
Gaussian Process Bandits for Top-k Recommendations

Anonymous Author(s)

Affiliation
Address
email

Abstract

1 Algorithms that utilize bandit feedback to optimize top-k recommendations are
2 vital for online marketplaces, search engines, and content platforms. However, the
3 combinatorial nature of this problem poses a significant challenge, as the possible
4 number of ordered top-k recommendations from n items grows exponentially with
5 k . As a result, previous work often relies on restrictive assumptions about the
6 reward or bandit feedback models, such as assuming that the feedback discloses
7 rewards for all recommended items rather than offering a single scalar feedback
8 for the entire set of top-k recommendations. We introduce a novel contextual
9 bandit algorithm for top-k recommendations, leveraging a Gaussian process with a
10 Kendall kernel to model the reward function. Our algorithm requires only scalar
11 feedback from the top-k recommendations and does not impose restrictive assump-
12 tions on the reward structure. Theoretical analysis confirms that the proposed
13 algorithm achieves sub-linear regret in relation to the number of rounds and arms.
14 Also, empirical results using a bandit simulator show that the proposed algorithm
15 surpasses other baselines across several scenarios.

16 **1 Introduction**

17 The top-k recommendation problem involves providing a ranked list of k items, such as news
18 articles or products, from a pool of n items [35, 13]. Online algorithms must adapt to dynamic user
19 preferences, making bandit algorithms suitable due to their use of limited feedback [1]. Developing
20 bandit algorithms is challenging due to limited feedback and the need for computational efficiency in
21 real-time recommendation environments. Recent research on user interfaces for recommendations
22 shows that the overall layout of the recommendation page is crucial for user appeal as designs
23 transition from simple dropdown lists to complex layouts [17, 13, 18]. As a result, bandit algorithms
24 must comprehensively select and display all top-k items jointly rather than merely selecting the most
25 relevant k items and displaying them in decreasing order of user relevance [32].

26 The joint consideration of top-k items makes the number of arms (possible actions for the bandit
27 algorithm) combinatorially large, i.e., $\Theta(n^k)$. Previous research on bandit algorithms often impose
28 strict assumptions about feedback models [31, 21], e.g., *semi-bandit* feedback gives a scalar value for
29 each of the top k items. Although semi-bandit feedback decomposes feedback from the combinatorial
30 number of arms to feedback for every recommended item, it is often unavailable [33]. Another
31 prevalent feedback assumption is *cascade* browsing [16], which posits that users examine items in a
32 pre-determined order and stop searching upon finding a desirable item, which provides item-specific
33 scalar feedback but does not fully capture possible non-linear interactions [27]. Figure 1 illustrates
34 the limitations of the cascade model in accurately representing user interactions within contemporary
35 top-k recommendation interfaces. These limitations motivate us to focus on a strictly more general
36 setting of the *full-bandit* feedback, where a single value for the entire top-k set is assumed [24].

37 Beyond feedback assumptions, the reward of bandit algorithms must be decomposable into scalars
38 over individual items to avoid a combinatorial explosion of arms, which is not always possible. For



Figure 1: A snapshot from Etsy showcases Father’s Day shopping recommendations. There is no obvious linear search order, which challenges the assumptions of the cascade model. Additionally, item proximity and arrangement are likely to influence clicks, suggesting a complex interaction pattern and advocating for full-bandit feedback without assumptions about user interaction with the recommended items.

Table 1: Compute and memory analysis for GP-TopK. Rows indicate different costs: overall *compute* and overall *memory* for T rounds, time for matrix-vector multiplication (**mvm**) with the kernel matrix K_{X_t} at time t , and time to update K_{X_t} . Columns indicate approaches: *kernel approach* uses full kernel matrices, while our novel *feature approach* performs the same operations via feature expansions and scales better with respect to T . The symbols c , k , and T represent the embedding size for contexts, the number of items, and the number of rounds, respectively.

Tasks	<i>kernel approach</i>	<i>feature approach</i>
<i>compute</i>	$\mathcal{O}(T^3)$	$\mathcal{O}(c \cdot k^2 \cdot T^2)$
<i>memory</i>	$\mathcal{O}(T^2)$	$\mathcal{O}(c \cdot k^2 \cdot T)$
mvm (K_{X_t})	$\mathcal{O}(t^2)$	$\mathcal{O}(c \cdot k^2 \cdot t)$
<i>compute</i> K_{X_t}	$\mathcal{O}((c + k^2) \cdot t)$	$\mathcal{O}(c \cdot k^2)$

39 instance, modern e-commerce platforms consider objectives such as diversity and fairness [1], which
40 cannot be measured by focusing solely on individual items [15]. This necessitates algorithms for
41 full-bandit feedback settings without assumptions about the objective or rewards [24].

42 This work develops a bandit algorithm that uses Gaussian processes (GPs) to model rewards under
43 full-bandit (i.e., one scalar value) feedback. GPs are selected for their flexibility in modeling feedback
44 for discrete, continuous, and mixed domains, such as continuous contexts and discrete rankings
45 [34]. Additionally, unlike parametric models that require optimization for accumulated feedback, GP
46 model updates are computationally inexpensive, involving only data updates [24]. While inference
47 for GPs may generally face computational limits, we will develop efficient inference methods tailored
48 to our proposed algorithm. Another challenge in developing GP-based bandit algorithms for top-k
49 recommendations is creating expressive positive-definite kernels that capture the similarity between
50 top-k recommendations [9].

51 GPs have been previously explored for bandit algorithms [28, 19]. Krause et al. [14] used GPs for
52 contextual bandits in continuous domains; we focus on the discrete domain of top-k recommendations.
53 Vanchinathan et al. [29] used GPs with a position-based feedback model, and Wang et al. [32] used
54 GPs with semi-bandit feedback for recommending top-k items. In contrast, our work does not focus
55 on a specific reward model or feedback assumption, and develops an efficient GP-based bandit
56 algorithm for top-k recommendations.

57 1.1 Contributions

58 Our primary contribution is the GP-TopK algorithm, a contextual bandit algorithm for recommending
59 top-k items. This algorithm operates in a full-bandit feedback setting without relying on assumptions
60 on reward, offering broader applicability than prior works. We leverage GPs with variants of the
61 *Kendall* kernel [12] to model the reward function and optimize the upper confidence bound (UCB) [28]
62 acquisition function to select the next arm. Further, we give a novel weighted convolutional Kendall
63 kernel for top-k recommendations that address pathologies in existing variants of the Kendall kernel
64 applied to top-k recommendations.

