

# ADAPTIVE QUALITY-DIVERSITY TRADE-OFFS FOR LARGE-SCALE BATCH RECOMMENDATION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

A core research question in recommender systems is to propose batches of highly relevant and diverse items, that is, items personalized to the user’s preferences, but which also might get the user out of their comfort zone. This diversity might induce properties of serendipity and novelty which might increase user engagement or revenue. However, many real-life problems arise in that case: e.g., avoiding to recommend distinct but too similar items to reduce the churn risk, and computational cost for large item libraries, up to millions of items. First, we consider the case when the user feedback model is perfectly observed and known in advance, and introduce an efficient algorithm called B-DivRec combining determinantal point processes and a fuzzy denuding procedure to adjust the degree of item diversity. This helps enforcing a quality-diversity trade-off throughout the user history. Second, we propose an approach to adaptively tailor the quality-diversity trade-off to the user, so that diversity in recommendations can be enhanced if it leads to positive feedback, and vice-versa. Finally, we illustrate the performance and versatility of B-DivRec in the two settings on synthetic and real-life data sets on movie recommendation and drug repurposing.

## 1 INTRODUCTION

Preserving user engagement, that is, the willingness of users to query a recommender system and to interact with recommended items, is crucial and yet a difficult task. It is well-known that, beyond recommending merely the most popular items or those closest to the estimated user’s interests, introducing diversity in recommendation is key to avoid the churn risk, *i.e.*, customer attrition (Poulain & Tarissan, 2020). This topic has been widely studied under the name of “diverse/novel recommendations” or “serendipity” (Abbassi et al., 2009; Kotkov et al., 2018; Ziarani & Ravanmehr, 2021). The rationale behind it appears in several real-life contexts: for instance, diversity might increase revenue by keeping user engagement high in the movie streaming or music industry (Van den Oord et al., 2013; Anderson et al., 2020); get a teenager out of their comfort zone and make them discover new cultural goods (Ibrahim et al., 2025); evaluate the global state of a student’s knowledge on a specific subject in education (Chavarriaga et al., 2014; Yanes et al., 2020); or discover a first-in-class drug treatment in the pharmaceutical industry, where it has been shown that first-in-class molecules might generate higher revenue compared to well-known classes of molecules with a therapeutic advantage (“best-in-class”) (Schulze & Ringel, 2013; Spring et al., 2023). All in all, the goal in recommender systems is to satisfy apparently contradictory objectives: to recommend user-personalized items which introduce diversity in the user’s history of recommended items.

Moreover, when implementing a recommender system for real-life applications, one must also face the problems of developing computationally tractable pipelines on large libraries of items (Cha et al., 2018). Last but not least, the definition of diversity and quantifying the diversity in the recommended (batch of) items is tricky by itself, as illustrated by the multiple definitions used in the literature (refer to Appendix A for an overview). Diversity might be understood in terms of “inbatch” (local) or “interbatch” (global) (Bederina & Vie, 2025), respectively meaning that the batch of recommended items at a specific round should be diverse, or the batch of recommended items should be diverse with respect to the user’s prior history of selected items. This paper also tackles a novel problem of the literature, that is, adaptively tuning the level of diversity in the recommendations to the user. The tolerance to diversity might indeed vary from one user to the other, and contributes to the user engagement in the recommender system (Xu & Matsumura, 2024). No other paper to the best of our

knowledge has attempted to propose an automated procedure for tuning the level of diversity at an early stage of the recommendation pipeline.

**Notation.** We denote  $N$  the total number of items (potentially of the order of millions) in the universe  $\Omega := \{1, 2, \dots, N\}$ . Recommendation rounds happen sequentially: (1) a (possibly new) user  $\mathbf{h}^t$  queries the recommender system at each time  $t$ , (2) the recommender system returns a fixed-size batch  $\mathcal{S}^t \subset \Omega$  of  $B$  items to present to the user, (3) user  $\mathbf{h}^t$  outputs feedback values  $\mathbf{y}^t := \{y_1^t, \dots, y_B^t\}$  for each recommendation. We choose to ignore the specific identity of items and users, and instead, respectively define them with the item embedding  $\phi^i \in \mathbb{R}^d$  for item  $i$ , or by the *summary* of user history  $\mathbf{h}^t \in \mathbb{R}^m$  at each recommendation round  $t$ . Ideally,  $\mathbf{h}^t$  captures all the information about the items the user has *previously interacted with* (e.g.,  $\mathbf{h}^t$  can describe how much one of the  $m$  categories is liked). It can be hard-coded (one-hot encoding of liked item categories), or an embedding learned over previously collected items, or the result of some dimension reduction algorithm applied to the embeddings of positively interacted with items. We define  $\mathcal{H}^t$  the user history of prior recommendations to user  $\mathbf{h}^t$  up to the round  $t$  (not included). To define the similarity between items in a flexible fashion, we select a kernel function  $k : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$ .  $I_n$  is the identity matrix of size  $n$ . Finally, to make good and *personalized* recommendations—regardless of diversity—we aim to learn the feedback model  $q_\Theta : \mathbb{R}^d \times \mathbb{R}^m \rightarrow \mathbb{R}$  where  $q_\Theta(\phi, \mathbf{h})$  is the expected feedback for user  $\mathbf{h}$  on item  $\phi$  which should be maximized. In the case of a known and noiseless feedback model, if  $\phi$  is the  $k^{\text{th}}$  recommended item to user  $\mathbf{h}^t$  at time  $t$ , then  $y_k^t = q_\Theta(\phi, \mathbf{h}^t)$ .

**Contributions.** Our objective is to find a principled and scalable approach to implement the quality-diversity trade-off—in the large sense: any kernel, any feedback model, or as general as possible—for a possibly large number of items up to millions. It means that anything beyond a time complexity linear in the number of items  $N$  will be intractable. First, we introduce an efficient algorithm called B-DivRec combining determinantal point processes and a fuzzy denuding procedure to adjust the degree of item diversity. This helps enforcing a quality-diversity trade-off throughout the user history when the user feedback model is known (Section 3). Second, we propose an approach to adaptively tailor the quality-diversity trade-off to the user, so that diversity in recommendations can be enhanced if it leads to positive feedback, and vice-versa (Section 4). Finally, we illustrate the performance and versatility of B-DivRec in the two settings on synthetic and real-life data sets on movie recommendation and drug repurposing (Section 5). We conclude the paper by discussing the potential and the limitations of B-DivRec in Section 6.

**Metrics.** To quantify the quality and the diversity of recommendations, we define the following pointwise metrics for any round  $t > 0$ . We are aware that an abundant literature on both quality and diversity in recommender systems exist, see Appendix A where they are discussed. We decided to separately assess quality and diversity. Quality (also called relevance in this paper and denoted  $\text{rel}$ ) is simply defined by the expected click-through rate. We also consider the precision metric (denoted  $\text{prec}$ ), that is, the ratio of positively rated items (i.e., such that the feedback value is higher than a given threshold  $\tau > 0$ ) over the total number of recommended items. Moreover, we define two types of diversity: the *inbatch/local diversity* (denoted  $\text{div}^L$ ), focusing on the diversity inside a batch of  $B$  items; and the *interbatch/global diversity* (denoted  $\text{div}^G$ ), which looks at the diversity of the previously and currently recommended items, meaning that it also takes into account the user history. An intuitive idea of how diverse a set  $\mathcal{S}$  of items is can be obtained by computing the volume of the parallelotope induced by the rows of the kernel matrix  $K_{\mathcal{S}, \mathcal{S}} := (k(\phi^i, \phi^j))_{i, j \in \mathcal{S}}$  built from the item embeddings in  $\mathcal{S}$ . The volume of a set  $\mathcal{S} := \{i_1, \dots, i_B\}$  of items is  $\text{vol}(\mathcal{S}) := (\det K_{\mathcal{S}, \mathcal{S}})^{1/2}$ .<sup>1</sup>

$$\forall t > 0, \text{rel}(\mathcal{S}^t) := \frac{1}{B} \sum_{i \in \mathcal{S}^t} q_\Theta(\phi^i, \mathbf{h}^t), \quad \text{div}^L(\mathcal{S}^t) := \text{vol}(\mathcal{S}^t), \quad \text{div}^G(\mathcal{S}^t) := \text{vol}(\mathcal{S}^t \cup \mathcal{H}^t). \quad (1)$$

To assess the performance accrued by a strategy across  $T$  recommendation rounds, we consider the average of  $\text{rel}$ ,  $\text{prec}$ ,  $\text{div}^L$  and  $\text{div}^G$  across rounds. Moreover, we also consider a summary metric which mixes quality and diversity at the end of  $T$  recommendation rounds. Given a threshold  $\tau$ , the effective diversity  $\text{div}^+$  contributed by positively-labeled items for a single user  $\mathbf{h}$  after  $T$  rounds is

$$\text{div}^+(\{\mathcal{S}^1, \dots, \mathcal{S}^T\}, \mathbf{h}) := \text{vol}\left(\left\{i \mid i \in \bigcup_{s < t} \mathcal{S}^s, q_\Theta(\phi^i, \mathbf{h}) \geq \tau\right\}\right). \quad (2)$$

<sup>1</sup>By definition of a kernel function, this definition of volume is well-defined for any  $\mathcal{S} \subseteq \Omega$ .

The final summary metric is the average  $\text{div}^+$  across encountered users.

## 2 RELATED WORKS

As mentioned in the introduction, the topics addressed in this paper are related to a large section of the literature on recommendation (serendipity, avoidance of the churn risk, tractable recommender systems). Due to the space limit, we refer the reader to Appendix A for a more comprehensive review of the literature. We focus here on fundamental notions in DPPs and baselines in Section 5.

**Determinantal Point Processes.** A point process is a distribution over finite subsets of a (finite) set  $\Omega$ . A determinantal point process (DPP) is a process where the probability of sampling a subset  $A$  is correlated to the determinant of the kernel function  $k$  applied to this subset, that is,  $\mathbb{P}(S = A) \propto \det k_{A,A}$  (Macchi, 1975). Here, we consider the definition of DPPs with  $L$ -ensembles, as described in Borodin & Rains (2005), and  $k_{\Omega,\Omega}$  is a positive semi-definite matrix. Intuitively, the higher the volume formed by the item embeddings in this subset, the higher the probability. Sampling algorithms for DPPs are naively in  $\mathcal{O}(N^3)$  (Hough et al., 2006, Theorem 7 and Algorithm 18), but their complexity may be reduced down to  $\mathcal{O}(\alpha N \cdot \text{poly}(B))$  when sampling a subset of fixed size  $B$  with an  $\alpha$ -DPP (Calandriello et al., 2020) where  $\alpha \leq 1$ . However, one might want to find the subset with highest probability—that is, the most diverse—instead of sampling according to the DPP. The associated maximization problem is NP-hard (Ko et al., 1995; Grigorescu et al., 2022), but there are greedy approximations in  $\mathcal{O}(B^2 N)$  (Gillenwater et al., 2012; Chen et al., 2018) for maximizing over subsets of size  $B$ . Finally, conditioning over another subset  $\mathcal{H}$  of items—that is, sampling a subset which is diverse compared to previously selected set of points—can be described with the following distribution (Borodin & Rains, 2005)  $\mathbb{P}(S = A \cup \mathcal{H} \mid S \supseteq \mathcal{H}) \propto \det(k_{A,A} - k_{A,\mathcal{H}} k_{\mathcal{H},\mathcal{H}}^{-1} k_{\mathcal{H},A}^\top)$ . This is a simple approach to integrating the user history to the recommendation. However, conditioning has a dependency in  $\Omega(|\mathcal{H}|^3)$  where  $\mathcal{H}$  is the user history, and inversion of the history-related kernel matrix might be expensive.<sup>2</sup>

To tackle the issue of recommending items with high quality/relevance and high “inbatch”/local diversity, Kulesza & Taskar (2010) introduced the quality-diversity (QD) decomposition of a DPP. Given  $N$  item embeddings  $\{\phi^i\}_{i \in \Omega}$  such that  $\|\phi^i\|_2 = 1$  for any  $i \in \Omega$ , and positive quality scores  $\{q_\Theta(\phi^i, \mathbf{h})\}_{i \in \Omega}$  for each item and a given user  $\mathbf{h}$ , the QD decomposition is the distribution  $\mathbb{P}(S = A) \propto \det(Q_A \Phi_A \Phi_A^\top Q_A)$ , where  $Q_A$  is a diagonal matrix of size  $B \times B$  which contains the quality scores for each item  $i_1, \dots, i_B$  in  $A$ , and  $\Phi_A := [\phi^{i_1}, \dots, \phi^{i_B}]^\top$  is the row-concatenation of all item embeddings in  $A$ . Many papers relied on this approach (Gong et al., 2014; Wilhelm et al., 2018; Zhan et al., 2021; Svensson et al., 2025; Xuan et al., 2025). However, this decomposition might be a bit restrictive, as the control of the quality-diversity trade-off is limited, and only linear or RBF (Affandi et al., 2014; Wilhelm et al., 2018) kernels are considered. Affandi et al. (2012) is closer to our objective of increasing the “interbatch” or global diversity across consecutive recommendation rounds, and introduces Markov DPPs applied to the daily recommendation of news headlines. The main idea is to condition the subset of items sampled at round  $t + 1$  on the subset sampled at round  $t$ . Their construction guarantees that the DPP marginals are maintained. Finally, the authors also study in their experiments the unknown feedback model setting, and tackle this problem by updating empirical quality scores for each item over time in the QD decomposition. However, their procedure is quite costly and not tractable in large data sets (see our experiments in Section 5).

**Non-DPP recommender systems for diversity.** Besides DPPs, there is a plethora of other approaches, including Maximal Marginal Relevance (MMR) (Carbonell & Goldstein, 1998). The problem originally tackled by MMR is slightly different from ours: given a library of items  $\Omega$ , an item-item similarity function  $\text{sim}$ , a relevance function  $\text{rel}$  between items, and a quality-diversity trade-off parameter  $\lambda \in [0, 1]$ , the goal is to return an item  $\phi^t$  at round  $t$  which is both relevant to the query item  $\mathbf{h}^{\text{query}}$  and diverse with respect to previously selected items  $i_1, i_2, \dots, i_{t-1}$ . The recommendation procedure is defined recursively. At time  $t$ , MMR selects item  $i_t := \arg \max_{i \in \Omega} (\lambda \text{rel}(\phi^i, \mathbf{h}^{\text{query}}) - (1 - \lambda) \max_{j \in \{i_1, \dots, i_{t-1}\}} \text{sim}(\phi^i, \phi^j))$  for each

<sup>2</sup>We abuse notation here, as it is clear that we refer to the complexity notation  $\Omega$  and not the universe of items  $\Omega$ .

