarantee the exploratory set for any step during the decision-making process.

H Experiments

H.1 Approximate algorithm

The approximate algorithm, SampleVIDS is given in Algorithm 2.

Algorithm 2 SampleVIDS (Sample Variance-based IDS)

Input: Unlabeled dataset S_t^n , prior distribution ϕ , total number of steps T , number of posterior samples M , constant λ .

Initialize history $\mathcal{F}_0 = \{\}$.

for $t = 1$ **to** T **do**

 Sample $\theta_1, \dots, \theta_M$ from $\phi(\cdot | \mathcal{F}_{t-1})$.

 Calculate instant regret by $\Delta_t(x) = \sum_{i=1}^M \max_{x' \in S_t^n} f_{\theta_i}(x') - f_{\theta_i}(x)$ for all $x \in S_t^n$.

 Let $\Theta(x) = \{\theta_i : x \in \arg \max_{x' \in S_t^n} f_{\theta_i}(x')\}$ and $\bar{f}(x) = \sum_{i=1}^M f_{\theta_i}(x)/M$.

 Calculate variance-based information ratio for all $x \in S_t^n$ by

$$v_t(x) = \sum_{x' \in S_t^n} \frac{|\Theta(x')|}{M} \left(\frac{1}{|\Theta(x')|} \sum_{\theta \in \Theta(x')} f_{\theta}(x) - \bar{f}(x) \right)^2.$$

 Calculate the information ratio $\Psi_t(x) = \Delta_t(x)^\lambda / v_t(x)$ and label $X_t = \arg \max_x \Psi_t(x)$.

 Update history $\mathcal{F}_t = \mathcal{F}_{t-1} \cup \{(X_t, Y_t)\}$.

end for

H.2 Complete graph for simulation studies

We provide the complete simulation results in Figure 3.

H.3 Additional information for reaction condition discovery

Additional information on two datasets. The complete library of informer molecules for Photoredox Nickel Dual-Catalysis (PNDC) and the structure of reactions for C-N Cross-Coupling with Isoxazoles (CNCCI) are provided in Figure 4 and 5. For PNDC, note that we only report the results