65 Our second key contribution is to improve the scalability of the GP-TopK algorithm for longer
66 time horizons. The initial computational demand for top-k ranking with GP-TopK is $\mathcal{O}(T^4)$ for T
67 rounds. We first reduce this to $\mathcal{O}(T^3)$ using iterative algorithms from numerical linear algebra [26].
68 Then, we derive sparse feature representations for the novel weighted convolutional Kendall kernel,
69 which, allows us to further improve the overall compute requirements from $\mathcal{O}(T^4)$ to $\mathcal{O}(T^2)$ and
70 memory requirements from $\mathcal{O}(T^2)$ to $\mathcal{O}(T)$. Table 1 summarizes these time and memory requirement
71 improvements, including their dependence on other parameters.

72 We also provide a theoretical analysis showing that GP-TopK’s regret is sub-linear in T and benefits
73 from the feature representations of the Kendall kernels we introduce. We show the regret’s upper

74 bound is almost quadratic in n , which improves significantly over the naive bound of $\Theta(n^k)$ for top-k
 75 recommendations without using specialized kernels [28]. Finally, we empirically validate GP-TopK's
 76 regret through simulations on real-world datasets and show improved regret compared to baselines.

77 **1.2 Organization**

78 The remainder of this paper is as follows: Section 2 presents Kendall kernels for full and top-k
 79 rankings, including the novel weighted convolutional Kendall kernel. Section 3 presents faster
 80 matrix-vector multiplication (MVM) algorithms for Kendall kernels, making the proposed bandit
 81 algorithm faster, as detailed later in Section 4, along with the regret analysis. Lastly, Sections 5 and 6
 82 present empirical results and discussion, respectively.

83 **2 Kendall Kernels for Full and Top-k Rankings**

84 This section overviews Kendall kernels and their extensions for top-k recommendations. First, we
 85 establish some notation. Let $[n] = \{1, 2, \dots, n\}$, and let π represent a top-k ranking, which is an
 86 ordered tuple of k distinct elements from $[n]$. We use σ to denote a full ranking ($k = n$) and let
 87 Π^k represent the set of all possible top-k rankings. Note that $|\Pi^k| = \Theta(n^k)$. The vector $\mathbf{p}^\sigma \in \mathbb{R}^n$
 88 corresponds to a full ranking σ with entry \mathbf{p}_i^σ giving the rank of item i . For top-k rankings, $\mathbf{p}^\pi \in \mathbb{R}^n$
 89 is constructed by arbitrarily assigning distinct ranks to items not in the top k . Indicator functions
 90 $\mathbf{p}_{i < j}^\sigma$ and $\mathbf{p}_{i > j}^\sigma$ indicate whether item i is ranked before or after item j , respectively in σ . Also, $\mathbf{p}_{i < j}^\pi$
 91 and $\mathbf{p}_{i > j}^\pi$ are indicator functions defined for top-k rankings.

92 **2.1 Kendall Kernels for Full Rankings**

93 Jiao et al. [9] showed that the Kendall tau rank correlation [12] is a positive definite (p.d.) kernel for
 94 full rankings, which we refer to as the standard Kendall (SK) kernel. The weighted Kendall (WK)
 95 kernel generalizes the SK kernel by differentially weighting item pairs [10]. Specifically, the SK and
 96 WK kernels for full rankings σ_1, σ_2 are defined as:

$$k^{sk}(\sigma_1, \sigma_2) := \frac{1}{\binom{n}{2}} \sum_{i < j} \eta_{i,j}(\sigma_1, \sigma_2) \quad (1)$$

$$k^{wk}(\sigma_1, \sigma_2) := \frac{1}{\binom{n}{2}} \sum_{i < j} w((\mathbf{p}_i^{\sigma_1}, \mathbf{p}_j^{\sigma_1}), (\mathbf{p}_i^{\sigma_2}, \mathbf{p}_j^{\sigma_2})) \cdot \eta_{i,j}(\sigma_1, \sigma_2), \quad (2)$$

97 where $\eta_{i,j}$ is 1 if the pair (i, j) is *concordant* (ordered the same in both rankings) and -1 otherwise; concretely, $\eta_{i,j}(\sigma_1, \sigma_2) := \mathbf{p}_{i < j}^{\sigma_1} \cdot \mathbf{p}_{i < j}^{\sigma_2} + \mathbf{p}_{i > j}^{\sigma_1} \cdot \mathbf{p}_{i > j}^{\sigma_2} - \mathbf{p}_{i < j}^{\sigma_1} \cdot \mathbf{p}_{i > j}^{\sigma_2} - \mathbf{p}_{i > j}^{\sigma_1} \cdot \mathbf{p}_{i < j}^{\sigma_2}$; and
 98 $ww((\mathbf{p}_i^{\sigma_1}, \mathbf{p}_j^{\sigma_1}), (\mathbf{p}_i^{\sigma_2}, \mathbf{p}_j^{\sigma_2}))$ is the value of a positive definite weighting kernel $w(\cdot, \cdot) : [n]^2 \times [n]^2 \mapsto$
 99 \mathbb{R} that operates on pairs of ranks. The $w_{i,j}$ allows flexibility and can assign varying importance
 100 to ranks, similar to the discounted cumulative gain (DCG) metric [7]. Note that both SK and WK
 101 kernels are p.d. and right-invariant with respect to Π^n [10]. In other words, they compute similarity
 102 based only on the relative ranks of pairs, not on the labels of items, as evident from Equations 1 and 2.
 103

104 **2.2 Kendall Kernels for Top-k Rankings**

105 **Weighted Kendall (WK) and Convolutional Kendall (CK) kernels.** To adapt the WK kernel
 106 from full rankings to top-k rankings, Jiao et al. [10] set the weighting function $w(i, j, \sigma_1, \sigma_2)$ to
 107 zero if either item is not in the top-k of either ranking. This scheme yields a p.d. kernel but does
 108 not consider items outside the intersection of top-k rankings. The convolutional operation offers an
 109 alternative for adapting the standard Kendall kernel to top-k rankings. Let B_π denote the set of full
 110 rankings consistent with the top-k ranking π (i.e., for every item i in π , $\forall \sigma \in B_\pi, \mathbf{p}_i^\pi = \mathbf{p}_i^\sigma$). The
 111 CK kernel is defined as:

$$k^{ck}(\pi_1, \pi_2) = \frac{1}{|B_{\pi_1}| \cdot |B_{\pi_2}|} \sum_{\sigma_1 \in B_{\pi_1}, \sigma_2 \in B_{\pi_2}} k^{sk}(\sigma_1, \sigma_2), \quad (3)$$

112 where k^{sk} is the standard Kendall kernel. The CK kernel is a p.d. kernel as it is a convolution
 113 of another p.d. kernel [5]. Unlike the WK kernel for top-k rankings, the CK kernel accounts for

114 items not in both top-k rankings. However, computing the CK kernel using Eq. (3) is expensive as it
 115 requires exponentially many evaluations of the kernel k^{sk} in the double summation. Therefore, Jiao
 116 et al. [9] developed an efficient algorithm to bypass this double summation and compute the kernel in
 117 $\mathcal{O}(k \log k)$ time.