$t \in \{1, \dots, B\}$ , where the maximum value over  $\emptyset$  is set to 0 by convention. However, it is easy to rewrite the expression in MMR to fit the setting where a user context and history is provided as query. Instead of a relevance function, we consider the feedback model and the kernel function as a similarity metric. Using notation from above, with  $b \in \{1, \dots, B\}$ , at time  $t$ , MMR selects item  $i^{t+b} := \arg \max_{i \in \Omega} (\lambda q_{\Theta}(\phi^i, \mathbf{h}^t) - (1 - \lambda) \max_{j \in \mathcal{H}^t \cup \{i^t, \dots, i^{t+b-1}\}} k(\phi^i, \phi^j))$ . Yet, it would be interesting to strengthen even further the constraint on diversity, and to leverage the literature on DPPs to sample a whole diverse subset instead of an iterative, greedy procedure.

### 3 A GENERIC APPROACH FOR THE QUALITY-DIVERSITY TRADE-OFF IN DPPS

In this section, we define our trade-off on relevance and diversity, instead of a multi-objective task (Abbassi et al., 2009) or a multi-step process for filtering relevance or diversity (Ibrahim et al., 2025; Yuan et al., 2016). We recall that  $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$  is a similarity kernel on items. To use a quality-diversity decomposition, we need:

**Assumption 3.1.** Positive-definite kernel. *The kernel  $k$  is a positive-definite kernel, that is, there exists  $d' \in \mathbb{N}^*$  and a feature map function  $\nu : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$  such that, for any  $\mathbf{x}, \mathbf{y} \in \Omega$ ,  $k(\mathbf{x}, \mathbf{y}) = \nu(\mathbf{x})^\top \nu(\mathbf{y})$ .  $\nu$  is extended to subsets of items:  $\nu(\mathcal{S}) := (\nu(\phi^i))_{i \in \mathcal{S}} \in \mathbb{R}^{B \times d'}$  for any  $\mathcal{S} \subseteq \Omega$ .*

**Assumption 3.2.** Positive unbounded feedback. *The feedback model  $q_{\Theta}$  has values in  $\mathbb{R}_+^*$ .*

**Assumption 3.3.** Unit embeddings.  *$\|\phi\|_2 = \|\mathbf{h}\|_2 = 1$  for all user embeddings  $\mathbf{h} \in \mathbb{R}^m$  and item embeddings  $\phi \in \Omega$ .*

Note that those assumptions are only moderately restrictive, as adequate kernel functions (linear, RBF, Matérn (Williams & Rasmussen, 2006; Duvenaud, 2014)) and classification or regression models with positive values (Lee & Seung, 1999; Wood, 2017) abound in the literature, and renormalizing embeddings is a common procedure. Then, we assume that the feedback model is known and perfect, and that the observed user feedback is noiseless.

**Assumption 3.4.** Noiseless observed feedback. *At a given recommendation round  $t > 0$  to user  $\mathbf{h}^t$ , if  $\phi$  is the  $k^{\text{th}}$  recommended item, then the observed feedback value  $y_k^t$  from user  $\mathbf{h}^t$  for the  $k^{\text{th}}$  item satisfies  $y_k^t = q_{\Theta}(\phi, \mathbf{h}^t)$ .*

**The Quality-Diversity Trade-off (QDT) family.** Under these assumptions, we extend the quality-diversity decomposition (Kulesza & Taskar, 2010) in determinantal point processes to any valid kernel, and we allow to interpolate from only quality-focused recommendation<sup>3</sup> to diversity-focused recommendation<sup>4</sup>. The likelihood matrix for this family of DPPs, called QDT, for a given user  $\mathbf{h}$ , their history  $\mathcal{H}$ , and item subset  $\mathcal{S}$ , is:

$$\forall \mathcal{S} \subseteq \Omega, \forall \mathbf{h} \in \mathbb{R}^m, L_f^\lambda(\mathcal{S}; \mathbf{h}) := (Q_{\mathbf{h}, \mathcal{S}})^{2\lambda} f(k, \mathcal{S}, \mathcal{H})^{2(1-\lambda)} (Q_{\mathbf{h}, \mathcal{S}})^{2\lambda}, \quad (3)$$

where  $\lambda \in [0, 1]$  controls the trade-off between quality ( $\lambda = 1$ ) and diversity ( $\lambda = 0$ ). The presence of  $\lambda$  allows us to implement continuously and explicitly the trade-off between relevance and diversity.  $Q_{\mathbf{h}, \mathcal{S}}$  is a diagonal matrix containing the expected rewards  $\{q_{\Theta}(\phi^i, \mathbf{h})\}$  for  $i \in \mathcal{S}$ .  $f$  is a function with values in the set of positive-definite matrices in  $\mathbb{R}^{B \times B}$ , and depends on kernels computed on  $\mathcal{S}, \Omega, \mathcal{H}$ .  $f$  should incorporate all information about the desired diversity in recommended items. Note that we completely disentangle quality and diversity in this definition.

Equation 3 offers a flexible definition of the quality-diversity trade-off, which recovers various known DPPs described in Section 2 by defining  $f$  and  $\lambda$ . For instance, traditional quality-diversity decomposition (Kulesza & Taskar, 2010) can be obtained by setting  $f^{\text{QDDcomp}}(k, \mathcal{S}, \mathcal{H}) := k_{\mathcal{S}, \mathcal{S}}$  with  $\lambda = 0.5$ , where  $k$  is the linear kernel  $k(\{\psi\}, \{\phi\}) = \psi^\top \phi$  or, equivalently,  $\nu : x \mapsto x$ . A conditional DPP (Borodin & Rains, 2005) can be described using  $\lambda = 0.5$  and  $f^{\text{CondDPP}}(k, \mathcal{S}, \mathcal{H}) := k_{\mathcal{S}, \mathcal{S}} - k_{\mathcal{S}, \mathcal{H}}(k_{\mathcal{H}, \mathcal{H}})^{-1}(k_{\mathcal{H}, \mathcal{S}})^\top$ . Finally, an  $\varepsilon$ -greedy approach can be obtained by setting  $f^{\varepsilon\text{-greedy}}(k, \mathcal{S}, \mathcal{H}) := I_B$  with  $\lambda = 0.5 \varepsilon\%$  of the rounds (greedy phase), and setting  $f^{\varepsilon\text{-greedy}}(k, \mathcal{S}, \mathcal{H}) := k_{\mathcal{S}, \mathcal{S}}$  with  $\lambda = 0$  the other  $(1 - \varepsilon)\%$  of rounds (exploratory phase). We now consider the log-determinant of  $L_f^\lambda(\mathcal{S}; \mathbf{h})$  as the score  $\text{scr}$  of any  $B$ -sized set  $\mathcal{S}$

$$\forall \mathcal{S} \subseteq \Omega, \forall \mathbf{h} \in \mathbb{R}^m, \text{scr}_{f, q_{\Theta}}^\lambda(\mathcal{S}; \mathbf{h}) := 4(1 - \lambda) \log \text{vol}(f(k, \mathcal{S}, \mathcal{H})) + 4\lambda \sum_{\phi \in \mathcal{S}} \log q_{\Theta}(\phi, \mathbf{h}), \quad (4)$$

<sup>3</sup>That is, we do not take into account the diversity described by the kernel.

<sup>4</sup>i.e., we ignore quality scores and only aim at being diverse

**Algorithm 1** Recommendation round at  $t > 0$  with B-DivRec and an adaptive Q-D trade-off.

---

```

1: Input  $\mathbf{h}^t$ : user context,  $\mathcal{H}^t$ : user history,  $q_\Theta$ : feedback model,  $\lambda^t$ : initial relevance weight
2: Parameters  $k$ : kernel function,  $f = f^{\text{B-DivRec}}$ ,  $\mathcal{A}$ : online learner
3: Output  $\mathcal{S}^t$ :  $B$  item recommendations,  $\lambda^{t+1}$ : updated relevance weight
4:  $Q_{\mathbf{h}^t, \Omega} \leftarrow \text{diag}(\{q_\Theta(\phi, \mathbf{h}^t)\}_{\phi \in \Omega})$  and initialize  $\mathcal{A}$  with  $\lambda^t$ 
5:  $L_{f^{\text{B-DivRec}}}^{\lambda^t}(\Omega; \mathbf{h}^t) \leftarrow (Q_{\mathbf{h}^t, \Omega})^{2\lambda^t} f(k, \Omega, \mathcal{H}^t)^{2(1-\lambda^t)} (Q_{\mathbf{h}^t, \Omega})^{2\lambda^t}$  # Likelihood matrix
6: Use SAMPLING Sample  $\mathcal{S}^t \subseteq \Omega$  proportionally to  $\det L_f^{\lambda^t}(\mathcal{S}^t; \mathbf{h}^t)$ 
7: or MAXIMIZATION Solve  $\mathcal{S}^t \in \arg \max_{\mathcal{S} \subseteq \Omega} \det L_f^{\lambda^t}(\mathcal{S}; \mathbf{h}^t)$ 
8:  $\mathbf{y}^t \leftarrow (y_1^t, y_2^t, \dots, y_B^t)$ , update  $\mathcal{A}$  with  $(1 - 2\lambda^t) \nabla_{\lambda} \text{scr}_{f, \mathbf{y}^t}^{\lambda^t}(\mathcal{S}^t; \mathbf{h}^t)$  and select  $\lambda^{t+1}$  from  $\mathcal{A}$ 
9: return  $\mathcal{S}^t, \lambda^{t+1}$ 

```

---

with the convention  $\text{scr}_{f, q_\Theta}^\lambda(\emptyset) := 0$ , provided that the volume is well-defined—that is,  $f(k, \Omega, \emptyset)$  is positive-definite for any universe of items  $\Omega$ . From the expression in Equation 4, the larger the score, the better the subset  $\mathcal{S}$  for the quality-diversity trade-off. We could stop at this point, and simply consider a (possibly conditional) DPP to obtain diversified recommendations (Affandi et al., 2012). Still, as discussed in the introduction, this approach might be expensive for users with long histories. We suggest another, more tractable, approach named B-DivRec below, which leverages the existence of kernel-associated feature maps (Assumption 3.1).

**The B-DivRec DPP.** B-DivRec belongs to the QDT family, and is defined with  $\alpha \in [0, 2]$

$$f^{\text{B-DivRec}}(k, \mathcal{S}, \mathcal{H}; \alpha) := \left( \nu(\mathcal{S}) - \nu(g(\mathcal{S}; \mathcal{H}, \alpha)) \right) \left( \nu(\mathcal{S}) - \nu(g(\mathcal{S}; \mathcal{H}, \alpha)) \right)^\top \text{ where} \quad (5)$$

$$\forall i \leq B, g(\mathcal{S}; \mathcal{H}, \alpha)_{i \cdot} := \begin{cases} \phi^{\ell_i} & \text{if } \max_{j \in \mathcal{H}} \cos(\nu(\phi^{\ell_i}), \nu(\phi^j)) \geq 1 - \alpha \\ \mathbf{0} & \text{otherwise} \end{cases} \in \mathbb{R}^{B \times d} \text{ if } \mathcal{S} = \{\ell_i\}_i,$$

and  $\cos(x, y)$  is the cosine similarity between vectors  $x$  and  $y$ .  $\alpha$  controls in a more subtle way than the kernel  $k$  the degree of diversity expected compared to the user history, by filtering out items too similar to the history. We discuss the choice of hyperparameter  $\alpha$  in Section 5. See Algorithm 1 for a pseudo-code of B-DivRec with a fixed value of  $\lambda \in [0, 1]$ . (Gartrell et al., 2017; Dupuy & Bach, 2018) also propose a family of low-rank factorizations of  $L$ -matrices. However, Gartrell et al. (2017) does not endorse a quality-diversity trade-off, whereas Dupuy & Bach (2018) requires two supplementary low-rank approximations. Conversely, B-DivRec straightforwardly uses the Nyström approximation to obtain computations in  $\mathbb{R}^{d'}$  where  $d' \ll d$ .

**Efficiency of B-DivRec.** Naively, the computation of the  $L$ -matrix in B-DivRec is in  $\mathcal{O}(N^3)$ . However, we leverage several approximations and methods to achieve a time complexity linear in  $N$  that we describe more thoroughly in Appendix B. In practice, we learn function  $\nu$  with a Nyström approximation (Nyström, 1930; Yang et al., 2012; Liu et al., 2021) of low rank  $d'$  to avoid computations on potentially large matrices of size  $N \times N$ . This approximation can be applied to any kernel function  $k$ . However, one should be careful to select  $d'$  such that  $N \gg d' \geq B + \max_{t \leq T} |\mathcal{H}^t|$  to ensure that the interbatch/global volume of any sampled batch of  $B$  items can be larger than 0 (see metrics in Section 1). The Nyström approximation of rank  $d'$  on a random selection of representative points, run once, has a time complexity of  $\mathcal{O}(N(d')^2 + (d')^3)$  (Williams & Seeger, 2000). The computation of  $g(\mathcal{S}; \mathcal{H}, \alpha)$  for any user  $\mathbf{h}$  of history  $\mathcal{H}$  and subset  $\mathcal{S}$  using a  $k$ -d tree leads to an average time complexity in  $\Omega(B \log |\mathcal{H}|)$ , where  $|\mathcal{H}| \ll N$  as a general rule, and  $\Omega(B|\mathcal{H}|)$  in the worst case of unbalanced trees (Arya & Mount, 1993). Retrieving the closest neighbor can also be computed on large sets of (feature maps of) items by considering an approximate nearest neighbor algorithm, e.g., FAISS (Douze et al., 2024) or LSH (Dasgupta et al., 2011). The computation of the matrix power (in  $\lambda$ ) can also be in linear time in  $N$ . Using the  $\alpha$ -DPP sampling procedure (Calandriello et al., 2020) for the SAMPLING strategy and the greedy algorithm in (Chen et al., 2018) for the MAXIMIZATION approach, with time complexities linear in  $N$  as described in Section 1, we confirm that B-DivRec remains tractable even when faced with millions of items, as demonstrated in the experiments. Moreover, we delve into supplementary details of implementation in Appendix B.

## 4 ADAPTIVE QUALITY-DIVERSITY (Q-D) TRADE-OFF

A frequent problem when dealing with the quality-diversity trade-off is to handle the fact that some users might be more receptive to diverse recommendations than others. We wish for an automated procedure for tuning the trade-off—in practice, changing the value of hyperparameter  $\lambda$ —which takes into account the user’s prior reactions to diversified items. We can frame this problem as an on-line learning game between a player—that is, the recommender system—and Nature—which is the interaction with *one* user  $(\mathbf{h}, \mathcal{H}^t)$  querying the recommender system at time  $t$  (Auer et al., 2002; De Rooij et al., 2014). At round  $t \leq T$  of the game, the recommender system chooses a value of  $\lambda^t \in [0, 1]$ . Then, the recommender system makes a recommendation  $\mathcal{S}^t$  to user  $\mathbf{h}$  and receives a loss (or gain (De Rooij et al., 2014)) value related to the quality-diversity trade-off achieved, depending on the vector of feedback  $\mathbf{y}^t$  from Nature. The recommender system should use this loss to update and use  $\lambda^{t+1} \in [0, 1]$  at the next round. The game ends after  $T$  interactions with the same user.

