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

Anonymous authors

000

001

002 003 004

010 011

012

013

014

016

017

018

019

021

025 026 027

028 029

031

033

034

037

040

041

042

043

044

046

047

048

051

052

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 serendipidity 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 "intrabatch" (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 h^t queries the recommender system at each time t, (2) the recommender system returns a fixed-size batch $S^t \subset \Omega$ of B items to present to the user, (3) user h^t outputs feedback values $m{y}^t := \{y_1^t, \dots, y_B^t\}$ 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 $h^t \in \mathbb{R}^m$ at each recommendation round t. Ideally, h^t captures all the information about the items the user has previously interacted with (e.g., 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 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 $q_{\Theta}: \mathbb{R}^d \times \mathbb{R}^m \to \mathbb{R}$ where $q_{\Theta}(\phi, h)$ is the expected feedback for user h on item ϕ which should be maximized. In the case of a known and noiseless feedback model, if ϕ is the k^{th} recommended item to user h^t at time t, then $y_k^t = q_{\Theta}(\phi, h^t)$.

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 rel) is simply defined by the expected click-through rate. We also consider the precision metric (denoted 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 *intrabatch/local diversity* (denoted div^L), focusing on the diversity inside a batch of B items; and the *interbatch/global diversity* (denoted 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 S of items is can be obtained by computing the volume of the parallelotope induced by the rows of the kernel matrix $K_{S,S}:=(k(\phi^i,\phi^j))_{i,j\in S}$ built from the item embeddings in S. The volume of a set $S:=\{i_1,\ldots,i_B\}$ of items is $vol(S):=(\det K_{S,S})^{1/2}$.

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

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

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

¹By definition of a kernel function, this definition of volume is well-defined for any $S \subseteq \Omega$.

The final summary metric is the average div⁺ across encountered users.

2 Related works

108

110

111112

113

114

115

116 117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

132

133

134

135

136

137 138

139

140

141

142

143

144

145

146

147

148

149

150

151 152

153

154

155

156

157

158

159

161

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 Ω . 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 α -DPP (Calandriello et al., 2020) where $\alpha < 1$. 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 $\mathcal{O}(B^2N)$ (Gillenwater et al., 2012; Chen et al., 2018) for maximizing over subsets of size B. Finally, conditioning over another subset \mathcal{H} of itemsthat is, sampling a subset which is diverse compared to previously selected set of points-can be described with the following distribution (Borodin & Rains, 2005) $\mathbb{P}(S = A \cup \mathcal{H} \mid S \supseteq \mathcal{H}) \propto$ $\det (k_{A,A} - k_{A,\mathcal{H}} k_{\mathcal{H},\mathcal{H}}^{-1} k_{A,\mathcal{H}}^{\mathsf{T}})$. 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.

To tackle the issue of recommending items with high quality/relevance and high "intrabatch"/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 $\{q_\Theta(\phi^i, h)\}_{i\in\Omega}$ for each item and a given user h, the QD decomposition is the distribution $\mathbb{P}(S = A) \propto \det(Q_A \Phi_A \Phi_A^{\mathsf{T}} Q_A)$, where Q_A is a diagonal matrix of size $B \times B$ which contains the quality scores for each item i_1, \ldots, i_B in A, and $\Phi_A := [\phi^{i_1}, \ldots, \phi^{i_B}]^{\mathsf{T}}$ 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 Ω , an item-item similarity function sim, a relevance function rel between items, and a quality-diversity trade-off parameter $\lambda \in [0,1]$, the goal is to return an item ϕ^t at round t which is both relevant to the query item h^{query} and diverse with respect to previously selected items i_1,i_2,\ldots,i_{t-1} . The recommendation procedure is defined recursively. At time t, MMR selects item $i_t := \arg\max_{i \in \Omega} \left(\lambda \operatorname{rel}(\phi^i, h^{\text{query}}) - (1-\lambda) \max_{j \in \{i_1,\ldots,i_{t-1}\}} \sin(\phi^i, \phi^j) \right)$ for each

²We abuse notation here, as it is clear that we refer to the complexity notation Ω and not the universe of items Ω .

 $t \in \{1,\ldots,B\}$, 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 $b \in \{1,\ldots,B\}$, at time t, MMR selects item $i^{t+b} := \arg\max_{i \in \Omega} \left(\lambda \ q_{\Theta}(\phi^i, \mathbf{h}^t) - (1-\lambda) \max_{j \in \mathcal{H}^t \cup \{i^t,\ldots,i^{t+b-1}\}} k(\phi^i, \phi^j)\right)$. 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 $k: \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ 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 $\nu : \mathbb{R}^d \to \mathbb{R}^{d'}$ such that, for any $\mathbf{x}, \mathbf{y} \in \Omega$, $k(\mathbf{x}, \mathbf{y}) = \nu(\mathbf{x})^\mathsf{T}\nu(\mathbf{y})$. ν 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_{Θ} *has values in* \mathbb{R}_{+}^{*} .