118 **Proposed Weighted Convolutional Kendall (WCK) Kernel.** To combine the strengths of the WK
 119 and CK kernels for top-k rankings, we propose the weighted convolutional Kendall kernel for top-k
 120 rankings π_1 and $\pi_2 \in \Pi^k$:

$$k^{wck}(\pi_1, \pi_2) := \frac{1}{|B_{\pi_1}| \cdot |B_{\pi_2}|} \sum_{\sigma_1 \in B_{\pi_1}, \sigma_2 \in B_{\pi_2}} k^{wk}(\sigma_1, \sigma_2), \quad (4)$$

121 where k^{wk} represents the weighted Kendall kernel for full rankings $\sigma_1, \sigma_2 \in \Pi^n$. The proposed WCK
 122 kernel combines the flexibility of weighting different ranks (among the top-k items) differently, like
 123 the WK kernel, with the ability to account for items outside the intersection of both top-k rankings,
 124 like the CK kernel. Also, since it's a convolution of a p.d. kernel it is also a p.d.. However, its
 125 computation is again challenging as the RHS of Equation 4 evaluates k^{wk} exponentially many times.
 126 To simplify, we focus on a specific form of rank weights for k^{wk} , which we call *product-symmetric*
 127 rank weights:

$$w_{ps}((i_1, j_1), (i_2, j_2)) := w_s(i_1, j_1) \cdot w_s(i_2, j_2), \quad (5)$$

128 where, $w_s(i, j) : [n] \times [n] \mapsto \mathbb{R}$ is a symmetric function, i.e., $w_s(i, j) = w_s(j, i)$. Notably, the WCK
 129 kernel can be computed efficiently for the case of w_{ps} weights (Claim 1 below).

130 The WCK kernel, even with the relatively
 131 simple w_{ps} weights, has notable properties
 132 as shown in Table 2. In this table we use
 133 $w_s(i, j) = \frac{1}{\log(i+1)} \cdot \frac{1}{\log(j+1)}$ inspired by the
 134 DCG metric used in recommendation sys-
 135 tems [7]. Note that the WK kernel ranks two
 136 rankings with no overlap (π_0 and π_1) as more
 137 similar than two rankings with the same items
 138 but reverses ordering (π_0 and π_2), a clear
 139 pathology. On the other hand, the CK kernel
 140 fails to differentiate between reversed pairs at
 141 different ranks ($k^{ck}(\pi_0, \pi_3) = k^{ck}(\pi_0, \pi_4)$),
 142 another clear limitation. The WCK kernel
 143 with product-symmetric ranking weights can
 144 address these shortcomings and provide a
 145 more nuanced similarity comparison for top-k rankings.

Table 2: Comparison of Kendall kernel similarities for top-k rankings. The table shows kernel values $k(\pi_0, \cdot)$ for the top-k ranking $\pi_0 = [1, 2, 3]$ with other rankings $(\pi_1, \pi_2, \pi_3, \pi_4)$ for $n = 7$ and $k = 3$. Rankings are ordered left to right by increasing similarity to π_0 . WCK kernel values, based on DCG rank weights, increase from left to right. All kernels are unit normalized.

Top-k Kernels	π_1 [4, 5, 6]	π_2 [3, 2, 1]	π_3 [2, 1, 3]	π_4 [1, 3, 2]
WK	0.00	-1.00	0.33	0.33
CK	-0.60	0.60	0.87	0.87
WCK	-0.38	0.09	0.46	0.87

Claim 1. *The weighted convolutional Kendall kernel (Equation 4) with product-symmetric rank weights (Equation 5) can be computed in $\mathcal{O}(k^2)$ time.*

146 Appendix A provides the proof, which exploits the structure of product-symmetric rank weights w_{ps}
 147 to establish the existence of a feature representation for the WCK kernel, as given in Claim 3 below.
 148 We then show that the inner product of these features, and thus the WCK kernel, can be computed
 149 in $\mathcal{O}(k^2)$ time (Algorithm 2 in the appendix). Similarly to the result of Jiao et al. [9], we can avoid
 150 exponentially many evaluations of k^{wk} on the RHS of Equation (4) by computing the WCK kernel
 151 directly by another means.

152 3 Fast Matrix-Vector Multiplication with Kendall Kernel Matrices

153 In GPs, inference can be accelerated by using iterative algorithms that take advantage of fast matrix-
 154 vector-multiplications (MVMs) with the kernel matrix [3]. This section focuses on fast algorithms
 155 for kernel MVMs that exploit the structure of Kendall kernel matrices. Specifically, let $\text{mvm}(K_{X_t})$
 156 denote the running time to multiply the $t \times t$ kernel matrix $K_{X_t} = (k(x_i, x_j))_{x_i, x_j \in X_t}$ by a vector.
 157 Naively, $\text{mvm}(K_{X_t}) = \mathcal{O}(t^2)$. However, if $k(x_i, x_j) = \phi^a(x_i)^T \phi^b(x_j)$ for vectors $\phi^a(x_i)$ and
 158 $\phi^b(x_j)$ with only z non-zero entries, then $\text{mvm}(K_{X_t})$ reduces to $\mathcal{O}(z \cdot t)$, which is much faster than

159 $\mathcal{O}(t^2)$ when $z \ll t$. When $\phi^a = \phi^b$, we call ϕ^a the *linear feature vector* for the kernel k . Before
160 focusing on top-k ranking kernels, we provide a linear feature vector for the WK kernel on full
161 rankings as defined in Equation 2.

Claim 2. Let $\phi^{wk}(\sigma) : \Pi^n \mapsto \mathbb{R}^{\binom{n}{2}}$ be a vector indexed by unique item pairs (i, j) , defined as:

$$\phi_{i,j}^{wk}(\sigma) := \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_i^\sigma, \mathbf{p}_j^\sigma) \cdot (\mathbf{p}_{i < j}^\sigma - \mathbf{p}_{i > j}^\sigma),$$

where w_s is the symmetric weighting function in product-symmetric weights. Then, ϕ^{wk} is a linear feature vector for the weighted Kendall kernel with product-symmetric weights w_{ps} .

162 Using Claim 2, the linear feature vector for the WK kernel can be extended to the WK top-k ranking
163 kernel utilizing the structure of the product-symmetric weights. Furthermore, the feature vector
164 $\phi^{wk}(\pi)$ contains only $\mathcal{O}(k^2)$ non-zero entries due to the WK kernel's focus on item pairs within the
165 top-k, resulting in $\text{mvm}(K_{X_t}) = \mathcal{O}(k^2 \cdot t)$ for the WK kernel matrix.

166 Next, we focus on fast MVMs with
167 the WCK kernel, which includes
168 the CK kernel as a particular case.
169 We observe that any convolutional
170 kernel inherits linear features from
171 its constituent kernel. Specifically,
172 $\sum_{\sigma \in B_\pi} \phi_{i,j}^{wk}(\sigma)$ forms a feature vector
173 for the WCK kernel, which follows
174 from Equation 4 and Claim 2.