The quality-diversity trade-off gain is given by the function  $\lambda, \mathcal{S}, \mathbf{y} \mapsto \text{scr}_{f^{\lambda}_{\text{B-DivRec}}, \mathbf{y}}^{\lambda}(\mathcal{S}; \mathbf{h})$  for all  $\lambda \in [0, 1]$ ,  $\mathcal{S} \subseteq \Omega$ ,  $|\mathcal{S}| = B$ , and  $\mathbf{y} \in (\mathbb{R}_+^*)^B$ .<sup>5</sup> The goal of the game, by selecting the  $\lambda_t$ ’s, is to maximize the cumulated quality-diversity trade-off, and alternatively, to minimize the cumulative regret compared to a deterministic oracle which knows in advance the pairs batches-feedbacks  $(\mathcal{S}^t, \mathbf{y}^t)_{t \leq T}$ :

$$\mathcal{R}^{\text{adapt}}(T; \mathbf{h}) := \max_{\lambda \in [0, 1]} \sum_{t \leq T} \text{scr}_{f^{\lambda}_{\text{B-DivRec}}, \mathbf{y}^t}^{\lambda}(\mathcal{S}^t; \mathbf{h}) - \text{scr}_{f^{\lambda^t}_{\text{B-DivRec}}, \mathbf{y}^t}^{\lambda^t}(\mathcal{S}^t; \mathbf{h}). \quad (6)$$

Many online learners have been introduced to solve this type of problem, *e.g.*, EXP3 (Auer et al., 2002) and AdaHedge (De Rooij et al., 2014). We select AdaHedge, as it is an online learner which is horizon  $T$ -agnostic, and does not require to know the scale of the gain function in advance. We further go into details as regards the implementation of this procedure in Appendix C. Algorithm 1 (purple lines) shows how we modify the initial recommendation algorithm with B-DivRec to adaptively choose  $\lambda$ .

## 5 EXPERIMENTAL STUDY

We consider several DPPs from the QDT family: QD decomposition, conditional DPP and the  $\varepsilon$ -greedy approach described in Section 3, along with baselines from the literature: Markov DPP (Affandi et al., 2012), MMR (Carbonell & Goldstein, 1998). To try to get as close as possible to a realistic online setting, we consider the following situation. At recommendation time  $t > 0$ , for a new user  $\mathbf{h}^t$  with ground-truth history  $\mathcal{H}^t$  from the initial data set, a recommender system will output a batch of recommendations for each round  $t, t+1, \dots, t+|\mathcal{H}^t|$  with respective user context-history pairs  $(\mathbf{h}^t, \emptyset), (\mathbf{h}^t, \{i_1\}), (\mathbf{h}^t, \{i_1, i_2\}), \dots, (\mathbf{h}^t, \mathcal{H}^t)$ , where  $\mathcal{H}^t := \{i_1, i_2, \dots, i_M\}$  if  $|\mathcal{H}^t| = M$  in the initial data set. It means that at each round, we incrementally increase the user history with true previously recommended items, and track the relevance and diversity of novel recommendations (see all metrics in Section 1). As the SAMPLING strategy is not deterministic, we iterate this process over 10 random seeds, and average all the metrics across those 10 iterations for each user. For pointwise metrics such as in Equation 1, we also average these metrics along the trajectory of recommendations of length  $|\mathcal{H}^t|$ . Finally, since we run this setting on several users, we aggregate these metrics across users by considering the average and the standard deviation. These are the final values shown in the tables in this section. Top ones are in bold type, second best one are underlined. For the sake of clarity, we only show the most relevant metrics and baselines in experiments. The full numerical results for each experiment can be found in Appendix E. Runtime is the time in seconds needed to output a single batch of recommendations. All values are rounded to the closest 2<sup>nd</sup> decimal place.

For all experiments, we consider a linear kernel function  $k : \phi, \psi \mapsto \phi^\top \psi$  to compare fairly with baselines for the quality-diversity trade-off. Similarly, we use as default value  $\lambda = 0.5$ : equal weight for the quality and diversity tasks;  $\varepsilon = 0.1$ : frequency of the greedy phase, to get a strong

<sup>5</sup>Analogously to a feedback model,  $\mathbf{y}$  is such that  $\mathbf{y}(\phi, \mathbf{h})$  is the (observed) feedback from the user  $\mathbf{h}$  if  $\phi$  has been recommended, and otherwise is equal to 1 (meaning that it is ignored for a non-visited  $\phi$ ).

baseline for diversity; and  $\alpha = 0$ : diversity hyperparameter for B-DivRec in Equation 5; unless otherwise specified. We consider the following data sets in our experiments: (1) Synthetic data sets of different sizes: their names are prefixed with SYNTHETIC, followed by the number of items  $N$ . Item and user embeddings are sampled at random from Gaussian distributions and the feedback model is  $q_\theta : \phi, \mathbf{h} \mapsto ((\phi)^\top \mathbf{h} + 1)/2 \in [0, 1]$ . Supplementary information can be found in Appendix D.1; (2) MovieLens data set (Harper & Konstan, 2015) for movie recommendation. Universal Sentence Encoder (Cer et al., 2018) was applied to each movie metadata, and a SVD model (Vie et al., 2015) has been trained on the existing ratings to yield a feedback model with values in  $1, 2, \dots, 5$ . Finally, (3) PREDICT data set (Réda, 2023) for drug repurposing: the item (drug) and user (disease) embeddings are provided, and the feedback model has been fit with a Heterogeneous Attention Network (HAN) (Wang et al., 2019) on known Phase 2 and 3 clinical trials, yielding rewards in  $[0, 1]$ . Further details about the real-life data sets are in Appendix D.2.




Figure 1: Sensitivity analysis for  $\lambda$  on SYNTHETIC750 (6 users,  $B = 3$ ,  $\tau = 0.5$ ), with the MAXIMIZATION strategy.

Table 1: Sensitivity analysis for  $\alpha$  on SYNTHETIC750 (6 users,  $B = 3$ ,  $\tau = 0.5$ ), with the MAXIMIZATION strategy.

$\alpha$	REL $\uparrow$	PREC $\uparrow$	DIV <sup>G</sup> $\uparrow$
0	0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	0.95 $\pm$ 0.0
0.1	0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	0.95 $\pm$ 0.0
0.5	0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	0.95 $\pm$ 0.0
0.8	0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	0.96 $\pm$ 0.0
0.90	0.6 $\pm$ 0.0	1.0 $\pm$ 0.0	0.97 $\pm$ 0.0
0.95	0.58 $\pm$ 0.0	1.0 $\pm$ 0.0	0.97 $\pm$ 0.0
0.98	0.57 $\pm$ 0.0	0.97 $\pm$ 0.0	0.97 $\pm$ 0.0
1	0.56 $\pm$ 0.0	0.92 $\pm$ 0.01	0.96 $\pm$ 0.0

**Known feedback setting with fixed  $\lambda$ .** First, we perform a sensitivity analysis on hyperparameters  $\lambda \in [0, 1]$  across the QDT family and baselines with the MAXIMIZATION strategy on the SYNTHETIC750 data set, and  $\alpha \in [0, 2]$  for B-DivRec (Equation 5). The trends for relevance and diversity depending on  $\lambda$  and  $\alpha$  are shown respectively in Figure 1 and Table 1. As expected, across all recommender systems, as  $\lambda$  increases, relevance increases, whereas (local or global) diversity globally decreases. The higher the area under the curve (AUC), the better the recommender system. Visually, the top two recommender systems are our contribution B-DivRec and MMR. Computing explicitly the areas under the curve for each metric (rounding to the closest third decimal place) shows that B-DivRec outperforms MMR on all metrics except for the global diversity (see Table 9 in Appendix). Overall, members of the QDT family outperform non-DPP baselines such as MMR and Markov DPP, especially for precision and relevance. Regarding  $\alpha$  in B-DivRec, unsurprisingly, the higher  $\alpha$ , the more (globally) diverse the recommendations, to the price of a loss in relevance and precision. Beyond a certain value of  $\alpha$  (for this data set,  $\alpha = 1$ ), the DPP can no longer sample enough diverse items, as  $\alpha$  filters out most of the items. We discuss this issue and potential solutions in Appendix B.

Now, we apply both the SAMPLING and MAXIMIZATION strategies for all recommender systems on the SYNTHETIC750 data set in Tables 2 and 3. The values of relevance for SAMPLING and of precision for MAXIMIZATION are not reported because they are similar across all algorithms but MMR (respectively,  $\text{rel} = 0.51$ , and  $\text{prec} = 1.0$  for all algorithms except for Markov DPP, where  $\text{prec} = 0.94$ , see Table 11 in Appendix). These results confirm that, for the quality-diversity trade-off, B-DivRec with the MAXIMIZATION strategy and MMR are the top contenders. Moreover, they show that the SAMPLING indeed encourages (global) diversity as evidenced by the  $\text{div}^G$  values and as reported in prior work (Kathuria et al., 2016). However, it results in a large loss in precision ( $\text{prec}$ ) and relevance ( $\text{rel}$ ), it is more time-consuming (see Appendix E), and does not quite achieve the best quality-diversity trade-off ( $\text{div}^+$ ). The SAMPLING strategy might be useful in the case when the recommender system must not recommend the same batch for users with the same embeddings and history, and when we are willing to recommend possibly irrelevant items. However, when the focus is on the optimization of the quality-diversity trade-off and computational efficiency, the MAXIMIZATION strategy might be more suitable.

Table 2: Benchmark on SYNTHETIC750 (6 users,  $B = 3$ ,  $\tau = 0.5$ ) with SAMPLING.

ALGO	PREC $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$
QDDec.	0.61 $\pm$ 0.01	0.95 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>
CondDPP	0.64 $\pm$ 0.01	0.94 $\pm$ 0.01	0.97 $\pm$ 0.01
$\varepsilon$ -Greedy	0.61 $\pm$ 0.01	0.95 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>
Markov	0.59 $\pm$ 0.01	0.17 $\pm$ 0.0	0.91 $\pm$ 0.01
MMR	<b>1.0 <math>\pm</math> 0.0</b>	<b>0.96 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>
<b>B-DivRec</b>	0.63 $\pm$ 0.01	0.95 $\pm$ 0.0	0.98 $\pm$ 0.01

Table 3: SYNTHETIC750 (6 users,  $B = 3$ ,  $\tau = 0.5$ ) with MAXIMIZATION.

REL $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$
<b>0.61 <math>\pm</math> 0.0</b>	0.9 $\pm$ 0.01	<b>1.01 <math>\pm</math> 0.0</b>
0.6 $\pm$ 0.0	0.89 $\pm$ 0.01	<b>1.01 <math>\pm</math> 0.0</b>
<b>0.61 <math>\pm</math> 0.0</b>	0.9 $\pm$ 0.01	<b>1.01 <math>\pm</math> 0.0</b>
0.58 $\pm$ 0.0	0.17 $\pm$ 0.0	0.98 $\pm$ 0.01
0.6 $\pm$ 0.0	<b>0.96 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>
<b>0.61 <math>\pm</math> 0.0</b>	0.95 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>

Table 4: Benchmark on SYNTHETIC3M (4 users,  $B = 3$ ) with MAXIMIZATION.

ALGO	REL $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$
CondDPP	0.55 $\pm$ 0.0	<b>0.97 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>
MMR	<b>0.73 <math>\pm</math> 0.0</b>	0.0 $\pm$ 0.0	0.02 $\pm$ 0.0
<b>B-DivRec</b>	0.55 $\pm$ 0.0	<b>0.97 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>

Table 5: SYNTHETIC15M (2 users,  $B = 5$ ) with MAXIMIZATION.

REL $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$
0.53 $\pm$ 0.0	<b>0.94 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>
<b>0.73 <math>\pm</math> 0.0</b>	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0
0.53 $\pm$ 0.0	<b>0.94 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>

We also ran the same type of experiments on SYNTHETIC3M and SYNTHETIC15M in Tables 4 and 5. Beyond 5 000 items, MarkovDPP (Affandi et al., 2012) is actually computationally intractable to run, hence we left it out of the benchmark for larger data sets. The observations made on SYNTHETIC750 (superiority of the MAXIMIZATION strategy over the SAMPLING one in terms of relevance) from those experiments are confirmed on larger synthetic data sets. Moreover, even if MMR might have better results on relevance-related metrics (rel, prec), this baseline clearly fails for diversity metrics (div<sup>G</sup>, div<sup>+</sup>), especially on larger data sets. As the number of items  $N$  increases, the difference in performance between algorithms from the QDT family decreases. This might be due to the fact that the more items there are, the less important is the impact of user history on the selection of items. We also applied recommender systems to MovieLens in Table 6, and PREDICT in Table 7. We only show in these tables the results with the MAXIMIZATION strategy, on account of the observations on the synthetic data sets. The full results can be found in Appendix E. On the MovieLens data set for movie recommendation, MMR is the top contender, whereas our contribution B-DivRec is the second best, improving significantly upon the conditional DPP in terms of relevance. The div<sup>G</sup> values are particularly small on MovieLens due to the fact that initial user histories are collinear: they feature similar movie embeddings, which leads to a small volume. However, on the PREDICT data set for drug repurposing, MMR is one of the worst recommender systems for all metrics, whereas B-DivRec succeeds in achieving a very good trade-off between quality and diversity, with a div<sup>+</sup> value larger than 1.

**Adaptive quality-diversity trade-off.** Finally, we evaluate our adaptive procedure for tuning the quality-diversity trade-off parameter  $\lambda \in [0, 1]$ . Note that  $\lambda$  is specific to one user: each online learner is initialized and updated along a trajectory corresponding to a single user. As such, we only consider one user, user with identifier 0, in all data sets. Then, we run B-DivRec (Equation 5) combined with the MAXIMIZATION strategy and the adaptive approach described in Section 4. We report the results in Table 8 for SYNTHETIC30k, MovieLens and PREDICT. To assess the goodness of our approach, the final tuned value  $\lambda_{\text{final}}$  is compared to the best *a posteriori* relevance

Table 6: Benchmark on MovieLens (4 users,  $B = 3$ ) with MAXIMIZATION.