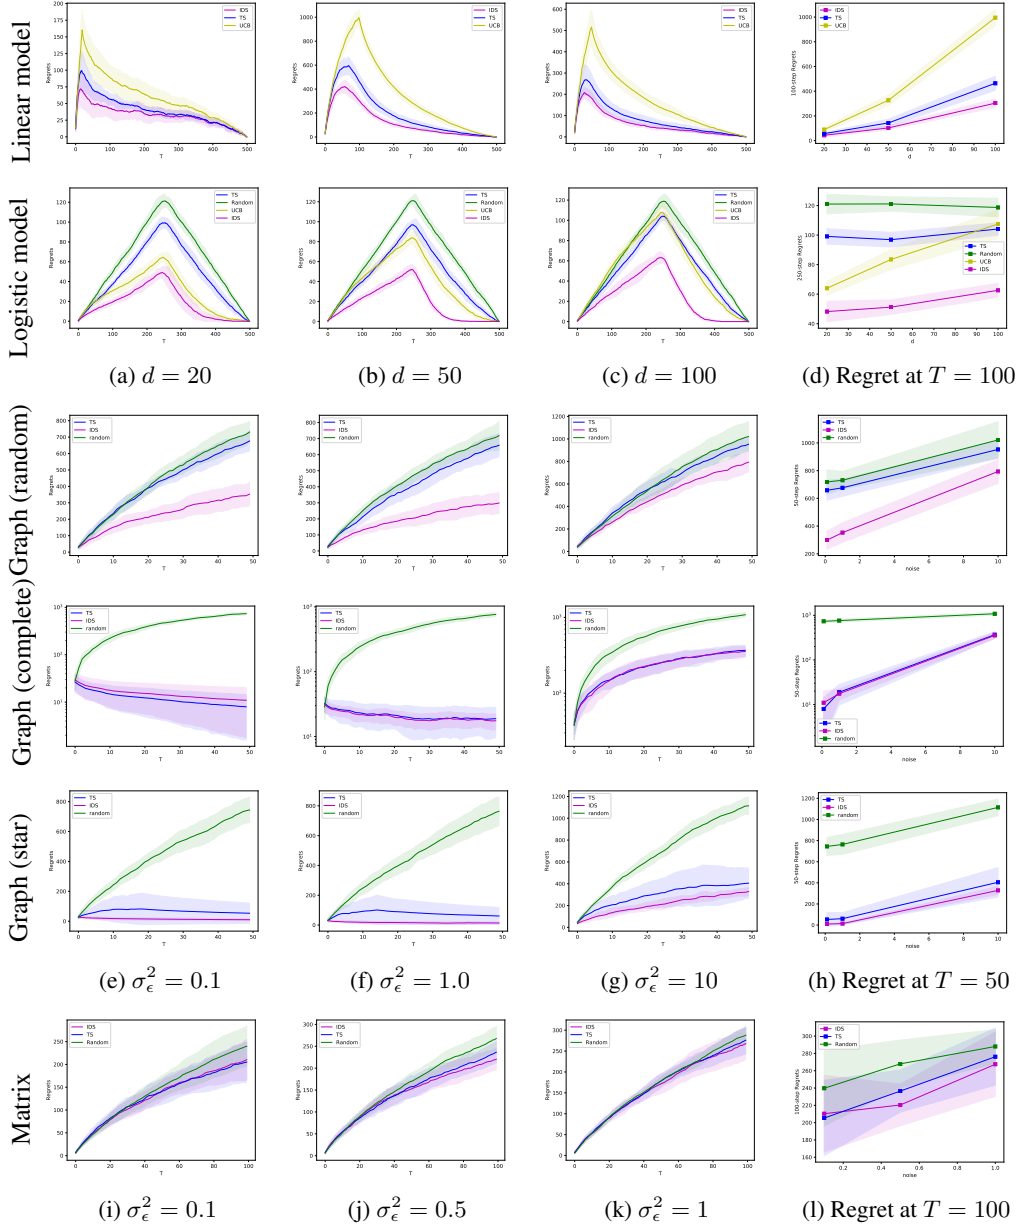


Figure 3: Cumulative regret curves for linear regression (row 1), logistic regression (row 2), random graph, star graph (row 3, 4) and low-rank matrix (row 5). The left three columns are the regret curves for different dimensions $d = 20, 50, 100$ in linear and logistic regression simulations and for different noise levels in graph and low-rank matrix simulations. The last column is the early-stage cumulative regret. The confidence ranges are given by the standard deviation of 10 independent runs.

for X2, X3, X4, X5, X6, X8, X11, X12, X13, X14, X15 in PNDC due to very low yields across all reaction conditions for the remaining molecules. Additionally, while the original dataset considers 12 photocatalysts (total 96 reaction conditions) for each molecule, we consider 10 photocatalysts (80 reaction conditions), due to the unavailability of descriptors that we intended to use. For CNCCI, only pairs of catalyst and base with datapoints of all 330 combinations of 15 aryl halides and 22 isoxazoles were considered. The features used for CNCCI is a subset of those prepared in [2], with highest feature importance values in models presented in the original work.

The two datasets are included in **PNDC.xlsx** and **CNCCI.xlsx** under data folder in the code. Each file has two sheets for yield and descriptors respectively.




Figure 4: Complete library for PNDC (Figure 2 in [11]).

Response distribution. Figure 6 provides the distributions of the response variables in two dataset.

Complete regret curves for PNDC. The complete regret curves for PNDC is given in Figure 7.

Complete regret curves for CNCCI. The complete regret curves for CNCCI is given in Figure 8.

I Comparison to ENS

In this section, we briefly compare IDS with ENS (efficient nonmyopic search).

We observed that ENS tends to over explore in the experiments on linear models. See our new Figure 9 in the appendix. We believe this is because ENS assumes that the labels of all remaining unlabeled points are conditionally independent. That is the extra gain by observing the new label y_t is uniform across all the remaining $T - t$ points. This is over estimating the gain, because in the later stage when the estimates on $\Pr(y = 1 | x, \mathcal{F}_t)$ are more accurate the extra gain from observing a single label is also much less. ENS thus weighs too much on the exploration side. We highly believe that this will lead to linear regret instead of T regret that can be achieved by IDS. This is also reflected in Figure 9, where ENS performs worse than IDS when noise level of the problem is low and we need more exploitation. In general, IDS provides a more flexible balance between exploration and exploitation.

We compare IDS with ENS on linear models with different level of noise. In general, a more noisy model requires more exploration. In Figure 9, ENS performs worse in the low-noise models while outperform IDS in higher noise settings.