175 However, computing this feature vector
176 explicitly requires the exponential
177 summation over all $\sigma \in B_\pi$. Claim 3
178 shows that the summation can be com-
179 puted analytically and provides ex-
180 plicit linear feature vectors for the
181 WCK and CK kernels. It also shows
182 that ϕ^{wck} has only $\mathcal{O}(k^2 + 2nk)$ non-
183 zeros among its $\mathcal{O}(n^2)$ entries. Con-
184 sequently, $\text{mvm}(K_{X_t})$ for the WCK
185 kernel requires $\mathcal{O}((k^2 + 2nk) \cdot t)$ operations, which improves from $\mathcal{O}(t^2)$ to linear in t . However,
186 this introduces a dependence on n and is beneficial only for $n \leq t$. We next leverage redundancy in
187 ϕ^{wck} to eliminate this dependence, leading to the following main theorem about the $\text{mvm}(K_{X_t})$.

Theorem 1. For the WCK kernel with product-symmetric weights w_{ps} , the computational complexity of multiplying the kernel matrix K_{X_t} with any admissible vector is $\mathcal{O}(k^2 t)$, i.e., $\text{mvm}(K_{X_t}) = \mathcal{O}(k^2 t)$, where X_t is any arbitrary set of t top-k rankings.

188 Appendix A provides the proof in two steps. First, we leverage the values of ϕ^{wck} from Claim 3 and
189 categorize $\phi^{wck}(\pi_1)^T \phi^{wck}(\pi_2)$ based on item pairs, as summarized in Table 4. Next, we show that
190 only five combinations yield non-zero values, i.e., $\phi^{wck}(\pi_1)^T \phi^{wck}(\pi_2) = \sum_{i=1}^5 s_i(\pi_1, \pi_2)$. Each
191 term $s_i(\pi_1, \pi_2)$ is a dot product of vectors $\phi^{a_i}(\pi_1)^T \phi^{b_i}(\pi_2)$, which contains at most $\mathcal{O}(k^2)$ non-zero
192 entries. Thus, for WCK and CK kernels, $\text{mvm}(K_{X_t}) = \mathcal{O}(k^2 t)$ as these vectors for all five terms
193 have only $\mathcal{O}(k^2)$ non-zero entries. Consequently, Theorem 1 demonstrates that using these vector
194 representations for top-k rankings yields faster MVMs, i.e., $\text{mvm}(K_{X_t}) = \mathcal{O}(k^2 t) \ll \mathcal{O}(t^2)$.

195 4 Proposed GP-TopK Bandit Algorithm

196 This section outlines the top-k recommendation problem and introduces a generic contextual bandit
197 algorithm for top-k recommendations. We then explain how the components of this algorithm

198 are instantiated using our GP approach. An analysis of the proposed algorithm's computational
 199 complexity and regret follows this.

200 Let T be the number of rounds. Contexts \mathcal{C} are in a finite c -dimensional space, $\mathcal{C} \subseteq \mathbb{R}^c$. In the t^{th}
 201 round, we receive a context $\mathbf{c}_t \in \mathcal{C}$ and select a top- k ranking $\pi_t \in \Pi^k$. We then obtain a noisy reward
 202 $y_t = \hat{f}(\mathbf{c}_t, \pi_t) + \epsilon_t$, where \hat{f} is the true reward function and ϵ_t is round-independent noise. The regret
 203 is $r_t := \max_{\pi' \in \Pi^k} \hat{f}(\mathbf{c}_t, \pi') - \hat{f}(\mathbf{c}_t, \pi_t)$, with cumulative regret $R_T := \sum_{t=1}^T r_t$. The accumulated
 204 data at the t^{th} round is $\mathcal{D}_t = (\mathbf{c}_i, \pi_i, y_i)_{i=1}^t$. Algorithm 1 presents the bandit algorithm's generic
 205 schematic, aiming to minimize cumulative regret while ensuring computational efficiency.

Algorithm 1 Contextual Bandit Algorithm for Top- k Recommendations

Input: Total rounds T , initial reward model \mathcal{M}_0 , and acquisition function \mathcal{AF} .

```

1: for  $t = 1, \dots, T$  do
2:   Observe a context  $\mathbf{c}_t$  from the context space  $\mathcal{C}$ .
3:   Select a top- $k$  ranking  $\pi_t$  that maximizes  $\mathcal{AF}(\mathcal{M}_{t-1}(\mathbf{c}_t, \pi))$  for the context  $\mathbf{c}_t$ .
4:   Obtain the scalar reward  $y_t$ .
5:   Update the reward model  $\mathcal{M}_t$  using the accumulated feedback  $\mathcal{D}_t$ .
6: end for
```

206 The above algorithm requires two components: a) a reward model \mathcal{M}_t that can estimate the reward
 207 value given any context and top- k ranking utilizing the accumulated feedback \mathcal{D}_t and b) an acquisition
 208 function \mathcal{AF} for selecting π_t given the current reward model \mathcal{M}_t and observed context \mathbf{c}_t .

209 **Reward model \mathcal{M} and acquisition function \mathcal{AF} .** Our proposed GP-TopK bandit algorithm uses
 210 GP regression to model rewards for contexts and top- k rankings. GP regression is briefly covered
 211 in Section B.1 for completeness. In essence, the reward model \mathcal{M} maintains a distribution over
 212 functions f , i.e., $f \sim \mathcal{N}(0, k(\cdot, \cdot))$, where k is a product kernel function over both contexts and top- k
 213 rankings ($\mathcal{C} \otimes \Pi^k$). Specifically, the kernel function k is defined as follows:

$$k((\mathbf{c}_1, \pi_1), (\mathbf{c}_2, \pi_2)) := k^c(\mathbf{c}_1, \mathbf{c}_2) \cdot k^r(\pi_1, \pi_2), \quad (6)$$

214 where $k^c(\mathbf{c}_1, \mathbf{c}_2) = \mathbf{c}_1^T \mathbf{c}_2$ is the dot-product kernel and k^r is a kernel for top- k rankings. We use
 215 variants of the Kendall kernel for k^r from Section 2. Updating the reward model \mathcal{M}_t at the t^{th}
 216 round involves adding new data points to our GP regression, which is computationally inexpensive
 217 compared to the fine-tuning steps required by parametric models to incorporate the latest feedback.
 218 We utilize the UCB function for the acquisition function as it effectively balances exploration and
 219 exploitation by selecting actions that maximize the upper confidence bound on the estimated reward
 220 [28]. The UCB acquisition function is $\mathcal{AF}(\mathcal{M}_t(\mathbf{c}_t, \pi)) := \mu_{f|\mathcal{D}}((\mathbf{c}_t, \pi)) + \beta^{\frac{1}{2}} \cdot \sigma_{f|\mathcal{D}}((\mathbf{c}_t, \pi))$,
 221 where $\sigma_{f|\mathcal{D}}((\mathbf{c}_t, \pi)) = \sqrt{k_{f|\mathcal{D}}((\mathbf{c}_t, \pi), (\mathbf{c}_t, \pi))}$, and β controls the trade-off between exploration
 222 and exploitation. $\mu_{f|\mathcal{D}}$ and $k_{f|\mathcal{D}}$ are the GP posterior mean and covariance functions, as detailed
 223 in Section B.1. At the t^{th} round, the algorithm selects the top- k ranking $\pi \in \Pi^k$ that maximizes
 224 $\mathcal{AF}(\mathcal{M}_t(\mathbf{c}_t, \pi))$, which is performed using local search [19], as detailed in Appendix B.