ALGO	REL $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$
CondDPP	3.02 $\pm$ 0.18	0.06 $\pm$ 0.01	0.64 $\pm$ 0.05
MMR	<b>3.78 <math>\pm</math> 0.12</b>	<b>0.07 <math>\pm</math> 0.01</b>	<b>0.77 <math>\pm</math> 0.04</b>
<b>B-DivRec</b>	3.48 $\pm$ 0.13	0.06 $\pm$ 0.01	0.59 $\pm$ 0.05

Table 7: PREDICT (9 diseases/users,  $B = 3$ ) with MAXIMIZATION.

REL $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$
0.75 $\pm$ 0.02	<b>0.63 <math>\pm</math> 0.02</b>	0.88 $\pm$ 0.01
0.66 $\pm$ 0.02	0.6 $\pm$ 0.02	0.7 $\pm$ 0.02
<b>0.76 <math>\pm</math> 0.02</b>	0.57 $\pm$ 0.02	<b>1.02 <math>\pm</math> 0.02</b>



Table 8: Benchmark with and without (“non adaptive”) the adaptive tuning procedure in Section 4, applied to B-DivRec (Equation 5) with the MAXIMIZATION strategy for user 0, starting with initial value  $\lambda_0 = 0.5$ .  $\lambda_{\text{final}}$ : final relevance weight after tuning.  $\lambda_*$ : best *a posteriori* relevance weight.

DATA SET	$\lambda_{\text{final}}$	$\lambda_*$	REL $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$	TIME $\downarrow$
SYNTHETIC750	0.18	0.10	0.61 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.5 $\pm$ 0.0
(non adaptive)	–	–	0.61 $\pm$ 0.0	0.96 $\pm$ 0.0	1.01 $\pm$ 0.0	0.16 $\pm$ 0.0
SYNTHETIC30k	0.19	0.10	0.65 $\pm$ 0.0	0.96 $\pm$ 0.0	1.0 $\pm$ 0.0	29.8 $\pm$ 0.06
(non adaptive)	–	–	0.56 $\pm$ 0.0	0.96 $\pm$ 0.0	1.01 $\pm$ 0.0	1.34 $\pm$ 0.02
MovieLens	1.0	1.0	4.49 $\pm$ 0.01	0.01 $\pm$ 0.0	0.77 $\pm$ 0.01	158.1 $\pm$ 0.84
(non adaptive)	–	–	4.38 $\pm$ 0.0	0.01 $\pm$ 0.0	0.78 $\pm$ 0.0	29.7 $\pm$ 0.82
PREDICT	0.35	0.10	0.79 $\pm$ 0.01	0.49 $\pm$ 0.0	0.72 $\pm$ 0.01	1.97 $\pm$ 0.05
(non adaptive)	–	–	0.79 $\pm$ 0.01	0.49 $\pm$ 0.0	0.72 $\pm$ 0.01	0.82 $\pm$ 0.0

weight  $\lambda_*$ . This value is computed at the end of a trajectory associated with one user by solving the maximization problem in  $\lambda$  in Equation 6, and corresponds to the best deterministic action that the player could have taken in the game with Nature explained in Section 4. First, the adaptive procedure seems to be able to roughly retrieve the oracle value  $\lambda_*$  across data sets—albeit it is unlikely that it can always find it, since the oracle value relies on the *a posteriori* knowledge of the user feedback. Second, the computational cost of using the adaptive procedure instead of a fixed value of  $\lambda = 0.5$  is moderate: the runtime with the adaptive procedure is multiplied by 8 in average across all data sets, due to the approximation of matrix power (see Appendix B). Third, on SYNTHETIC30k, adaptively selecting  $\lambda$  throughout the trajectory allows us to noticeably increase relevance while trading off some diversity. On SYNTHETIC750 and PREDICT, the diversity and relevance metrics remain globally unchanged after the adaptive procedure. On MovieLens, since the users seem to be more biased towards popular movies compared to novel recommendations—as illustrated by the low value of  $\text{div}^G$  in the non-adaptive setting—the recommender system is leaning toward making popular recommendations, leading to a higher  $\lambda$ , higher relevance (rel +2.5%) to the price of some of effective diversity ( $\text{div}^+ -1.3\%$ ).

## 6 DISCUSSION

In this paper, we introduced a versatile and flexible approach for diverse recommendation. We proposed a general family of DPPs, named the QDT family, embedding the quality-diversity trade-off which includes several well-known recommender systems. Building on this family, we introduced B-DivRec which is a novel approach to integrating the user history in a computationally tractable fashion. We illustrated the versatility of our contributions by applying them in several settings: when the feedback model is known, and when the recommender system must adapt the diversity of its recommendations to the user based on user interaction with the recommended items. Our approach has been tested on both synthetic and real-life data sets for movie recommendation and drug repurposing, and performs consistently well across settings and data sets. We released as open source our algorithm and code for the experiments (see supplementary material).

Yet, many venues for research are still unresolved in the field. Recommender systems often face a pervasive issue related to missing feedback. How can one leverage information from the fact that a user has not visited nor rated a recommended item (also known as the cold-start problem (Briand et al., 2021; Wang et al., 2025) in the literature)? This issue is connected with many topics, *e.g.*, when the feedback needs to be learned on the fly (Radlinski et al., 2008; Zenati et al., 2022; Hikmawati et al., 2024; Park & Jia, 2025). However, as we have been able to successfully extend the QDT family to solve the problem of adaptively tuning the quality-diversity trade-off, improving upon this framework might be the key to solve the issue of missing feedback while still ensuring the diversity of recommended items. A first approach could implement the optimism principle as in multi-armed bandits (Auer et al., 2002; Abbasi-Yadkori et al., 2011) and building confidence intervals on the expected feedback on items in the likelihood matrix of a DPP from the QDT family.

## REFERENCES

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. {TensorFlow}: a system for {Large-Scale} machine learning. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, pp. 265–283, 2016.
- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011.
- Zeinab Abbassi, Sihem Amer-Yahia, Laks VS Lakshmanan, Sergei Vassilvitskii, and Cong Yu. Getting recommender systems to think outside the box. In *Proceedings of the third ACM conference on Recommender systems*, pp. 285–288, 2009.
- Raja Hafiz Affandi, Alex Kulesza, and Emily B Fox. Markov determinantal point processes. *arXiv preprint arXiv:1210.4850*, 2012.
- Raja Hafiz Affandi, Emily Fox, Ryan Adams, and Ben Taskar. Learning the parameters of determinantal point process kernels. In *International Conference on Machine Learning*, pp. 1224–1232. PMLR, 2014.
- Ashton Anderson, Lucas Maystre, Ian Anderson, Rishabh Mehrotra, and Mounia Lalmas. Algorithmic effects on the diversity of consumption on spotify. In *Proceedings of the web conference 2020*, pp. 2155–2165, 2020.
- Sunil Arya and David M Mount. Approximate nearest neighbor queries in fixed dimensions. In *SODA*, volume 93, pp. 271–280, 1993.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multi-armed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- Hiba Bederina and Jill-Jênn Vie. Bayesian-guided diversity in sequential sampling for recommender systems. *arXiv preprint arXiv:2506.21617*, 2025.
- Omer Ben-Porat, Lee Cohen, Liu Leqi, Zachary C Lipton, and Yishay Mansour. Modeling attrition in recommender systems with departing bandits. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pp. 6072–6079, 2022.
- Alexei Borodin and Eric M Rains. Eynard–mehta theorem, schur process, and their pfaffian analogs. *Journal of statistical physics*, 121(3):291–317, 2005.
- Léa Briand, Guillaume Salha-Galvan, Walid Bendada, Mathieu Morlon, and Viet-Anh Tran. A semi-personalized system for user cold start recommendation on music streaming apps. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 2601–2609, 2021.
- Adrian Burian, Jarmo Takala, and Mikko Ylinen. A fixed-point implementation of matrix inversion using cholesky decomposition. In *2003 46th Midwest Symposium on Circuits and Systems*, volume 3, pp. 1431–1434. IEEE, 2003.
- Daniele Calandriello, Michal Dereziński, and Michal Valko. Sampling from a k-dpp without looking at all items. *Advances in Neural Information Processing Systems*, 33:6889–6899, 2020.
- Jaime Carbonell and Jade Goldstein. The use of mmr, diversity-based reranking for reordering documents and producing summaries. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 335–336, 1998.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*, 2018.
- Y Cha, T Erez, IJ Reynolds, D Kumar, J Ross, G Koytiger, R Kusko, B Zeskind, S Risso, E Kagan, et al. Drug repurposing from the perspective of pharmaceutical companies. *British journal of pharmacology*, 175(2):168–180, 2018.

- Wei-Lun Chao, Boqing Gong, Kristen Grauman, and Fei Sha. Large-margin determinantal point processes. In *UAI*, pp. 191–200, 2015.
- Oscar Chavarriaga, Beatriz Florian-Gaviria, and Oswaldo Solarte. A recommender system for students based on social knowledge and assessment data of competences. In *European conference on technology enhanced learning*, pp. 56–69. Springer, 2014.
- Laming Chen, Guoxin Zhang, and Eric Zhou. Fast greedy map inference for determinantal point process to improve recommendation diversity. *Advances in neural information processing systems*, 31, 2018.
- Yifan Chen and Yun Yang. Fast statistical leverage score approximation in kernel ridge regression. In *International Conference on Artificial Intelligence and Statistics*, pp. 2935–2943. PMLR, 2021.
- Anirban Dasgupta, Ravi Kumar, and Tamás Sarlós. Fast locality-sensitive hashing. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1073–1081, 2011.
- Steven De Rooij, Tim Van Erven, Peter D Grünwald, and Wouter M Koolen. Follow the leader if you can, hedge if you must. *The Journal of Machine Learning Research*, 15(1):1281–1316, 2014.
- Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. The faiss library. *arXiv preprint arXiv:2401.08281*, 2024.
- Iain S Duff and Bora Uçar. Combinatorial problems in solving linear systems. *Combinatorial Scientific Computing*, 09061, 2009.
- Christophe Dupuy and Francis Bach. Learning determinantal point processes in sublinear time. In *International Conference on Artificial Intelligence and Statistics*, pp. 244–257. PMLR, 2018.
- David Duvenaud. The kernel cookbook: Advice on covariance functions. URL <https://www.cs.toronto.edu/duvenaud/cookbook>, 2014.
- Mike Gartrell, Ulrich Paquet, and Noam Koenigstein. Low-rank factorization of determinantal point processes. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.
- Guillaume Gautier, Guillermo Polito, Rémi Bardenet, and Michal Valko. Dppy: Dpp sampling with python. *J. Mach. Learn. Res.*, 20:180–1, 2019. URL <https://dppy.readthedocs.io/en/latest/>.
- Mouzhi Ge, Carla Delgado-Battenfeld, and Dietmar Jannach. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *Proceedings of the fourth ACM conference on Recommender systems*, pp. 257–260, 2010.
- Jennifer Gillenwater, Alex Kulesza, and Ben Taskar. Near-optimal map inference for determinantal point processes. *Advances in Neural Information Processing Systems*, 25, 2012.
- Boqing Gong, Wei-Lun Chao, Kristen Grauman, and Fei Sha. Diverse sequential subset selection for supervised video summarization. *Advances in neural information processing systems*, 27, 2014.
- Elena Grigorescu, Brendan Juba, Karl Wimmer, and Ning Xie. Hardness of maximum likelihood learning of dpps. In *Conference on learning theory*, pp. 3800–3819. PMLR, 2022.
- Danil Gusak, Nikita Sukhorukov, and Evgeny Frolov. Recommendation is a dish better served warm. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems*, pp. 1267–1272, 2025.
- Insu Han, Prabhanjan Kambadur, Kyoungsoo Park, and Jinwoo Shin. Faster greedy map inference for determinantal point processes. In *International Conference on Machine Learning*, pp. 1384–1393. PMLR, 2017.
- F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.

- Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pp. 639–648, 2020.
- Erna Hikmawati, Heru Nugroho, and Kridanto Surendro. Improve the quality of recommender systems based on collaborative filtering with missing data imputation. In *Proceedings of the 2024 13th International Conference on Software and Computer Applications*, pp. 75–80, 2024.
- J. Ben Hough, Manjunath Krishnapur, Yuval Peres, and Bálint Virág. Determinantal processes and independence. *Probability Surveys*, 3, January 2006. doi: 10.1214/154957806000000078.
- Leo Iaquina, Marco De Gemmis, Pasquale Lops, Giovanni Semeraro, Michele Filannino, and Piero Molino. Introducing serendipity in a content-based recommender system. In *2008 eighth international conference on hybrid intelligent systems*, pp. 168–173. IEEE, 2008.
- Carole Ibrahim, Hiba Bederina, Daniel Cuestal, Laurent Montier, Cyrille Delabre, and Jill-Jênn Vie. Diversified recommendations of cultural activities with personalized determinantal point processes. In *Second International Workshop on Recommender Systems for Sustainability and Social Good*, 2025.
- Marius Kaminskis and Derek Bridge. Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 7(1):1–42, 2016.
- Tarun Kathuria, Amit Deshpande, and Pushmeet Kohli. Batched gaussian process bandit optimization via determinantal point processes. *Advances in neural information processing systems*, 29, 2016.
- Takahiro Kawashima and Hideitsu Hino. A family of distributions of random subsets for controlling positive and negative dependence. *arXiv preprint arXiv:2408.01022*, 2024.
- David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146, 2003.
- Chun-Wa Ko, Jon Lee, and Maurice Queyranne. An exact algorithm for maximum entropy sampling. *Operations Research*, 43(4):684–691, 1995.
- Denis Kotkov, Joseph A Konstan, Qian Zhao, and Jari Veijalainen. Investigating serendipity in recommender systems based on real user feedback. In *Proceedings of the 33rd annual acm symposium on applied computing*, pp. 1341–1350, 2018.
- Denis Kotkov, Alan Medlar, Triin Kask, and Dorota Glowacka. The dark matter of serendipity in recommender systems. In *Proceedings of the 2024 Conference on Human Information Interaction and Retrieval*, pp. 108–118, 2024.
- Alex Kulesza and Ben Taskar. Structured determinantal point processes. *Advances in neural information processing systems*, 23, 2010.
- Alex Kulesza, Ben Taskar, et al. Determinantal point processes for machine learning. *Foundations and Trends® in Machine Learning*, 5(2–3):123–286, 2012. URL <https://arxiv.org/abs/1207.6083>.
- Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *nature*, 401(6755):788–791, 1999.
- Fan Li, Xu Si, Shisong Tang, Dingmin Wang, Kunyan Han, Bing Han, Guorui Zhou, Yang Song, and Hechang Chen. Contextual distillation model for diversified recommendation. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 5307–5316, 2024.