I.1 Computation resources and implementation assets.

All the computation are done on MacBook Pro with 1.4 GHz Quad-Core Intel Core i5 Processor and 16GB Memory. Part of the code is from [39].




Figure 5: Complete library for CNCCI (Figure 3 in [2]).




Figure 6: The distribution of the response variables in PNDC and CNCCI dataset. x-axes are response variables (yield rate) and each panel corresponds to one target molecule.

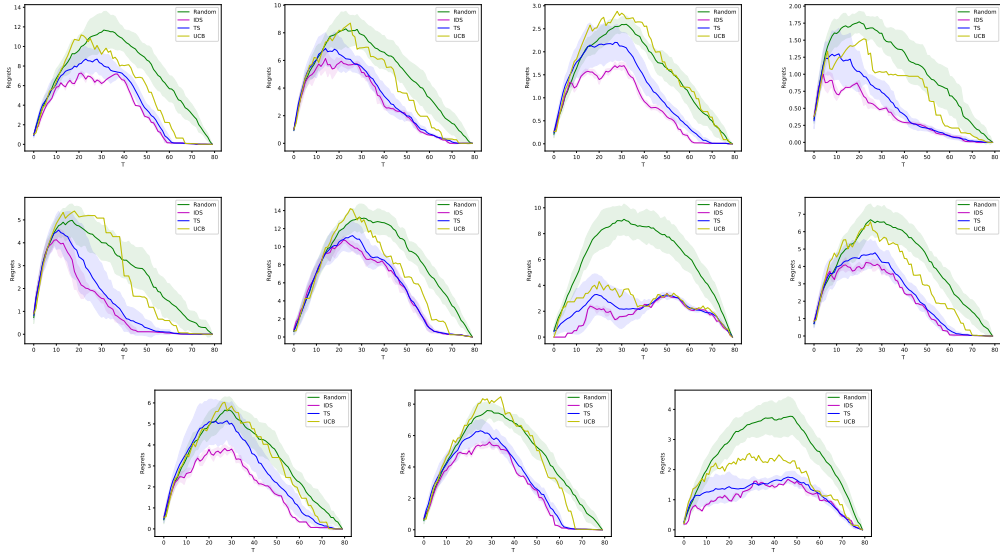


Figure 7: The whole horizon regret curves for PNDC dataset that corresponds to Table 1. The confidence interval are the standard deviation calculated from 10 independent runs.

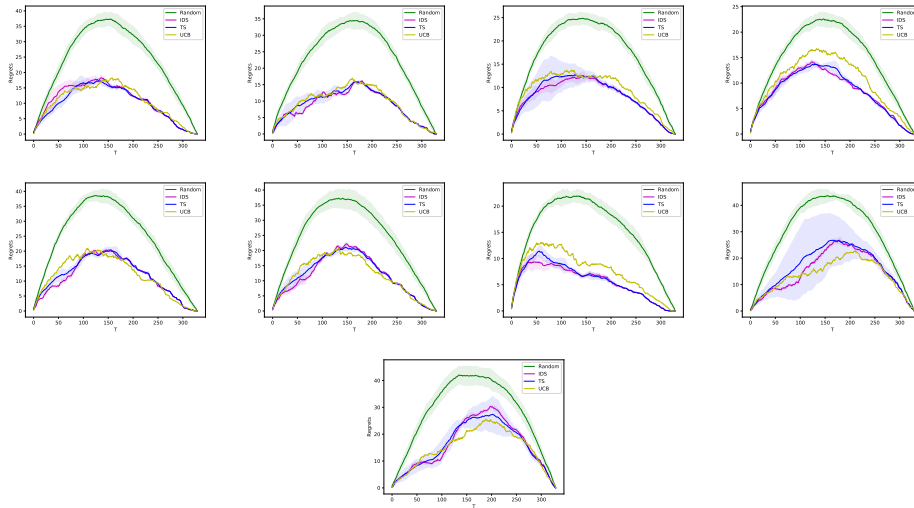


Figure 8: The whole horizon regret curves for CNCCI dataset that corresponds to Table 1. The confidence interval are the standard deviation calculated from 10 independent runs.




Figure 9: Comparing IDS with ENS on linear models with $\sigma = 0.1, 1, 5, 10$.