225 **Computational complexity.** The GP-TopK bandit algorithm requires no compute for model updates.
 226 i.e., updating \mathcal{M}_t in the Line 5 of the Algorithm 1 requires only updating \mathcal{D}_t . The GP-TopK relies
 227 on local search for optimizing \mathcal{AF} , so the compute requirements arise only from \mathcal{AF} evaluations in
 228 the local search. As shown in Section B.1, computing the GP variance term for evaluating \mathcal{AF} , i.e.,
 229 $\sigma_{f|\mathcal{D}}((\mathbf{c}_t, \pi))$ involves solving $[K_{X_t} + \sigma^2 I]^{-1} \mathbf{v}$ for a vector \mathbf{v} , where $X_t = [(\mathbf{c}_1, \pi_1), \dots, (\mathbf{c}_t, \pi_t)]$.
 230 Naively, this requires $\mathcal{O}(t^3)$ time per round, leading to $\mathcal{O}(T^4)$ over T rounds. Iterative algorithms can
 231 expedite the process using our results on fast MVMs with kernel matrices, as discussed previously in
 232 Section 3 [25]. Theorem 2 formalizes the computational demands of the GP-TopK algorithm.

Theorem 2. Assuming a fixed number of iterations required by the iterative algorithms, the total computational time for running the GP-TopK bandit algorithm for T rounds of top- k recommendations, using the contextual product kernel (Equation 6), is $\mathcal{O}(k^2 c \ell T^2)$. This applies to WK, CK, and WCK top- k ranking kernels, where ℓ is the number of local search evaluations.

233 The proof of Theorem 2 is provided in Appendix B. It demonstrates efficiency gains from integrating
234 feature representation with iterative algorithms, reducing the computational time from $\mathcal{O}(T^3)$ to
235 $\mathcal{O}(T^2)$. This is a significant improvement, as even one MVM with K_{X_t} using the full kernel matrix
236 at each round requires $\mathcal{O}(T^3)$ time. Furthermore, the theorem also shows that the running time of the
237 GP-TopK algorithm does not explicitly depend on the number of items n .

238 **Regret analysis of the GP-TopK algorithm.** The cumulative regret of the proposed algorithm is
239 $R_T = \sum_{t=1}^T \max_{\pi' \in \Pi^k} \hat{f}(\mathbf{c}_t, \pi') - \hat{f}(\mathbf{c}_t, \pi_t)$, where π_t is the ranking chosen at round t . Optimizing
240 cumulative regret for top- k recommendations is challenging, requiring learning the context-arm
241 relationship and outperforming the best mapping. To bound cumulative regret, regularity assumptions
242 on the reward function \hat{f} are necessary [28, 14]. We consider the following two assumptions, either
243 of which suffices.

244 **Assumption 1.** \mathcal{X} is finite, meaning that only finite contexts are considered ($|\mathcal{C}| < \infty$), and the
245 reward function \hat{f} is sampled from the GP prior with a noise variance of ξ^2 .

246 **Assumption 2.** \mathcal{X} is arbitrary and the reward function \hat{f} has a bounded RKHS norm for the kernel k ,
247 i.e., $\|f\|_k \leq B$. The reward noises ϵ_t form an arbitrary martingale difference sequence (i.e., reward
248 noise does not systematically depend on its past values) and are uniformly bounded by ξ .

249 Under either Assumption 1 or 2, we prove the following regret bound for GP-TopK:

Theorem 3. If either Assumptions 1 or 2 hold, setting β_t as $2 \log \left(\frac{|\mathcal{C}| \cdot |\Pi^k| \cdot t^2 \cdot \pi^2}{6\delta} \right)$
and $300\gamma_t \ln^3 \left(\frac{t}{\delta} \right)$ respectively, the cumulative regret \mathcal{R}_T of the GP-TopK bandit algorithm for top- k recommendations can, with at least $1 - \delta$ probability, be bounded by $\tilde{\mathcal{O}}(n\sqrt{C_1 T c (\log |\mathcal{C}| + k + \log(T^2 \pi^2 / 6\delta))})$ under Assumption 1, and $\tilde{\mathcal{O}}(n\sqrt{C_1(2B^2 c + 300n^2 c^2 \ln^3(T/\delta))T})$ under Assumption 2. Here, $C_1 = \frac{8}{\log(1 + \xi^{-2})}$, and $\tilde{\mathcal{O}}$ excludes logarithmic factors related to n , k , and T .

250 Appendix B.4 provides the proof, leveraging the insight that $\log \det |I + \sigma^{-2} K_{X_T}|$ for any set X_T
251 can be effectively bounded using the finite-dimensional feature vectors introduced in this work.
252 Specifically, Proposition 2 utilizes the feature vectors from Section 2. Theorem 3 establishes that
253 the cumulative regret of the GP-TopK bandit algorithm grows sublinearly in T with high probability
254 under both assumptions. It also highlights the importance of using top- k ranking kernels, which
255 improve the asymptotic order $\tilde{\mathcal{O}}$ concerning n to $n^{k/2-1}$ and n^{k-1} under Assumptions 1 and 2,
256 respectively, as compared with Srinivas et al. [28]. This improvement is significant even for small
constant values of k , such as $k = 6$, as detailed below in Table 3.

Table 3: Comparison with Srinivas et al. (2010) for regret bounds of the bandit algorithm under both assumptions.

	Srinivas et al. (2010)	This work
Assumption 1	$\tilde{\mathcal{O}} \left(n^{\frac{k}{2}} \sqrt{C_1 T c (\log \mathcal{C} + k + \log(T^2 \pi^2 / 6\delta))} \right)$	$\tilde{\mathcal{O}} \left(n\sqrt{C_1 T c (\log \mathcal{C} + k + \log(T^2 \pi^2 / 6\delta))} \right)$
Assumption 2	$\tilde{\mathcal{O}} \left(n^{\frac{k}{2}} \sqrt{C_1 T c (2B^2 + 300n^2 c \ln^3(T/\delta))} \right)$	$\tilde{\mathcal{O}} \left(n\sqrt{C_1 T c (2B^2 + 300n^2 c \ln^3(T/\delta))} \right)$

257

258 5 Experiments

259 This section empirically evaluates the proposed GP-TopK bandit algorithms for the top- k recom-
260 mendations using a simulation based on the MovieLens dataset [4]. The reliance on simulation for

261 evaluating bandit algorithms is prevalent in the literature. It stems from the difficulty of conducting
 262 online evaluations in real-world bandit scenarios, mainly when there are combinatorial arms [29].
 263 Next, we provide details of the simulation setup and considered reward settings. Following that, we
 264 present results for the empirical regret for small and large numbers of arms below, respectively.