- Fanghui Liu, Xiaolin Huang, Yudong Chen, and Johan AK Suykens. Random features for kernel approximation: A survey on algorithms, theory, and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):7128–7148, 2021.
- Odile Macchi. The coincidence approach to stochastic point processes. *Advances in Applied Probability*, 7(1):83–122, 1975.
- Vered Madar. Direct formulation to cholesky decomposition of a general nonsingular correlation matrix. *Statistics & probability letters*, 103:142–147, 2015.
- Zelda E Mariet, Yaniv Ovadia, and Jasper Snoek. Dppnet: Approximating determinantal point processes with deep networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Mikhail Mironov and Liudmila Prokhorenkova. Measuring diversity: Axioms and challenges. *arXiv preprint arXiv:2410.14556*, 2024.
- Cameron Musco and Christopher Musco. Recursive sampling for the nystrom method. *Advances in neural information processing systems*, 30, 2017.
- Elvis Nava, Mojmir Mutny, and Andreas Krause. Diversified sampling for batched bayesian optimization with determinantal point processes. In *International Conference on Artificial Intelligence and Statistics*, pp. 7031–7054. PMLR, 2022.
- Evert J Nyström. Über die praktische auflösung von integralgleichungen mit anwendungen auf randwertaufgaben. *Acta Math.*, 54:185–204, 1930.
- Jiho Park and Huiwen Jia. Multi-user contextual cascading bandits for personalized recommendation. *arXiv preprint arXiv:2508.13981*, 2025.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- Rémy Poulain and Fabien Tarissan. Investigating the lack of diversity in user behavior: The case of musical content on online platforms. *Information processing & management*, 57(2):102169, 2020.
- Filip Radlinski, Robert Kleinberg, and Thorsten Joachims. Learning diverse rankings with multi-armed bandits. In *Proceedings of the 25th international conference on Machine learning*, pp. 784–791, 2008.
- Clémence Réda, Jill-Jënn Vie, and Olaf Wolkenhauer. stanscofi and benchscofi: A new standard for drug repurposing by collaborative filtering. *Journal of Open Source Software*, 9(93):5973, 2024.
- Clémence Réda. Predict drug repurposing dataset, May 2023. URL <https://doi.org/10.5281/zenodo.7983090>.
- Ulrik Schulze and Michael Ringel. What matters most in commercial success: first-in-class or best-in-class? *Nature Reviews Drug Discovery*, 12(6):419–420, 2013.
- Lindsay Spring, Kunle Demuren, Michael Ringel, and John Wu. First-in-class versus best-in-class: an update for new market dynamics. *Nat Rev Drug Discov*, 22(7):531–532, 2023.
- Hampus Gummesson Svensson, Ola Engkvist, Jon Paul Janet, Christian Tyrchan, and Morteza Haghir Chehreghani. Diverse mini-batch selection in reinforcement learning for efficient chemical exploration in de novo drug design. *arXiv preprint arXiv:2506.21158*, 2025.
- Nicolas Tremblay, Simon Barthelmé, and Pierre-Olivier Amblard. Determinantal point processes for coresets. *Journal of Machine Learning Research*, 20(168):1–70, 2019.
- Aaron Van den Oord, Sander Dieleman, and Benjamin Schrauwen. Deep content-based music recommendation. *Advances in neural information processing systems*, 26, 2013.
- Jill-Jënn Vie, Camille Lailly, and Solene Pichereau. Mangaki: an anime/manga recommender system with fast preference elicitation. *Tech. Rep.*, 2015.

- Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020.
- Dong Wang, Junyi Jiao, Arnab Bhadury, Yaping Zhang, Mingyan Gao, and Onkar Dalal. Item-centric exploration for cold start problem. *arXiv preprint arXiv:2507.09423*, 2025.
- Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. Heterogeneous graph attention network. In *The world wide web conference*, pp. 2022–2032, 2019.
- Mark Wilhelm, Ajith Ramanathan, Alexander Bonomo, Sagar Jain, Ed H Chi, and Jennifer Gillenwater. Practical diversified recommendations on youtube with determinantal point processes. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pp. 2165–2173, 2018.
- Christopher Williams and Matthias Seeger. Using the nyström method to speed up kernel machines. *Advances in neural information processing systems*, 13, 2000.
- Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- Simon N Wood. *Generalized additive models: an introduction with R*. chapman and hall/CRC, 2017.
- Zhelin Xu and Atsushi Matsumura. A serendipitous recommendation system considering user curiosity. In *International Conference on Information Integration and Web Intelligence*, pp. 33–48. Springer, 2024.
- Yueqing Xuan, Kacper Sokol, Mark Sanderson, and Jeffrey Chan. Diverse negative sampling for implicit collaborative filtering. *arXiv preprint arXiv:2508.14468*, 2025.
- Nacim Yanes, Ayman Mohamed Mostafa, Mohamed Ezz, and Saleh Naif Almuayqil. A machine learning-based recommender system for improving students learning experiences. *IEEE Access*, 8:201218–201235, 2020.
- Tianbao Yang, Yu-Feng Li, Mehrdad Mahdavi, Rong Jin, and Zhi-Hua Zhou. Nyström method vs random fourier features: A theoretical and empirical comparison. *Advances in neural information processing systems*, 25, 2012.
- Dongjin Yu, Ting Yu, Dongjing Wang, and Sixuan Wang. Long tail service recommendation based on cross-view and contrastive learning. *Expert Systems with Applications*, 238:121957, 2024.
- Yuan Yuan, Xiangtao Zheng, and Xiaoqiang Lu. Discovering diverse subset for unsupervised hyperspectral band selection. *IEEE Transactions on Image Processing*, 26(1):51–64, 2016.
- Houssam Zenati, Alberto Bietti, Eustache Diemert, Julien Mairal, Matthieu Martin, and Pierre Gailard. Efficient kernelized ucb for contextual bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 5689–5720. PMLR, 2022.
- Xueying Zhan, Qing Li, and Antoni B Chan. Multiple-criteria based active learning with fixed-size determinantal point processes. *arXiv preprint arXiv:2107.01622*, 2021.
- Reza Jafari Ziarani and Reza Ravanmehr. Serendipity in recommender systems: a systematic literature review. *Journal of Computer Science and Technology*, 36(2):375–396, 2021.

## A RELATED WORKS (CONTINUED)

Due to space constraints in the main text, we write here a more comprehensive section than in Section 2 about prior works on recommender systems with diversity.

### A.1 DIVERSITY METRICS

In Section 1, we focused on a geometrically intuitive definition of diversity relying on the volume of the parallelotope built by the columns of the kernel function applied to a given subset of items. That definition is versatile enough that, not only could we use it to define intrabatch/local and interbatch/global diversity metrics (Bederina & Vie, 2025) (see Equation 1), but it also enabled us to produce a single metric for the quality-diversity trade-off (that is, the  $\text{div}^+$  metric in Equation 2). These metrics hold for any feedback model and kernel function, hence they are very flexible. As described in Section 2, prior works (Carbonell & Goldstein, 1998; Kaminskas & Bridge, 2016) also proposed a linear or convex combination of relevance and diversity metrics, controlled by a parameter  $\lambda \leq 0$ , or a quality-diversity trade-off induced by a matrix decomposition (Kulesza & Taskar, 2010). In our paper, we actually combine the two approaches. We discuss below other metrics mentioned in the literature. Note that we make a distinction between metrics which are only evaluated in the experimental studies (*e.g.*, intralist average distance or category coverage (Ge et al., 2010; Kaminskas & Bridge, 2016; Chen et al., 2018; Li et al., 2024)) from the metrics which are optimized upon and actually used in the recommender systems. We focus on the latter in this section.

A strong competitor of our metrics are ridge leverage scores. Given the ridge hyperparameter  $\zeta \geq 0$ , for a given kernel function  $k$ , Musco & Musco (2017, Definition 1) define the ridge leverage score of item  $\phi^i \in \Omega := \{\phi^1, \phi^2, \dots, \phi^N\}$  as  $r_k^\zeta(\phi^i) := (L(L + \zeta I_N)^{-1})_{i,i}$ , where  $L := k_{\Omega, \Omega}$ . Alternatively, defining  $U \in \mathbb{R}^{N \times d}$  such that  $K = UU^\top$ , the ridge leverage score for  $\phi^i$  is the value of the ridge regression problem  $\min_{\mathbf{y} \in \mathbb{R}^N} \|\mathbf{U}_{i,\cdot} - \mathbf{y}U\|_2^2 + \zeta \|\mathbf{y}\|_2^2 \in [0, 1]$ . In other words,  $\mathbf{y}$  is the mixing vector that achieves the linear combination of rows in  $U$  closest to the  $i^{\text{th}}$  row of  $U$  (which is the row associated with  $\phi^i$ ). Then the ridge leverage score is a measure of the unicity of the vector  $\phi^i$  among elements in  $\Omega$ : the smaller it is, the least “unique”  $\phi^i$  is. However, computing the exact ridge leverage scores requires potentially inverting a matrix of size  $N \times N$  which itself has a time complexity of  $\mathcal{O}(N^3)$ . To get a measure of diversity for a whole set of items  $\mathcal{S}$ , summing all leverage scores across  $\mathcal{S}$  yields the “effective dimension”, also called “degrees of freedom”:  $d_{\text{eff},k}^\zeta(\mathcal{S}) := \sum_{\phi^i \in \mathcal{S}} r_k^\zeta(\phi^i) = \text{tr}(L(L + \zeta I_N)^{-1})$ . This measure of diversity is quite intuitive, and the time complexity is in  $\tilde{\mathcal{O}}(N)$ , where  $\tilde{\mathcal{O}}$  hides logarithmic terms in  $N$  (Chen & Yang, 2021). Then, similarly to determinants and volumes where Cholesky decompositions on sparse or low-rank matrices can be leveraged (see Appendix B), the time complexity of these measures might be linear in the number of items. However, since the determinant is directly optimized (in the MAXIMIZATION strategy) or considered (in the SAMPLING strategy) in our algorithms, an evaluation of recommender systems by the volume might be more consistent than by leverage scores.

As mentioned in Section 1, the topic of the quality-diversity trade-off is related to serendipity—which is also called out-of-the-box (Abbassi et al., 2009), unexpectedness (Xu & Matsumura, 2024), surprise (Ziarani & Ravanmehr, 2021)—where an unexpectedly good item is recommended to a user, contrary to simply recommending items which are very different from those in the user history or cover different categories (which would be diversity). Looking for serendipitous items allows us to overcome the popularity bias in recommendation (Yu et al., 2024). Kaminskas & Bridge (2016); Poulain & Tarissan (2020); Ziarani & Ravanmehr (2021); Kotkov et al. (2024) show that there is no consensus regarding the concept of serendipity, nor a single metric. The same holds for diversity metrics: recently, Mironov & Prokhorenkova (2024) tried to formalize the expected properties of a good diversity metric, showed that known diversity metrics do not match them, and proposed two matching metrics. These properties are monotonicity (the diversity of the union of two sets should be higher than the maximum diversity for each of these sets), uniqueness (replacing an element from a set with a copy of an already present element should decrease the diversity of the modified set), and continuity of the diversity metric. In particular, the determinant does not satisfy the monotonicity property. Yet the metrics that Mironov & Prokhorenkova (2024) proposed are too computationally expensive in practice. In Yu et al. (2024), authors built a recommender system for long-tail Web

services/items likely to be queried by applications/users, which is based on linear propagation of information on graphs of interconnected services with a LightGCN (He et al., 2020). This allows them to learn item and user embeddings. The predicted probability of querying a service by an application is the inner product of the corresponding embeddings. A recommender system is then trained on those predicted scores with a pairwise Bayesian Personalized Ranking (BPR) loss, which is appropriate with implicit and sparse feedback data. Other papers also study diversity through random walks on graphs (Poulain & Tarissan, 2020). Contrary to our paper, most of the reported metrics combine metrics of item unpopularity (related to relevance) and item dissimilarity (Iaquinta et al., 2008; Abbassi et al., 2009) to be applied at a reranking stage where a subset of relevant items has already been retrieved. However, as reported in Li et al. (2024), filtering too early for relevant items might cut down diversity in an irreversible way. Moreover, albeit the presence of numerous novelty-related metrics, the selection of a single metric is subjective, whereas the quality and diversity metrics are usually straightforwardly guided by the recommendation task.

## A.2 ABOUT DISTRIBUTIONS ON RANDOM SUBSETS OF POINTS

This section focuses on theoretical developments in DPPs which are relevant to our paper.

We recall that selecting  $B$  elements among  $N$  using a determinantal point process (DPP) has a time complexity of  $\mathcal{O}(NB^3)$ , after finding the eigenvalues of the likelihood matrix, which has a time complexity of  $\mathcal{O}(N^3)$ . In the case of a linear kernel on  $d$ -dimensional item embeddings, finding the eigenvalues has a time complexity of  $\mathcal{O}(Nd^2)$ . Moreover, using the dual representation of DPPs, sampling  $B$  elements among  $N$  has a time complexity of  $\mathcal{O}(NdB^2 + d^2B^3)$  (Kulesza et al., 2012), whereas the greedy approximation used for the maximization strategy (which is a NP-hard problem otherwise) has a naive time complexity of  $\mathcal{O}(B^2N)$ . However, several works (Gillenwater et al., 2012; Han et al., 2017; Chen et al., 2018) improve upon this time complexity, by leveraging the fact that the logarithm of the determinant of the likelihood matrix in the DPP should be maximized (see Equation 4) and using linear algebra approximations.  $\log \det$  is a submodular function but not monotone; instead of obtaining a  $1 - 1/e$  approximation of the exact solution (Kempe et al., 2003), proposed approaches yield an  $1/4$  approximation (Gillenwater et al., 2012), and run with a time complexity in  $\mathcal{O}(NB^2)$  (Chen et al., 2018). Moreover, Mariet et al. (2019) introduce an approximation of a DPP—without the quality-diversity trade-off—by a deep learning model, named DPPNet, which could also be used to further empirically speed up the sampling and maximization of a DPP. Implementing our contributions with a DPPNet instead of a DPP might be a future venue for enhancing their use in real-life applications.

A recent paper by Kawashima & Hino (2024) introduces a family of distributions on sets of points which includes DPPs, called discrete kernel point processes (DKPPs), where the attraction and the repulsion of sampled points can be controlled explicitly. Given a continuous function  $\psi : \mathbb{R}^+ \mapsto \mathbb{R}$  and  $k_{\Omega, \Omega}$  a positive-semidefinite Hermitian matrix, the probability of sampling the subset  $A \subseteq \Omega$  is given by  $\mathbb{P}(S = A) \propto \exp(\text{tr}(\psi(k_{A,A})))$ .  $\psi$  controls the parametrization between positive and negative correlations. DKPPs enable maximization, *i.e.*, finding the subset with highest probability, and sampling as well. However, using DPPs allows us to leverage existing implementations and algorithms. Furthermore, positive correlations between items are not necessarily linked to relevance, as the user information or context might also intervene in the quality score.