Assumption 3.3. Unit embeddings. $\|\phi\|_2 = \|h\|_2 = 1$ for all user embeddings $h \in \mathbb{R}^m$ 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 h^t , if ϕ is the k^{th} recommended item, then the observed feedback value y_k^t from user h^t for the k^{th} item satisfies $y_k^t = q_{\Theta}(\phi, h^t)$.

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 3 to diversity-focused recommendation 4 . The likelihood matrix for this family of DPPs, called QDT, for a given user h, their history \mathcal{H} , and item subset \mathcal{S} , is:

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

where $\lambda \in [0,1]$ controls the trade-off between quality $(\lambda=1)$ and diversity $(\lambda=0)$. The presence of λ allows us to implement continuously and explicity the trade-off between relevance and diversity. $Q_{h,S}$ is a diagonal matrix containing the expected rewards $\{q_{\Theta}(\phi^i,h)\}$ 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 S,Ω,\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 λ . For instance, traditional quality-diversity decomposition (Kulesza & Taskar, 2010) can be obtained by setting $f^{\text{QDDecomp}}(k, \mathcal{S}, \mathcal{H}) := k_{\mathcal{S},\mathcal{S}}$ with $\lambda = 0.5$, where k is the linear kernel $k(\{\psi\}, \{\phi\}) = \psi^{\mathsf{T}}\phi$ or, equivalently, $\nu: x \mapsto x$. A conditional DPP (Borodin & Rains, 2005) can be described using $\lambda = 0.5$ and $f^{\text{CondDPP}}(k, \mathcal{S}, \mathcal{H}) := k_{\mathcal{S},\mathcal{S}} - k_{\mathcal{S},\mathcal{H}}(k_{\mathcal{H},\mathcal{H}})^{-1}(k_{\mathcal{S},\mathcal{H}})^{\mathsf{T}}$. Finally, an ε -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 $L_f^\lambda(\mathcal{S}; \boldsymbol{h})$ as the score scr of any B-sized set \mathcal{S}

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

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

⁴*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 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: $Q_{h^t,\Omega} \leftarrow \text{diag}(\{q_{\Theta}(\phi, h^t)\}_{\phi \in \Omega})$ and initialize \mathcal{A} with λ^t

5: $L_{f\text{B-DivRec}}^{\lambda t}(\Omega; \boldsymbol{h}^t) \leftarrow (Q_{\boldsymbol{h}^t,\Omega})^{2\lambda^t} f(k,\Omega,\mathcal{H}^t)^{2(1-\lambda^t)} (Q_{\boldsymbol{h}^t,\Omega})^{2\lambda^t}$ # Likelihood matrix

6: Use SAMPLING Sample $S^t \subseteq \Omega$ proportionally to $\det L_f^{\lambda^t}(S^t; h^t)$

7: or MAXIMIZATION Solve $S^t \in \arg \max_{S \subseteq \Omega} \det L_f^{\lambda^t}(S; h^t)$

8: $\boldsymbol{y}^t \leftarrow (y_1^t, y_2^t, \dots, y_B^t)$, update \mathcal{A} with $(1 - 2\lambda^t) \nabla_{\lambda} \mathrm{scr}_{f, \boldsymbol{y}^t}^{\lambda_t^t}(\mathcal{S}^t; \boldsymbol{h}^t)$ and select λ^{t+1} from \mathcal{A} 9: **return** $\mathcal{S}^t, \lambda^{t+1}$

216

217

218 219 220

221

222

223

224

230

231

232

233

234

235

236 237

238

244

245

246

247

248

249

250

251 252

253

254

255

256

257

258

259

260

261

262

263

264

265

266

267

268

269

with the convention $\mathrm{scr}_{f,q_{\Theta}}^{\lambda}(\emptyset) := 0$, provided that the volume is well-defined—that is, $f(k,\Omega,\emptyset)$ is positive-definite for any universe of items Ω . From the expression in Equation 4, the larger the score, the better the subset 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 $\alpha \in [0, 2]$

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

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

and $\cos(x,y)$ is the cosine similarity between vectors x and y. α 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 α 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 ν 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 $N \gg d' \geq B +$ $\max_{t < T} |\mathcal{H}^t|$ 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 $q(S; \mathcal{H}, \alpha)$ for any user h of history \mathcal{H} and subset 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 λ) can also be in linear time in N. Using the α -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 λ -which takes into account the user's prior reactions to diversified items. We can frame this problem as an online learning game between a player-that is, the recommender system-and Nature-which is the interaction with *one* user (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 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 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 $\lambda, \mathcal{S}, \boldsymbol{y} \mapsto \operatorname{scr}_{f^{\operatorname{B-DivRec}}, \boldsymbol{y}}^{\lambda}(\mathcal{S}; \boldsymbol{h})$ for all $\lambda \in [0,1], \mathcal{S} \subseteq \Omega, |\mathcal{S}| = B$, and $\boldsymbol{y} \in (\mathbb{R}_+^*)^B$. The goal of the game, by selecting the λ_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 $(\mathcal{S}^t, \boldsymbol{y}^t)_{t \leq T}$:

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

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 λ .