265 **Simulation setup and reward settings.** The bandit simulation setup follows the framework outlined
 266 by Jeunen et al. [8], utilizing real-world datasets on user-item interactions. Specifically, we train user
 267 and item embeddings using a collaborative filtering approach [6]. The user embeddings are accessed
 268 by the bandit algorithms as context embeddings, while the item embeddings remain hidden. In the
 269 non-contextual setup, the first user from the dataset is chosen as a fixed context throughout the bandit
 270 algorithm run, allowing us to use the same reward functions as the contextual bandit algorithm.

271 For setting up the reward functions, we utilize a similarity function $s(\mathbf{c}, \theta) := \sigma(a \cdot (\mathbf{c}^T \theta) - b)$ to
 272 measure similarity between any user and item embeddings, where a and b are similarity score and
 273 shift scalars, respectively. The sigmoid function σ maps similarity scores to a range between 0 and 1,
 274 enhancing the interpretability of the reward signal [32]. We set a and b to 6 and 0.3, respectively, to
 275 fully utilize the range of the similarity function, as assessed by evaluating its value for many arms.

276 We set up two preliminary reward functions based on the similarity function s . The first is the
 277 DCG metric, $\hat{f}_{\text{dcg}}(\mathbf{c}, \pi) = \sum_{i=1}^k \frac{1}{\log_2(i+1)} s(\mathbf{c}, \theta_{\pi_i})$, where \mathbf{c} and θ_{π_i} represent the context and item
 278 embeddings, respectively. The second is the diversity measure, $\hat{f}_{\text{div}}(\pi) = \frac{1}{k^2} \sum_{i=1}^k \sum_{j=1}^k \theta_{\pi_i}^T \theta_{\pi_j}$.
 279 These metrics quantify the relevance and diversity of top-k recommendations, respectively.

280 We use these functions in two contextual reward settings. The first setting focuses on normalized-
 281 DCG (nDCG), $\hat{f}_{\text{ndcg}}(\mathbf{c}, \pi) = \frac{\hat{f}_{\text{dcg}}(\mathbf{c}, \pi)}{\max_{\pi'} \hat{f}_{\text{dcg}}(\mathbf{c}, \pi')}$ [7]. The second setting combines \hat{f}_{ndcg} and \hat{f}_{div} as
 282 $\hat{f}_{\text{ndcgdiv}}(\mathbf{c}, \pi) = \lambda \cdot \hat{f}_{\text{ndcg}}(\mathbf{c}, \pi) + (1 - \lambda) \cdot \hat{f}_{\text{div}}(\pi)$, evaluating the aggregate effect of relevance and
 283 diversity. We set $\lambda = 0.25$ to emphasize relevance over diversity.

284 **Evaluation for small arm space.** This section presents empirical results for the cumulative regret
 285 of bandit algorithms with a limited number of arms. Specifically, with $n = 20$ and $k = 3$, there are
 286 6,840 top-k rankings, allowing for an exhaustive search to optimize the acquisition function. All
 287 bandit algorithms run in batch mode, updating every five rounds. We consider both reward settings
 288 for contextual and non-contextual scenarios, using a subset of five users for the contextual setting.
 289 Several baselines are set to assess the benefits of ranking (Kendall) kernels. Section C details the
 290 remaining hyper-parameter configurations and details of other baseline bandit algorithms.

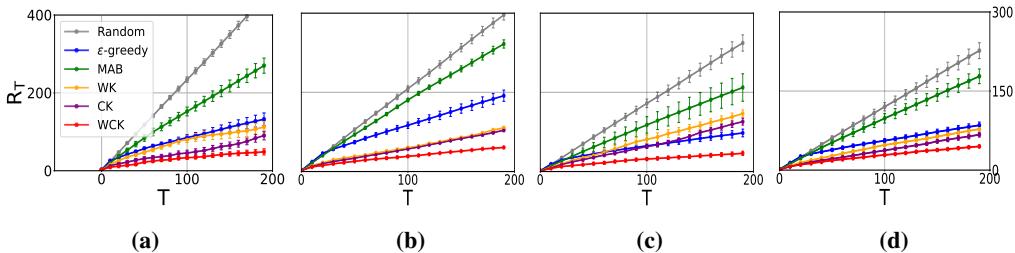


Figure 2: Comparative evaluation of bandit algorithms: The cumulative regret R_T over T rounds is shown.
 Lower values indicate better performance. Plots (a) and (b) represent non-contextual settings for nDCG (\hat{f}_{ndcg})
 and nDCG + diversity (\hat{f}_{ndcgdiv}) rewards, respectively. Plots (c) and (d) show results for contextual settings for
 five users using the same rewards. The y-axis for (a) and (b) is on the left, and for (c) and (d) on the right.
 The GP-TopK algorithm with Kendall kernels, especially the weighted convolutional Kendall (WCK) kernel,
 outperforms others. Details on other algorithms are in the text. Results are averaged over six trials.

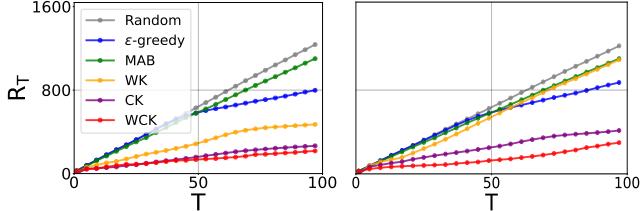
291 The *Random* algorithm randomly recommends any k items. The ϵ -*greedy* algorithm alternates
 292 between recommending a random top-k ranking with a probability of ϵ and choosing the top-k
 293 ranking with the highest observed mean reward. In contextual settings, ϵ -*greedy* differentiates arms
 294 for each unique context. Similarly, *MAB-UCB* conceptualizes each ranking as an independent arm,
 295 an equivalent of using a direct delta kernel approach for GPs along with UCB \mathcal{AF} . In contextual
 296 scenarios, *MAB-UCB* also treats arms distinctly per context. Each variant of the top-k ranking kernel
 297 yields one variation of the proposed GP-TopK algorithm, namely, WK, CK, and WCK. Figure 2

298 presents empirical values of the cumulative regrets for the above baseline and the proposed GP-TopK
 299 algorithms. In all cases, across both reward settings and in both contextual and non-contextual setups,
 300 the variants of the proposed GP-TopK algorithm outperform baselines that do not use Kendall kernels,
 301 highlighting the significance of top-k ranking kernels for full bandit feedback. Specifically, the CK
 302 and WCK kernels significantly outperform the WK kernel regarding the converged values of the
 303 regret, with the WCK kernel further improving on the CK kernel variant.