## A.3 BASELINES FOR DIVERSIFIED RECOMMENDATIONS

This section gives an overview of the state-of-the-art on diversified recommender systems. We focused on Determinantal Point Processes (DPPs) in our paper to leverage the literature regarding fast implementation of sampling and maximization algorithms, and to incorporate some flexibility in the description of item similarity thanks to the use of kernel functions. Here, we do not consider prior works where the feedback model needs to be learned on the fly (Radlinski et al., 2008; Chao et al., 2015; Kathuria et al., 2016; Nava et al., 2022; Hikmawati et al., 2024; Park & Jia, 2025).

As mentioned in Section 1, DPPs are popular in the field of recommender systems. Recently, Ibrahim et al. (2025) applied two DPPs (a traditional quality-diversity decomposition, and one with a linear kernel, without quality scores nor trade-off parameter) to output relevant and diverse recommendations of cultural goods to teenagers, again, in the sense of “intra-batch” diversity. They reported an improved diversity (volume) in the recommendation with DPPs, to the price of a notable



part of the relevance, both in offline and online/live experiments. Wilhelm et al. (2018) (using a quality-diversity decomposition with a RBF kernel) also applied DPPs on video recommendation. Authors showed that DPPs, contrary to all other baselines, yielded an increase in the number of long user sessions, which are indicative of user satisfaction with the recommendations.

Recommender systems that build upon MMR, which is a strong baseline in the field of diverse recommendations, have also been investigated in the literature. Recently, Li et al. (2024) introduced Contextual Distillation Model (CDM), which trains a surrogate attention-based deep model of Maximal Marginal Relevance (MMR), for a diversity metric correlated to the scalar product of item and user vectors. The surrogate model estimates the probability of a given item of being ranked among the top- $B$  items by the MMR score (see Section 2). Authors compared their contribution to a DPP (Chen et al., 2018) and MMR, showing that the DPP had a better performance diversity-wise, but arguing that the latter have a quadratic time complexity in the number of items  $N$ . However, we showed in our work that, provided some approximations (see Appendix B), we can actually run DPPs (B-DivRec and the conditional DPP) and MMR on libraries of dozens of millions of items. Note that Li et al. (2024) only focused on what we call “intra-batch diversity”.

In addition to (Affandi et al., 2012), there are other papers regarding the improvement of the diversity compared to a past sequence of recommended items. In addition to proposing a faster algorithm for the MAXIMIZATION strategy for DPPs, Chen et al. (2018) also studies the case when the diversity is only required among items in a sliding time window. They use a standard quality-diversity decomposition for the likelihood matrix, with the MAXIMIZATION strategy, and modify the maximization problem such that only the last  $\omega - 1$  recommended items are taken into account in the greedy algorithm. However, their approach is suitable when a long series of interactions with the user occurs and when we allow the recommender system to forget the oldest recommendations. Yet, in practice, users interact only few times with the recommender system before settling on a recommendation or dropping out (Ben-Porat et al., 2022; Gusak et al., 2025).

Finally, few papers on the literature deal with the issue of tuning the level of diversity in the recommendations to the user. Given a notion of usefulness (*i.e.*, relevance) and unexpectedness, which is somewhat related to diversity (see the previous paragraph). Xu & Matsumura (2024) studies the problem of finding the good proportions of usefulness and unexpectedness in recommendations for each user after a step of retrieval of relevant items. The latter is closely related to the challenge of addressing the adaptive quality-diversity trade-off in Section 4, albeit our procedure operates at any stage of recommendation (including retrieval) and does not feature a serendipity metric. The metric in Xu & Matsumura (2024) incorporates for each user a convex combination of long-term and short-term preferences as the usefulness-unexpectedness trade-off (“curiosity”) parameter. In that paper, this parameter is precomputed for each user, and cannot change anymore, unlike our approach.

## B PRACTICAL IMPLEMENTATION

We resorted to the Python package DPPy (Gautier et al., 2019; Calandriello et al., 2020) for the implementation of the SAMPLING strategy, whereas we used the greedy algorithm in (Chen et al., 2018) for the MAXIMIZATION approach. Please refer to the `requirements.txt` file in the code for the package versions. This section is motivated by the fact that, in Algorithm 1, the step of computation of the likelihood matrix for the DPP in Line 5 and of the recommendation set in Lines 6-7 might be expensive (at least in  $\mathcal{O}(N^3)$ ) when  $N$  is large when implemented naively.

We also used sparse matrices as implemented in the Python package SciPy (Virtanen et al., 2020) with a denuding procedure: meaning that we could make the matrix sparser by rounding up values in matrices up to the  $q^{\text{th}}$  decimal place to save memory. In practice, for our experiments, we did not have to round up values. However, for even larger libraries of items, this technique might prove useful. Below, we list the most computationally expensive parts of the code and explain the solutions considered to get a tractable implementation.

### B.1 NYSTRÖM APPROXIMATIONS OF THE KERNEL FUNCTION

One important element of the code is that a dense matrix of size  $N \times N$  should never be stored in full in the RAM. To comply with this rule, we extensively relied on the feature map  $\nu$  associated with a kernel function  $k$  (which also facilitates the computations as if  $k$  were a linear function). To compute  $\nu$ , we use a Nyström approximation (Nyström, 1930) that assumes that the rank of  $k_{\Omega, \Omega}$  is actually  $C \ll N$ . Given a dimension  $d' \ll d$ , the Nyström approximation starts by subsampling a set of  $C$  items  $\mathcal{I}$  in  $\Omega$  at random (without replacement), and then builds the eigendecomposition of the real-valued matrix  $k_{\Omega, \Omega}$ , that is, the evaluation of the kernel function on the whole universe.

$$k_{\Omega, \Omega} = U \Lambda U^\top = \begin{bmatrix} U_{\mathcal{I}} \Lambda U_{\mathcal{I}}^\top & U_{\mathcal{I}} \Lambda U_{-\mathcal{I}}^\top \\ U_{-\mathcal{I}} \Lambda U_{\mathcal{I}}^\top & U_{-\mathcal{I}} \Lambda U_{-\mathcal{I}}^\top \end{bmatrix} = \begin{bmatrix} k_{\mathcal{I}, \mathcal{I}} & k_{\mathcal{I}, -\mathcal{I}} \\ k_{-\mathcal{I}, \mathcal{I}}^\top & k_{-\mathcal{I}, -\mathcal{I}} \end{bmatrix},$$

where  $U_{\mathcal{I}}$  (respectively,  $U_{-\mathcal{I}}$ ) is the part of the orthonormal basis of the eigendecomposition associated with items in  $\mathcal{I}$  (resp., not in  $\mathcal{I}$ ), and  $\Lambda$  is the diagonal matrix of eigenvalues of  $k_{\Omega, \Omega}$ . When one needs to compute  $k_{\mathcal{S}, \mathcal{S}}$  for any subset  $\mathcal{S} \subset \Omega$  of size  $B$ , using the fact that  $k_{\mathcal{S}, \mathcal{S}} = (k_{\mathcal{S}, \mathcal{I}}^\top U_{\mathcal{I}} \Lambda^{-1/2}) (k_{\mathcal{S}, \mathcal{I}}^\top U_{\mathcal{I}} \Lambda^{-1/2})^\top$  by replacing  $U_{-\mathcal{I}}$  by  $k_{-\mathcal{I}, \mathcal{I}} U_{\mathcal{I}} \Lambda^{-1}$ .  $\Lambda^{-1/2}$  consists in simply applying the square root function to the eigenvalues of  $k_{\Omega, \Omega}$ , whereas  $k_{\mathcal{S}, \mathcal{I}} \in \mathbb{R}^{B \times d'}$  where  $d' \ll N$  can be explicitly evaluated. All in all, the corresponding  $\nu$  function is  $\nu(\mathcal{S}) = k_{\mathcal{S}, \mathcal{I}}^\top U_{\mathcal{I}} \Lambda^{-1/2}$  for any  $\mathcal{S} \subset \Omega$ .

As previously mentioned in the main text, the time complexity of the Nyström approximation is  $\mathcal{O}(N(d')^2 + (d')^3)$ . In practice, we use the implementation of the Nyström approximation in the Python package scikit-learn (Pedregosa et al., 2011). The computation of  $\mathcal{I}$ ,  $U_{\mathcal{I}}$  and  $\Lambda$  happens only once before any recommendation is made. We use  $d' = 100$  for the rank in the Nyström approximation. We note that, while we use random elements to build the approximation, other papers suggest to consider more representative points (Tremblay et al., 2019), for instance, those with maximum leverage score which would be more representative of the set. However, this approach can be costly (see Appendix A) which is why we elected to stick to random sampling.

Note that the computation of the volume (see metrics in Equation 1) also relies on the Nyström approximation, as the full matrix  $k_{\mathcal{S}, \mathcal{S}}$  is not built. This accounts for the fact that, although Assumption 3.3 mentions that the item embeddings have a  $\ell_2$ -norm equal to 1, the actual computed volumes in the experiments in Section 5 might exceed 1. The Nyström approximation does not preserve the property on the norm of the item embeddings.

### B.2 APPROXIMATE MATRIX INVERSIONS AND DETERMINANTS

For the conditional and the k-Markov (Affandi et al., 2012) DPPs, there is a step of inversion of a positive definite matrix of potentially size  $N \times N$ , which is extremely costly in a naive implementation. As suggested by prior works (Burian et al., 2003), we resort to Cholesky decompositions of the matrices that we wish to inverse. The Cholesky decomposition of a positive-definite matrix  $M$  of size  $n$  is  $M = R^\top R$  where  $R$  is an upper (or lower) triangular matrix of size  $n$ . Then, as

$M^{-1} = R^{-1}(R^{-1})^\top$ , it suffices to solve  $RX = I_n$  in  $X$  by forward substitution, because  $R$  is a triangular matrix, to obtain the inverse of  $M$ . To compute the (log-)determinant of  $M$  of size  $n$ , the following relation is true:  $\log \det(M) = 2 \sum_{i=1}^n \log R_{i,i}$ , because  $R$  is a triangular matrix with positive diagonal elements (Madar, 2015) and  $\det R = \det R^\top$ .

In our implementation, we use the Python package Scikit-Sparse. For the conditional DPP, the Nyström approximation of the kernel function can be leveraged on top of the Cholesky decomposition (decomposing  $\tilde{A}^{1/2} \tilde{A}^{1/2} \approx A$  instead of  $A$ , which is less expensive due to the low-rank assumption for  $\tilde{A}^{1/2}$ ), which is why the conditional DPP can be run on large data sets contrary to Markov DPP. Indeed, the theoretical time complexity of the Cholesky decomposition is in  $\mathcal{O}(N^3)$  as a general rule. However, for sparse matrices, a fill-reducing Cholesky decomposition can be computed by reordering rows to restrict the creation of new nonzero elements (Duff & Uçar, 2009), and for sparse matrices, the practical time complexity is typically much smaller than cubic in  $N$ .

### B.3 APPROXIMATE MATRIX POWER

Matrix power intervenes in Equation 3 which describes the likelihood matrix for a DPP of the QDT family for  $\lambda \neq 1/2$ , where the power is possibly real-valued. The exact computation would require computing the whole set of eigenvalues and eigenvectors, which is quickly intractable as  $N$  increases. Similarly to Appendix B.1, we make a low-rank assumption and consider the truncated Singular Value Decomposition (SVD) of rank  $r \ll N$  of a real-valued matrix  $M$  where  $M$  is a square real matrix of size  $N$ . The SVD yields  $M = U\Lambda V^\top$ , where  $U, V \in \mathbb{R}^{r \times N}$  are orthogonal matrices. Then we compute  $V\Lambda^p V^\top$  as a proxy for  $M^p$ , that is,  $M$  multiplied  $p$  times where  $p \in \mathbb{R}^+$ . Indeed,  $M^p \approx (V^\top \Lambda V)^p = \underbrace{(V^\top \Lambda V)(V^\top \Lambda V) \dots (V^\top \Lambda V)}_{p \text{ times}} = V^\top \Lambda^p V$ , since  $VV^\top = I_r$

and  $V^\top V = I_N$ . This operation has a time complexity in  $\mathcal{O}(Nr^2)$  instead of the naive time complexity  $\mathcal{O}(N^3)$  which is dominated by the computation of the full eigendecomposition. Moreover, if we need to find  $L$  instead, where  $LL^\top = M^p$ , we can output  $L = V\sqrt{\Lambda^p}$  where  $\sqrt{\Lambda^p}$  is such that  $\sqrt{\Lambda^p}\sqrt{\Lambda^p} = \Lambda^p$ . Note that, for the diagonal matrix  $\Lambda^p$ ,  $\sqrt{\Lambda^p}$  is the result of the square root function applied element-wise to the diagonal elements of  $\Lambda^p$ . We choose  $r = N - 1$  if  $N \leq 1000$  and  $r = 100$  otherwise.

### B.4 APPROXIMATE CLOSEST-NEIGHBOR FINDING ALGORITHMS

B-DivRec requires the computation of the maximum cosine distance between feature maps of items from the current batch and the user history (see Equation 5). Obviously, computing all cosine distances between any item and all items in the user history would be expensive since it is performed for all items in  $\Omega$ . As mentioned in Section 3, many approaches allow us to retrieve the closest neighbor of a point in a set. Our implementation uses FAISS Douze et al. (2024), which is suitable for large data sets. In practice, we do not even need to build a new tree for each user and each recommendation round, but only a single tree on all items in  $\Omega$ . Indeed, queries to FAISS can be made while ignoring a subset of items when querying for the closest neighbor (in our case, ignoring all items which are not in the current user’s history).

### B.5 DIVERSITY ISSUES IN DPPs

As mentioned in Section 5, beyond a certain value of  $\alpha$ , B-DivRec can no longer sample enough diverse items, as  $\alpha$  filters out most of the items. This issue is not specific to our contribution, as Ibrahim et al. (2025) has also reported it for other conditional DPPs. It arises when the rank of the likelihood matrix ( $L$ -matrix) of the DPP is strictly smaller than  $B$ .

Possible solutions to bypass this issue would be (1) to return a batch of recommended items of size smaller than  $B$  (with a risk of no longer recommending any item at some point), (2) to only focus on quality and set the matrix  $f(k, \mathcal{S}, \mathcal{H}) \leftarrow I_N$ , (3) to forget some of the user history, and only preserve either the most representative points or the latest recommendations. Future work on this topic would be of interest for practical recommender systems.

## B.6 RETRIEVING EMBEDDINGS

Finally, our implementation cannot afford to store the full item embedding matrix  $\Phi = [\phi^i, i \in \Omega] \in \mathbb{R}^{N \times d}$  in the RAM. First, we store the user histories in files, and update them in-place. Second, we store, access and read the item embedding matrix by batches in memory with .csv files. Moreover, for the largest data set of our paper (the synthetic data set SYNTHETIC15M with five millions items), we use the .npy binary format instead of .csv for faster input/output operations.