5 EXPERIMENTAL STUDY

We consider several DPPs from the QDT family: QD decomposition, conditional DPP and the ε 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 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 $t, t+1, \ldots, t+|\mathcal{H}^t|$ with respective user context-history pairs $(h^t, \emptyset), (h^t, \{i_1\}), (h^t, \{i_1, i_2\}), \dots, (h^t, \mathcal{H}^t)$, 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 2nd decimal place.

For all experiments, we consider a linear kernel function $k:\phi,\psi\mapsto\phi^{\mathsf{T}}\psi$ 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

⁵Analoguously to a feedback model, y is such that $y(\phi, h)$ is the (observed) feedback from the user h if ϕ has been recommended, and otherwise is equal to 1 (meaning that it is ignored for a non-visited ϕ).

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 $q_{\Theta}: \phi, h \mapsto ((\phi)^{\mathsf{T}}h + 1)/2 \in [0,1]$. 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 $1, 2, \ldots, 5$. 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.

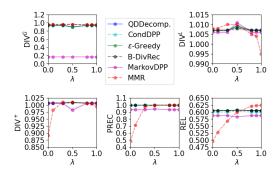


Table 1: Sensitivity analysis for α on SYN-THETIC750 (6 users, $B=3, \tau=0.5$), with the MAXIMIZATION strategy.

α	REL ↑	PREC ↑	DIV ^G ↑
0	0.61 ± 0.0	1.0 ± 0.0	0.95 ± 0.0
0.1	0.61 ± 0.0	1.0 ± 0.0	0.95 ± 0.0
0.5	0.61 ± 0.0	1.0 ± 0.0	0.95 ± 0.0
0.8	0.61 ± 0.0	1.0 ± 0.0	0.96 ± 0.0
0.90	0.6 ± 0.0	1.0 ± 0.0	0.97 ± 0.0
0.95	0.58 ± 0.0	1.0 ± 0.0	0.97 ± 0.0
0.98	0.57 ± 0.0	0.97 ± 0.0	0.97 ± 0.0
1	0.56 ± 0.0	0.92 ± 0.01	0.96 ± 0.0

Figure 1: Sensitivity analysis for λ on SYN-THETIC750 (6 users, $B=3,\,\tau=0.5$), with the MAXIMIZATION strategy.

Known feedback setting with fixed λ . 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 λ and α are shown respectively in Figure 1 and Table 1. As expected, across all recommender systems, as λ 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 α in B-DivRec, unsurprisingly, the higher α , the more (globally) diverse the recommendations, to the price of a loss in relevance and precision. Beyond a certain value of α (for this data set, $\alpha = 1$), the DPP can no longer sample enough diverse items, as α 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, rel = 0.51, and prec = 1.0 for all algorithms except for Markov DPP, where 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 div values and as reported in prior work (Kathuria et al., 2016). However, it results in a large loss in precision (prec) and relevance (rel), it is more time-consuming (see Appendix E), and does not quite achieve the best quality-diversity trade-off (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.

379

380 381 382

390

396 397 398

399

400

401

402

403

404

405

406

407

408

409

410

411

412

413

414

415

416 417

418

419

420

421

422

423 424 425

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

ALGO	PREC ↑	$\mathbf{DIV^G} \uparrow$	$\mathbf{DIV}^+ \uparrow$	REL ↑	$\mathrm{DIV}^{\mathrm{G}}\uparrow$	$\mathbf{DIV}^+\uparrow$
QDDec.	0.61 ± 0.01	0.95 ± 0.0	$\boldsymbol{1.01 \pm 0.0}$	0.61 ± 0.0	0.9 ± 0.01	$\boldsymbol{1.01 \pm 0.0}$
CondDPP	0.64 ± 0.01	0.94 ± 0.01	0.97 ± 0.01	0.6 ± 0.0	0.89 ± 0.01	$\textbf{1.01} \pm \textbf{0.0}$
ε -Greedy	0.61 ± 0.01	0.95 ± 0.0	$\textbf{1.01} \pm \textbf{0.0}$	$\overline{0.61\pm0.0}$	0.9 ± 0.01	$\textbf{1.01} \pm \textbf{0.0}$
Markov	0.59 ± 0.01	$\overline{0.17 \pm 0.0}$	0.91 ± 0.01	0.58 ± 0.0	0.17 ± 0.0	0.98 ± 0.01
MMR	$\textbf{1.0} \pm \textbf{0.0}$	0.96 ± 0.0	$\textbf{1.01} \pm \textbf{0.0}$	0.6 ± 0.0	0.96 ± 0.0	1.01 ± 0.0
B-DivRec	0.63 ± 0.01	0.95 ± 0.0	0.98 ± 0.01	$\overline{0.61\pm0.0}$	0.95 ± 0.0	$\textbf{1.01} \pm \textbf{0.0}$