304 **Evaluation for large arm space.**

305 We evaluate bandit algorithms in a
 306 large arm space scenario with $n =$
 307 50 and $k = 3$ and $k = 5$, resulting
 308 in 1.1×10^5 and 1.1×10^{10} possible
 309 top-k rankings, respectively. Using
 310 local search, we focus on the nDCG
 311 reward. The remaining configuration
 312 is consistent with the small arm
 313 space setup. We use 10 restarts and
 314 5 steps in each search direction for
 315 the local search, starting with 1000
 316 initial candidates.

317 Figure 3 shows that the regret for
 318 the GP-TopK variants remains
 319 consistently lower even with a large arm
 320 space, despite the use of local search. The WCK approach significantly outperforms the CK, es-
 321 pecially for $k = 5$, as illustrated in the right plot of Figure 3. Additional empirical results on the
 322 effectiveness of local search in a large arm space and other rewards are given in Appendix C.



323 **Figure 3:** Comparative evaluation of bandit algorithms for large arm
 324 spaces, with $> 1.1 \times 10^5$ arms for the left plot and $> 1.1 \times 10^{10}$ arms for the right plot. Cumulative regret with respect to the rounds of
 325 the bandit algorithm is depicted. Results are averaged over six trials.
 326 In both settings, the WCK approach outperforms other baselines. For
 327 more details, see the textual description.

328 **6 Discussion**

329 This work develops a contextual bandit algorithm for top-k recommendations using Gaussian pro-
 330 cesses with Kendall kernels in a full-bandit feedback setting. We make no restrictive assumptions
 331 about feedback or reward models. Gaussian processes allow computationally free model updates for
 332 accumulated feedback data, although inference remains challenging. We address this by providing
 333 features for Kendall kernels for top-k rankings, resulting in a faster inference algorithm that reduces
 334 the complexity from $\mathcal{O}(T^4)$ to $\mathcal{O}(T^2)$. Additionally, we address issues with known variants and
 335 propose a more expressive Kendall kernel for top-k recommendations. Finally, we present theoretical
 336 and empirical results on cumulative regret to evaluate the proposed GP-TopK bandit algorithm.

337 **Future Directions and Limitations.** This work opens several research avenues. Efficient matrix-
 338 vector multiplication with Kendall kernel matrices can enable faster bandit algorithms with various
 339 acquisition functions, such as Thompson sampling and expected improvement. Exploring other
 340 kernels, like Mallow kernels, for top-k rankings and developing efficient algorithms for them is
 341 an intriguing direction, especially since the effectiveness of our algorithm depends on the function
 342 space induced by the RKHS of the underlying kernel. Assessing how well these kernels approximate
 343 various reward functions for top-k recommendations would provide valuable insights.

344 Exploring other bandit problem settings, such as stochastic item availability or delayed feedback,
 345 would enhance the applicability of this work to more complex scenarios. Extending the finite-
 346 dimensional GP framework to other acquisition functions using local search is another promising
 347 direction. One limitation of our regret analysis is that it does not account for approximations in the
 348 arm selection step due to local search [20]. This limitation is common in continuous domains, where
 349 optimizing acquisition functions often involves non-convex optimization [28].

350 **Impact.** This research advances bandit algorithms for top-k item recommendations. By improving
 351 recommendation efficiency and accuracy, our algorithms can enhance user experiences across plat-
 352 forms, promoting content relevancy and engagement. However, they may reinforce implicit biases in
 353 training data, limiting content diversity and entrenching prejudices. Therefore, monitoring over time
 354 is essential when deploying these algorithms in real-world environments.

350 **References**

351 [1] Róbert Busa-Fekete, Balázs Szörényi, Paul Weng, and Shie Mannor. Multi-objective bandits:
352 Optimizing the generalized gini index. In *Proceedings of the 34th International Conference on*
353 *Machine Learning (ICML)*, pages 625–634. PMLR, 2017.

354 [2] Aryan Deshwal, Syrine Belakaria, Janardhan Rao Doppa, and Dae Hyun Kim. Bayesian
355 optimization over permutation spaces. In *Proceedings of the 25th Conference on Artificial*
356 *Intelligence (AAAI)*, pages 6515–6523, 2022.

357 [3] Jacob R. Gardner, Geoff Pleiss, David Bindel, Kilian Q. Weinberger, and Andrew Gordon
358 Wilson. GPyTorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration.
359 In *Advances in the 31st Neural Information Processing Systems (NeurIPS)*, pages 7576–7586,
360 2018.

361 [4] F Maxwell Harper and Joseph A Konstan. The MovieLens datasets: History and context. In
362 *Transactions on Interactive Intelligent Systems*, volume 5, pages 1–19. ACM, 2015.

363 [5] David Haussler. Convolution kernels on discrete structures. Technical report, Technical report,
364 Department of Computer Science, University of California Santa Cruz, 1999.

365 [6] Yifan Hu, Yehuda Koren, and Chris Volinsky. Collaborative filtering for implicit feedback
366 datasets. In *8th IEEE International Conference on Data Mining (ICDM)*, pages 263–272. IEEE,
367 2008.

368 [7] Kalervo Jarvelin and Jaana Kekalainen. Cumulated gain-based evaluation of ir techniques. In
369 *ACM Transactions of Information Systems*, volume 20, pages 422–446, 2002.

370 [8] Olivier Jeunen and Bart Goethals. Top-k contextual bandits with equity of exposure. In
371 *Proceedings of the 15th ACM Conference on Recommender Systems (RecSys)*, pages 310–320,
372 2021.

373 [9] Yunlong Jiao and Jean-Philippe Vert. The Kendall and Mallows kernels for permutations.
374 In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pages
375 1935–1944. PMLR, 2015.

376 [10] Yunlong Jiao and Jean-Philippe Vert. The weighted Kendall and high-order kernels for permutations.
377 In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80, pages 2314–2322. PMLR, 2018.

378 [11] Michael N. Katehakis and Arthur F. Veinott. The multi-armed bandit problem: Decomposition
379 and computation. In *Mathematics of Operations Research*, volume 12, pages 262–268, 1987.

380 [12] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.

381 [13] Bart P. Knijnenburg, Martijn C. Willemsen, Zeno Gantner, Hakan Soncu, and Chris Newell.
382 Explaining the user experience of recommender systems. In *User Modeling and User-Adapted*
383 *Interaction*, volume 22, pages 441–504, 2011.

384 [14] Andreas Krause and Cheng Ong. Contextual gaussian process bandit optimization. *Advances in*
385 *the 24th Neural Information Processing Systems (NeurIPS)*, page 2447–2455, 2011.