## C ADAPTIVE QUALITY-DIVERSITY TRADE-OFF

We recall here the online learning game for the adaptive quality-diversity trade-off (see Section 4). The player is the recommender system, and Nature is the interaction with *one* user  $(\mathbf{h}, \mathcal{H}^t)$  querying the recommender system at time  $t$  (Auer et al., 2002; De Rooij et al., 2014). At round  $t \leq T$  of the game, the recommender system chooses a value of  $\lambda^t \in [0, 1]$ . Then, the recommender system makes a recommendation  $\mathcal{S}^t$  to user  $\mathbf{h}$  and receives a loss value related to the quality-diversity trade-off achieved, depending on the vector of feedback  $\mathbf{y}^t$  from Nature. The recommender system should use this loss to update and use  $\lambda^{t+1} \in [0, 1]$  at the next round. The game ends after  $T$  interactions with the same user.

Applying the AdaHedge algorithm (De Rooij et al., 2014) to our problem, we consider two experts. One favors diversity, whereas the other prefers quality. We denote the posterior weight vector associated with those two experts  $\boldsymbol{\lambda}^t := [1 - \lambda^t, \lambda^t]$  at time  $t$ . In our case, we consider a gain instead of a loss function—since the optimization problem in Equation 6 is a maximization problem. The function (to maximize) at time  $t$  with  $\lambda^t$  is

$$\begin{aligned} \text{scr}_{f, \mathbf{y}^t}^{\lambda^t}(\mathcal{S}^t; \mathbf{h}^t) &= 4(1 - \lambda^t) \log \text{vol}(f(k, \mathcal{S}^t, \mathcal{H}^t)) + 4\lambda^t \log \det(\text{diag}(\{y_k^t\}_{k \leq B})) \\ &= 4(1 - \lambda^t) \log \text{vol}(f(k, \mathcal{S}^t, \mathcal{H}^t)) + 4\lambda^t \sum_{k \leq B} \log y_k^t, \end{aligned} \quad (7)$$

where  $f = f^{\text{B-DivRec}}$  (note however that this approach straightforwardly works for any DPP of the QDT family as described in Equation 3). Now, since  $\text{scr}_{f, \mathbf{y}^t}^{\lambda^t}(\mathcal{S}^t; \mathbf{h}^t)$  is clearly linear in  $\lambda^t$  (hence concave), we can apply the gradient trick as follows

$$\begin{aligned} \mathcal{R}^{\text{adapt}}(T; \mathbf{h}) &:= \max_{\lambda \in [0, 1]} \sum_{t \leq T} \text{scr}_{f, \mathbf{y}^t}^{\lambda}(\mathcal{S}^t; \mathbf{h}) - \text{scr}_{f, \mathbf{y}^t}^{\lambda^t}(\mathcal{S}^t; \mathbf{h}) \\ &\leq \max_{\lambda \in [0, 1]} \sum_{t \leq T} (\lambda - \lambda^t) \nabla_{\lambda} \text{scr}_{f, \mathbf{y}^t}^{\lambda^t}(\mathcal{S}^t; \mathbf{h}). \end{aligned}$$

The gradient  $\nabla_{\lambda} \text{scr}_{f, \mathbf{y}^t}^{\lambda}(\mathcal{S}^t; \mathbf{h})$  with respect to  $\lambda \in [0, 1]$  is straightforward to compute

$$\nabla_{\lambda} \text{scr}_{f, \mathbf{y}^t}^{\lambda}(\mathcal{S}^t; \mathbf{h}^t) := -4 \log \text{vol}(f(k, \mathcal{S}^t, \mathcal{H}^t)) + 4 \sum_{k \leq B} \log y_k^t. \quad (8)$$

Then, we obtained an upper bound on the quantity that we want to optimize in the form of a inner product. This is the motivation for using as a gain function at time  $t$

$$g_t : \boldsymbol{\lambda} \in \triangle_2 \mapsto [-\nabla_{\lambda} \text{scr}_{f, \mathbf{y}^t}^{\lambda}(\mathcal{S}^t; \mathbf{h}^t), \nabla_{\lambda} \text{scr}_{f, \mathbf{y}^t}^{\lambda}(\mathcal{S}^t; \mathbf{h}^t)],$$

where  $\triangle_2$  is the simplex of dimension 2. Note that the first coordinate of  $g_t(\boldsymbol{\lambda})_1$  is actually the gradient of  $\text{scr}_{f, \mathbf{y}^t}^{\lambda}(\mathcal{S}^t; \mathbf{h})$  with respect to  $1 - \lambda$  (proven with a change of variable). Then in the AdaHedge learner, at time  $t$ , we update the posterior weight with

$$-\langle \boldsymbol{\lambda}^t, g_t(\boldsymbol{\lambda}) \rangle = -((1 - \lambda^t)(-g_t(\boldsymbol{\lambda})_2) + \lambda^t g_t(\boldsymbol{\lambda})_2) = (1 - 2\lambda^t)g_t(\boldsymbol{\lambda})_2 = (1 - 2\lambda^t)\nabla_{\lambda} \text{scr}_{f, \mathbf{y}^t}^{\lambda}(\mathcal{S}^t; \mathbf{h}^t),$$

which corresponds to Line 8 in Algorithm 1. Then, we obtain a new value  $\lambda^{t+1} \in \triangle_2$ , from which we extract the second coordinate  $\lambda^{t+1}$  for the next round.

## D SUPPLEMENTARY MATERIAL ON THE EXPERIMENTAL STUDY

All experiments have been run on a remote server (configuration: processor Intel Xeon Processor (Skylake, IBRS), 12 cores @2.4GHz, RAM 23GB).

### D.1 GENERATION OF SYNTHETIC DATA SETS

The synthetic data sets aim at providing an easy approach to test the scalability of the algorithms and the diversity in recommendations. The key idea is that  $d$ -dimensional item embeddings  $(\phi^1, \dots, \phi^N)$  and user contexts  $(\mathbf{h}^1, \dots, \mathbf{h}^{N_u})$  are sampled at random (hence  $m = d$ ), and the feedback value for the item-user pair  $(\phi, \mathbf{h})$  is  $((\phi)^\top \mathbf{h} + 1)/2 \in [0, 1]$ . We actually generate  $B$  groups of items, with the expectation that, when outputting a batch of  $B$  recommendations, a good recommender system for diversity would sample elements from each group.

Given a desired number of items  $N$  and a batch size  $B$ , we first sample at random and independently  $\lceil N/B \rceil \times d$  values from  $\mathcal{N}(0, 2)$  to build a first matrix of  $\ell_2$ -normalized item embeddings  $\Phi^1 := [\phi^1, \phi^2, \dots, \phi^{\lceil N/B \rceil}] \in \mathbb{R}^{\lceil N/B \rceil \times d}$ . Then, for the  $\ell^{\text{th}}$  group,  $\ell = 2, \dots, B$ , we construct the matrix  $\Phi^\ell := \Phi^1 + 0.01\ell \mathbf{1}_{\lceil N/B \rceil \times d}$ , where  $\mathbf{1}_{n \times d}$  is the matrix of size  $n \times d$  which coefficients are ones. Finally, we concatenate all (renormalized) matrices  $\Phi^1, \Phi^2, \dots, \Phi^B$  to obtain  $\Phi \in \mathbb{R}^{N \times d}$ . For users, given a maximum number of users  $N_u$ ,  $N_u \times d$  values are sampled from  $\mathcal{N}(0, 1)$  to build the user contexts, after  $\ell_2$ -normalization.

### D.2 TRAINING ON REAL-LIFE DATA SETS

In this section, we delve more into detail concerning the two real-life data sets for movie recommendation and drug repurposing. In both cases, the important part is to learn the feedback model  $q_\Theta$  based on prior user-item interactions (e.g., previously rated movies or terminated clinical trials, collected offline by browsing clinical trial registries). The feedback model is then used in the computation of the likelihood matrix in Line 5 in Algorithm 1. See the `README.md` file for instructions to download the files related to those data sets.

**Movie recommendation: MovieLens.** The MovieLens data set (Harper & Konstan, 2015) that we considered has 9,725 items, 610 users and 100,837 nonzero ratings in  $\{1, 2, \dots, 5\}$ . The higher the rating, the more the user liked the movie. Metadata about the items/movies are the movie title, year of release, and the genre tags (e.g., comedy, action). First, 512-dimensional item embeddings of norm 1 are built by applying Universal Sentence Encoder (Cer et al., 2018) to the movie metadata, using Python package Tensorflow Hub (Abadi et al., 2016). Second, we split the list of item-user pairs with nonzero ratings into a training set (75%) and a testing set (25%) at random. We fit a SVD model on the training set using the Python package Mangaki (Vie et al., 2015), which fills the zeroes in the ratings matrix of MovieLens. We computed the Root Mean Squared Error on the testing set, yielding RMSE= 0.96. Knowing that the ratings are in  $\{1, 2, \dots, 5\}$ , it means that, in average, the mistake made by the SVD model on the predicted rating is not enough to accidentally classify a very bad movie as a good one. We considered a classification threshold at  $\tau = 2.5$ .

**Drug repurposing: PREDICT.** The PREDICT data set (Réda, 2023) comprises information about 1,577 drugs and 1,070 diseases, along with the status of Phase 2 or 3 clinical trials involving each drug-disease pair: 0 means that the pair has not been tested in a Phase 2 or 3 clinical trial, 1 means that the clinical trial was successful in showing that the drug has a therapeutic effect on the disease, and -1 means that the clinical trial failed (e.g., low accrual, emergence of adverse side effects). There are six types of drug features and five types of disease features, corresponding to the similarity of a drug (respectively, a disease) to another drug (resp., disease) in the data set with respect to each type of feature. We fitted the Heterogeneous Attention Network (HAN) algorithm (Wang et al., 2019) as feedback model to this data set, using the Python package stanscofi Réda et al. (2024), splitting nonzero ratings at random into training (80%) and testing (20%) sets. The Area Under the Curve on the testing set was 0.92.

Table 9: Sensitivity analysis for  $\lambda$  on SYNTHETIC750 (6 users,  $B = 3$ ) with the MAXIMIZATION strategy. Area Under the Curve for the function of the metric depending on the value of  $\lambda \in [0, 1]$ : REL: relevance. PREC: precision with threshold  $\tau = 0.5$ .  $DIV^L$ : intrabatch/local diversity.  $DIV^G$ : interbatch/global diversity.  $DIV^+$ : effective diversity. All values are rounded to the closest third decimal place.

	AUC	REL $\uparrow$	PREC $\uparrow$	$DIV^L$ $\uparrow$	$DIV^G$ $\uparrow$	$DIV^+$ $\uparrow$
QDDecomp.		<b>0.606</b>	<b>1.000</b>	<b>1.008</b>	0.928	<b>1.008</b>
CondDPP		0.601	<b>1.000</b>	<u>1.007</u>	0.920	<u>1.007</u>
$\varepsilon$ -Greedy		<b>0.606</b>	<b>1.000</b>	<b>1.008</b>	0.928	<b>1.008</b>
MarkovDPP		0.586	<u>0.935</u>	1.000	0.165	1.000
MMR		0.585	<u>0.926</u>	0.998	<b>0.960</b>	0.998
<b>B-DivRec</b>		<u>0.605</u>	<b>1.000</b>	<u>1.007</u>	<u>0.954</u>	<u>1.007</u>

Table 10: Sensitivity analysis for  $\lambda$  and  $\alpha$  on SYNTHETIC750 (6 users,  $B = 3$ ) with B-DivRec and the MAXIMIZATION strategy. REL: relevance. PREC: precision with threshold  $\tau = 0.5$ .  $DIV^L$ : intrabatch/local diversity.  $DIV^G$ : interbatch/global diversity.  $DIV^+$ : effective diversity. TIME: runtime. The difference in runtime for  $\lambda = 0.5$  and  $\lambda \neq 0.5$  is due to the approximation for the matrix power (see Appendix B).

	LAMBDA	REL $\uparrow$	PREC $\uparrow$	$DIV^L$ $\uparrow$	$DIV^G$ $\uparrow$	$DIV^+$ $\uparrow$	TIME $\downarrow$
$\lambda = 0.0$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.45 $\pm$ 0.0
$\lambda = 0.1$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.46 $\pm$ 0.0
$\lambda = 0.3$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.45 $\pm$ 0.0
$\lambda = 0.5$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.12 $\pm$ 0.0
$\lambda = 0.8$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.54 $\pm$ 0.0
$\lambda = 0.9$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.54 $\pm$ 0.0
$\lambda = 1.0$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.60 $\pm$ 0.0
<b>ALPHA</b>							
$\alpha = 0.0$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.11 $\pm$ 0.0
$\alpha = 0.1$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.12 $\pm$ 0.0
$\alpha = 0.5$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.12 $\pm$ 0.0
$\alpha = 0.8$		0.61 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.96 $\pm$ 0.0	1.01 $\pm$ 0.0	0.12 $\pm$ 0.0
$\alpha = 0.9$		0.6 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.97 $\pm$ 0.0	1.01 $\pm$ 0.0	0.12 $\pm$ 0.0
$\alpha = 0.95$		0.58 $\pm$ 0.0	1.0 $\pm$ 0.0	1.01 $\pm$ 0.0	0.97 $\pm$ 0.0	1.01 $\pm$ 0.0	0.12 $\pm$ 0.0
$\alpha = 0.98$		0.57 $\pm$ 0.0	0.97 $\pm$ 0.0	1.01 $\pm$ 0.0	0.97 $\pm$ 0.0	1.01 $\pm$ 0.0	0.12 $\pm$ 0.0
$\alpha = 1.0$		0.56 $\pm$ 0.0	0.92 $\pm$ 0.01	1.01 $\pm$ 0.0	0.96 $\pm$ 0.0	1.0 $\pm$ 0.0	0.12 $\pm$ 0.0

## E SUPPLEMENTARY NUMERICAL RESULTS

We report here the full tables of numerical results presented in Section 5 in the main text.

Table 11: Benchmark on SYNTHETIC750 (6 users,  $B = 3$ ). REL: relevance. PREC: precision with  $\tau = 0.5$ . DIV<sup>L</sup>: intrabatch/local diversity. DIV<sup>G</sup>: interbatch/global diversity. DIV<sup>+</sup>: effective diversity. TIME: runtime.