Table 4: Benchmark on SYNTHETIC3M (4 users, Table 5: SYNTHETIC15M (2 users, B = 3) with MAXIMIZATION.

B = 5) with MAXIMIZATION.

ALGO	REL ↑	DIV ^G ↑	DIV ⁺ ↑	REL↑	DIV ^G ↑	DIV ⁺ ↑
CaradDDD	0.55 0.0	0.07 +0.0	1.01 +0.0	0.52 0.0	0.04 ± 0.0	1.01 + 0.0
CondDPP	0.55 ± 0.0	0.97 ± 0.0	1.01 ± 0.0	0.53 ± 0.0	0.94 ± 0.0	1.01 ± 0.0
MMR	0.73 ± 0.0	0.0 ± 0.0	0.02 ± 0.0	0.73 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
B-DivRec	0.55 ± 0.0	$\overline{0.97\pm0.0}$	$\overline{1.01\pm0.0}$	0.53 ± 0.0	0.94 ± 0.0	$\overline{1.01\pm0.0}$

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 SYN-THETIC750 (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^G , div^+), 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^G 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⁺ 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 λ 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 λ_{final} is compared to the best a posteriori relevance

Table 6: Benchmark on MovieLens (4 users, B=3) with Table 7: PREDICT (9 diseases/users, B=3) MAXIMIZATION. 3) with MAXIMIZATION.

ALGO	REL ↑	DIV ^G ↑	DIV ⁺ ↑	REL↑	DIV ^G ↑	DIV ⁺ ↑
CondDPP MMR B-DivRec	3.78 ± 0.12	$\begin{array}{c} 0.06 \pm 0.01 \\ \hline \textbf{0.07} \pm \textbf{0.01} \\ 0.06 \pm 0.01 \end{array}$	0.77 ± 0.04	0.66 ± 0.02		0.7 ± 0.02

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 $\lambda_0=0.5$. λ_{final} : final relevance weight after tuning. λ_{\star} : best *a posteriori* relevance weight.

DATA SET	$\lambda_{ ext{final}}$	λ_{\star}	$\mathbf{REL}\uparrow$	$\mathbf{DIV^G} \uparrow$	$\mathbf{DIV}^+\uparrow$	TIME ↓
SYNTHETIC750	0.18	0.10	0.61 ± 0.0	0.95 ± 0.0	1.01 ± 0.0	0.5 ± 0.0
(non adaptive)	_	_	0.61 ± 0.0	0.96 ± 0.0	1.01 ± 0.0	0.16 ± 0.0
SYNTHETIC30k	0.19	0.10	0.65 ± 0.0	0.96 ± 0.0	1.0 ± 0.0	29.8 ± 0.06
(non adaptive)	_	_	0.56 ± 0.0	0.96 ± 0.0	1.01 ± 0.0	1.34 ± 0.02
MovieLens	1.0	1.0	4.49 ± 0.01	0.01 ± 0.0	0.77 ± 0.01	158.1 ± 0.84
(non adaptive)	_	_	4.38 ± 0.0	0.01 ± 0.0	0.78 ± 0.0	29.7 ± 0.82
PREDICT	0.35	0.10	0.79 ± 0.01	0.49 ± 0.0	0.72 ± 0.01	1.97 ± 0.05
(non adaptive)	_	_	0.79 ± 0.01	0.49 ± 0.0	0.72 ± 0.01	0.82 ± 0.0

weight λ_{\star} . This value is computed at the end of a trajectory associated with one user by solving the maximization problem in λ 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 λ_{\star} 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 λ 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 div^G in the non-adaptive setting–the recommender system is leaning toward making popular recommendations, leading to a higher λ , higher relevance (rel +2.5%) to the price of some of effective diversity (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 semipersonalized 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 Derezinski, 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 Iaquinta, 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 Kaminskas 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 multiarmed bandits. In *Proceedings of the 25th international conference on Machine learning*, pp. 784–791, 2008.
 - Clémence Réda, Jill-Jênn 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-Jênn Vie, Camille Laıly, 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 Gaillard. 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.