SAMPLING	REL $\uparrow$	PREC $\uparrow$	DIV <sup>L</sup> $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$	TIME $\downarrow$
QDDecomp.	0.51 $\pm$ 0.0	0.61 $\pm$ 0.01	<b>1.01 <math>\pm</math> 0.0</b>	0.95 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>	0.06 $\pm$ 0.0
CondDPP	0.51 $\pm$ 0.0	0.64 $\pm$ 0.01	1.0 $\pm$ 0.0	0.94 $\pm$ 0.01	0.97 $\pm$ 0.01	0.07 $\pm$ 0.0
$\varepsilon$ -Greedy	0.51 $\pm$ 0.0	0.61 $\pm$ 0.01	<b>1.01 <math>\pm</math> 0.0</b>	0.95 $\pm$ 0.0	1.01 $\pm$ 0.0	0.06 $\pm$ 0.0
MarkovDPP	0.51 $\pm$ 0.0	0.59 $\pm$ 0.01	1.0 $\pm$ 0.0	0.17 $\pm$ 0.0	0.91 $\pm$ 0.01	0.08 $\pm$ 0.0
<b>B-DivRec</b>	0.51 $\pm$ 0.0	0.63 $\pm$ 0.01	<b>1.01 <math>\pm</math> 0.0</b>	0.95 $\pm$ 0.0	0.98 $\pm$ 0.01	0.14 $\pm$ 0.0
<b>MAXIMIZ.</b>						
QDDecomp.	<b>0.61 <math>\pm</math> 0.0</b>	<b>1.0 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	0.9 $\pm$ 0.01	<b>1.01 <math>\pm</math> 0.0</b>	0.03 $\pm$ 0.0
CondDPP	0.6 $\pm$ 0.0	<b>1.0 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	0.89 $\pm$ 0.01	<b>1.01 <math>\pm</math> 0.0</b>	0.04 $\pm$ 0.0
$\varepsilon$ -Greedy	<b>0.61 <math>\pm</math> 0.0</b>	<b>1.0 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	0.9 $\pm$ 0.01	<b>1.01 <math>\pm</math> 0.0</b>	0.03 $\pm$ 0.0
MarkovDPP	0.58 $\pm$ 0.0	0.94 $\pm$ 0.01	<b>1.01 <math>\pm</math> 0.0</b>	0.17 $\pm$ 0.0	0.98 $\pm$ 0.01	0.06 $\pm$ 0.0
<b>B-DivRec</b>	<b>0.61 <math>\pm</math> 0.0</b>	<b>1.0 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	0.95 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>	0.12 $\pm$ 0.0
<b>NO DPP</b>						
MMR	0.6 $\pm$ 0.0	<b>1.0 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	<b>0.96 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	0.08 $\pm$ 0.0

Table 12: Benchmark on SYNTHETIC30k (4 users,  $B = 3$ ). REL: relevance. PREC: precision with  $\tau = 0.5$ . DIV<sup>L</sup>: intrabatch/local diversity. DIV<sup>G</sup>: interbatch/global diversity. DIV<sup>+</sup>: effective diversity. TIME: runtime.

SAMPLING	REL $\uparrow$	PREC $\uparrow$	DIV <sup>L</sup> $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$	TIME $\downarrow$
QDDecomp.	0.5 $\pm$ 0.0	0.33 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>	0.96 $\pm$ 0.0	1.0 $\pm$ 0.0	47.05 $\pm$ 0.15
CondDPP	0.5 $\pm$ 0.0	0.33 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>	0.96 $\pm$ 0.0	1.0 $\pm$ 0.0	47.19 $\pm$ 0.12
$\varepsilon$ -Greedy	0.5 $\pm$ 0.0	0.33 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>	0.96 $\pm$ 0.0	1.0 $\pm$ 0.0	47.49 $\pm$ 0.32
<b>B-DivRec</b>	0.5 $\pm$ 0.0	0.33 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>	0.96 $\pm$ 0.0	1.0 $\pm$ 0.0	48.1 $\pm$ 0.12
<b>MAXIMIZ.</b>						
QDDecomp.	0.54 $\pm$ 0.0	0.92 $\pm$ 0.02	<b>1.01 <math>\pm</math> 0.0</b>	0.96 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>	0.79 $\pm$ 0.0
CondDPP	0.54 $\pm$ 0.0	0.92 $\pm$ 0.02	<b>1.01 <math>\pm</math> 0.0</b>	0.96 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>	1.05 $\pm$ 0.01
$\varepsilon$ -Greedy	0.54 $\pm$ 0.0	0.92 $\pm$ 0.02	<b>1.01 <math>\pm</math> 0.0</b>	0.96 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>	0.79 $\pm$ 0.0
<b>B-DivRec</b>	0.54 $\pm$ 0.0	0.92 $\pm$ 0.02	<b>1.01 <math>\pm</math> 0.0</b>	0.96 $\pm$ 0.0	<b>1.01 <math>\pm</math> 0.0</b>	1.45 $\pm$ 0.0
<b>NO DPP</b>						
MMR	<b>0.68 <math>\pm</math> 0.0</b>	<b>1.0 <math>\pm</math> 0.0</b>	1.0 $\pm$ 0.0	0.96 $\pm$ 0.0	1.0 $\pm$ 0.0	1.73 $\pm$ 0.0



Table 13: Benchmark on SYNTHETIC3M (4 users,  $B = 3$ ). REL: relevance. PREC: precision with  $\tau = 0.5$ . DIV<sup>L</sup>: intrabatch/local diversity. DIV<sup>G</sup>: interbatch/global diversity. DIV<sup>+</sup>: effective diversity. TIME: runtime.

SAMPLING	REL $\uparrow$	PREC $\uparrow$	DIV <sup>L</sup> $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$	TIME $\downarrow$
QDDDecomp.	0.5 $\pm$ 0.0	<u>0.58 <math>\pm</math> 0.03</u>	<b>1.01 <math>\pm</math> 0.0</b>	<b>0.97 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	248.8 $\pm$ 10.44
CondDPP	0.5 $\pm$ 0.0	<u>0.58 <math>\pm</math> 0.03</u>	<b>1.01 <math>\pm</math> 0.0</b>	<b>0.97 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	331.23 $\pm$ 6.84
$\epsilon$ -Greedy	0.5 $\pm$ 0.0	<u>0.58 <math>\pm</math> 0.03</u>	<b>1.01 <math>\pm</math> 0.0</b>	<b>0.97 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	235.06 $\pm$ 4.81
<b>B-DivRec</b>	0.5 $\pm$ 0.0	<u>0.58 <math>\pm</math> 0.03</u>	<b>1.01 <math>\pm</math> 0.0</b>	<b>0.97 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	383.04 $\pm$ 16.32
<b>MAXIMIZ.</b>						
QDDDecomp.	<u>0.55 <math>\pm</math> 0.0</u>	<b>1.0 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	<b>0.97 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	232.0 $\pm$ 4.95
CondDPP	<u>0.55 <math>\pm</math> 0.0</u>	<b>1.0 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	<b>0.97 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	297.51 $\pm$ 6.56
$\epsilon$ -Greedy	<u>0.55 <math>\pm</math> 0.0</u>	<b>1.0 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	<b>0.97 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	224.08 $\pm$ 7.87
<b>B-DivRec</b>	<u>0.55 <math>\pm</math> 0.0</u>	<b>1.0 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	<b>0.97 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	372.8 $\pm$ 18.67
<b>NO DPP</b>						
MMR	<b>0.73 <math>\pm</math> 0.0</b>	<b>1.0 <math>\pm</math> 0.0</b>	<u>0.02 <math>\pm</math> 0.0</u>	<u>0.0 <math>\pm</math> 0.0</u>	<u>0.02 <math>\pm</math> 0.0</u>	569.72 $\pm$ 12.82

Table 14: Benchmark on SYNTHETIC15M (2 users,  $B = 5$ ). REL: relevance. PREC: precision with  $\tau = 0.5$ . DIV<sup>L</sup>: intrabatch/local diversity. DIV<sup>G</sup>: interbatch/global diversity. DIV<sup>+</sup>: effective diversity. TIME: runtime. As explained in Appendix B, we further worked on the input/output operations to speed up even more the algorithms, hence the difference in runtime compared to all other experiments. This change has only an impact on the runtime and not on other metrics.

MAXIMIZ.	REL $\uparrow$	PREC $\uparrow$	DIV <sup>L</sup> $\uparrow$	DIV <sup>G</sup> $\uparrow$	DIV <sup>+</sup> $\uparrow$	TIME $\downarrow$
CondDPP	<u>0.53 <math>\pm</math> 0.0</u>	<u>0.8 <math>\pm</math> 0.04</u>	<b>1.0 <math>\pm</math> 0.0</b>	<b>0.94 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	63.03 $\pm$ 0.58
<b>B-DivRec</b>	<u>0.53 <math>\pm</math> 0.0</u>	<u>0.8 <math>\pm</math> 0.04</u>	<b>1.0 <math>\pm</math> 0.0</b>	<b>0.94 <math>\pm</math> 0.0</b>	<b>1.01 <math>\pm</math> 0.0</b>	251.66 $\pm$ 0.73
<b>NO DPP</b>						
MMR	<b>0.73 <math>\pm</math> 0.0</b>	<b>1.0 <math>\pm</math> 0.0</b>	<u>0.0 <math>\pm</math> 0.0</u>	<u>0.0 <math>\pm</math> 0.0</u>	<u>0.0 <math>\pm</math> 0.0</u>	98.2 $\pm$ 0.21

Table 15: Benchmark on MovieLens (4 movie streamers/users,  $B = 3$ ). REL: relevance. PREC: precision with  $\tau = 2.5$ .  $DIV^L$ : intrabatch/local diversity.  $DIV^G$ : interbatch/global diversity.  $DIV^+$ : effective diversity. TIME: runtime.

SAMPLING	REL $\uparrow$	PREC $\uparrow$	$DIV^L$ $\uparrow$	$DIV^G$ $\uparrow$	$DIV^+$ $\uparrow$	TIME $\downarrow$
QDDDecomp.	$1.0 \pm 0.0$	$0.0 \pm 0.0$	$0.84 \pm 0.01$	$0.05 \pm 0.01$	$0.0 \pm 0.0$	$74.88 \pm 2.09$
CondDPP	$1.01 \pm 0.0$	$0.0 \pm 0.0$	$0.85 \pm 0.0$	$0.06 \pm 0.01$	$0.01 \pm 0.0$	$77.39 \pm 3.05$
$\varepsilon$ -Greedy	$1.0 \pm 0.0$	$0.0 \pm 0.0$	$0.84 \pm 0.01$	$0.05 \pm 0.01$	$0.0 \pm 0.0$	$66.82 \pm 2.25$
<b>B-DivRec</b>	$1.0 \pm 0.0$	$0.0 \pm 0.0$	<b><math>0.86 \pm 0.0</math></b>	$0.06 \pm 0.01$	$0.0 \pm 0.0$	$67.5 \pm 2.29$
<b>MAXIMIZ.</b>						
QDDDecomp.	$3.48 \pm 0.13$	$0.75 \pm 0.06$	$0.81 \pm 0.01$	$0.06 \pm 0.01$	$0.59 \pm 0.05$	$24.36 \pm 0.44$
CondDPP	$3.02 \pm 0.18$	$0.65 \pm 0.06$	$0.83 \pm 0.01$	$0.06 \pm 0.01$	$0.64 \pm 0.05$	$25.48 \pm 0.52$
$\varepsilon$ -Greedy	$3.48 \pm 0.13$	$0.75 \pm 0.06$	$0.81 \pm 0.01$	$0.06 \pm 0.01$	$0.59 \pm 0.05$	$25.25 \pm 0.49$
<b>B-DivRec</b>	$3.48 \pm 0.13$	$0.75 \pm 0.06$	$0.81 \pm 0.01$	$0.06 \pm 0.01$	$0.59 \pm 0.05$	$26.65 \pm 0.54$
<b>NO DPP</b>						
MMR	<b><math>3.78 \pm 0.12</math></b>	<b><math>0.84 \pm 0.04</math></b>	$0.74 \pm 0.03$	<b><math>0.07 \pm 0.01</math></b>	<b><math>0.77 \pm 0.04</math></b>	$26.93 \pm 0.56$

Table 16: Benchmark on PREDICT (9 diseases/users,  $B = 3$ ). REL: relevance. PREC: precision with  $\tau = 0.5$ .  $DIV^L$ : intrabatch/local diversity.  $DIV^G$ : interbatch/global diversity.  $DIV^+$ : effective diversity. TIME: runtime.

SAMPLING	REL $\uparrow$	PREC $\uparrow$	$DIV^L$ $\uparrow$	$DIV^G$ $\uparrow$	$DIV^+$ $\uparrow$	TIME $\downarrow$
QDDDecomp.	$0.51 \pm 0.01$	$0.52 \pm 0.02$	$0.68 \pm 0.01$	$0.44 \pm 0.02$	$0.77 \pm 0.03$	$1.02 \pm 0.01$
CondDPP	$0.53 \pm 0.01$	$0.57 \pm 0.01$	$0.69 \pm 0.01$	$0.45 \pm 0.02$	$0.86 \pm 0.02$	$1.25 \pm 0.01$
$\varepsilon$ -Greedy	$0.51 \pm 0.01$	$0.52 \pm 0.02$	$0.68 \pm 0.01$	$0.44 \pm 0.02$	$0.77 \pm 0.03$	$1.01 \pm 0.0$
<b>B-DivRec</b>	$0.52 \pm 0.01$	$0.54 \pm 0.02$	$0.77 \pm 0.01$	$0.45 \pm 0.02$	$0.87 \pm 0.03$	$1.29 \pm 0.01$
<b>MAXIMIZ.</b>						
QDDDecomp.	<b><math>0.76 \pm 0.02</math></b>	<b><math>0.93 \pm 0.01</math></b>	<b><math>1.02 \pm 0.02</math></b>	$0.57 \pm 0.02$	<b><math>1.02 \pm 0.02</math></b>	$0.87 \pm 0.01$
CondDPP	$0.75 \pm 0.02$	<b><math>0.93 \pm 0.01</math></b>	$0.87 \pm 0.01$	<b><math>0.63 \pm 0.02</math></b>	$0.88 \pm 0.01$	$1.07 \pm 0.01$
$\varepsilon$ -Greedy	<b><math>0.76 \pm 0.02</math></b>	<b><math>0.93 \pm 0.01</math></b>	<b><math>1.02 \pm 0.02</math></b>	$0.57 \pm 0.02$	<b><math>1.02 \pm 0.02</math></b>	$0.87 \pm 0.01$
<b>B-DivRec</b>	<b><math>0.76 \pm 0.02</math></b>	<b><math>0.93 \pm 0.01</math></b>	<b><math>1.02 \pm 0.02</math></b>	$0.57 \pm 0.02$	<b><math>1.02 \pm 0.02</math></b>	$1.15 \pm 0.01$
<b>NO DPP</b>						
MMR	$0.66 \pm 0.02$	$0.86 \pm 0.01$	$0.69 \pm 0.02$	$0.6 \pm 0.02$	$0.7 \pm 0.02$	$1.88 \pm 0.01$