386 [15] Matevz Kunaver and Tomaz Povzrl. Diversity in recommender systems - a survey. In
387 *Knowledge Based Systems*, volume 123, pages 154–162, 2017.

388 [16] Branislav Kveton, Csaba Szepesvari, Zheng Wen, and Azin Ashkan. Cascading bandits:
389 Learning to rank in the cascade model. In *Proceedings of the 32nd International Conference on*
390 *Machine Learning (ICML)*, pages 767–776. PMLR, 2015.

391 [17] E. Lex, Dominik Kowald, Paul Seitlinger, Thi Ngoc Trang Tran, Alexander Felfernig, and
392 Markus Schedl. Psychology-informed recommender systems. In *Foundations and Trends in*
393 *Information Retrieval*, volume 15, pages 134–242, 2021.

394 [18] Zeyang Liu, Yiqun Liu, Ke Zhou, Min Zhang, and Shaoping Ma. Influence of vertical result in
395 web search examination. In *Proceedings of the 38th International ACM SIGIR Conference on*
396 *Research and Development in Information Retrieval*, pages 193–202, 2015.

397 [19] Changyong Oh, Jakub Tomczak, Efstratios Gavves, and Max Welling. Combinatorial bayesian
398 optimization using the graph Cartesian product. *Advances in the 32nd Neural Information*
399 *Processing Systems (NeurIPS)*, 2019.

400

401 [20] My Phan, Yasin Abbasi Yadkori, and Justin Domke. Thompson sampling and approximate
 402 inference. *Advances in the 32nd Neural Information Processing Systems (NeurIPS)*, 32, 2019.

403 [21] Lijing Qin, Shouyuan Chen, and Xiaoyan Zhu. Contextual combinatorial bandit and its applica-
 404 tion on diversified online recommendation. In *In Proceedings of the 2014 SIAM International*
 405 *Conference on Data Mining*, pages 461–469, 2014.

406 [22] Carl Edward Rasmussen. *Evaluation of Gaussian processes and other methods for non-linear*
 407 *regression*. PhD thesis, University of Toronto Toronto, Canada, 1997.

408 [23] Carl Edward Rasmussen. Gaussian Processes in Machine Learning. In *Advanced Lectures on*
 409 *Machine Learning*, pages 63–71. Springer, 2004.

410 [24] Idan Rejwan and Yishay Mansour. Top- k combinatorial bandits with full-bandit feedback. In
 411 *Algorithmic Learning Theory*, pages 752–776. PMLR, 2020.

412 [25] Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.

413 [26] Yunus Saatçi. *Scalable inference for structured Gaussian process models*. PhD thesis, University
 414 of Cambridge, 2012.

415 [27] Dong-Her Shih, Kuan-Chu Lu, and Po-Yuan Shih. Exploring shopper’s browsing behavior and
 416 attention level with an eeg biosensor cap. In *Brain Sciences*, volume 9, page 301, 2019.

417 [28] Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process
 418 optimization in the bandit setting: No regret and experimental design. In *Proceedings of the*
 419 *27th International Conference on Machine Learning (ICML)*, pages 1015–1022, 2010.

420 [29] Hastagiri P Vanchinathan, Isidor Nikolic, Fabio De Bona, and Andreas Krause. Explore-
 421 exploit in top- n recommender systems via Gaussian processes. In *Proceedings of the 8th ACM*
 422 *Conference on Recommender Systems (RecSys)*, pages 225–232, 2014.

423 [30] Richard S. Varga. Geršgorin and his circles. 2004.

424 [31] Siwei Wang and Wei Chen. Thompson sampling for combinatorial semi-bandits. In *Proceedings*
 425 *of the 35th International Conference on Machine Learning (ICML)*, pages 5114–5122, 2018.

426 [32] Yingfei Wang, Hua Ouyang, Chu Wang, Jianhui Chen, Tsvetan Asamov, and Yi Chang. Efficient
 427 ordered combinatorial semi-bandits for whole-page recommendation. In *Proceedings of the*
 428 *20th Conference on Artificial Intelligence (AAAI)*, volume 31, page 2746–2753, 2017.

429 [33] Yue Wang, Dawei Yin, Luo Jie, Pengyuan Wang, Makoto Yamada, Yi Chang, and Qiaozhu
 430 Mei. Beyond ranking: Optimizing whole-page presentation. *Proceedings of the Ninth ACM*
 431 *International Conference on Web Search and Data Mining*, pages 103–112, 2016.

432 [34] Christopher K. I. Williams and Carl Edward Rasmussen. Gaussian Processes for Regression.
 433 *Advances in the 9th Neural Information Processing Systems (NeurIPS)*, pages 514–520, 1996.

434 [35] Dong Xin, Jiawei Han, and Kevin Chen-Chuan Chang. Progressive and selective merge:
 435 computing top- k with ad-hoc ranking functions. In *ACM SIGMOD Conference*, pages 103–114,
 436 2007.

437 **A Kendall Kernels for Full and Top-k Rankings – Omitted Details**

438 This section includes the proofs that were omitted from Section 2, presented in the following order:

439 • In Section A.1, we present proofs for Claims 2 and 3, which concern the feature representations
440 of Kendall kernels.

441 • In Section A.2, we provide Algorithms 2 and a proof of its correctness for computing the
442 WCK kernel in $\mathcal{O}(k^2)$ time, thereby proving Claim 1. Additionally, we extend this proof
443 to cover the proof of correctness for Algorithm 3, which can compute the CK kernel in
444 $\mathcal{O}(k \log k)$, initially introduced by Jiao et al. [9]. The original paper presented the algorithm
445 without a formal proof of correctness, a gap we address and fill in this section.

446 • Section A.3 details the proof for Theorem 1, discussing the matrix-vector multiplications
447 with the Kendall kernel matrix for top-k rankings. This proof builds on the Algorithm 2
448 given for computing the WCK kernel for top-k rankings.

449 **A.1 Feature Representation for Kendall Kernels for Top-k Rankings**

450 This section revisits the claims regarding the feature representations of the weighted Kendall kernel
451 and the weighted convolutional Kendall kernel, subsequently providing the proofs for these claims
452 mentioned earlier.

Claim 2. *Let $\phi^{wk}(\sigma) : \Pi^n \mapsto \mathbb{R}^{\binom{n}{2}}$ be a vector indexed by unique item pairs (i, j) , defined as:*

$$\phi_{i,j}^{wk}(\sigma) := \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_i^\sigma, \mathbf{p}_j^\sigma) \cdot (\mathbf{p}_{i < j}^\sigma - \mathbf{p}_{i > j}^\sigma),$$

where w_s is the symmetric weighting function in product-symmetric weights. Then, ϕ^{wk} is a linear feature vector for the weighted Kendall kernel with product-symmetric weights w_{ps} .

453 *Proof.* Following the definition of linear feature representation, we need to prove that $k^{wk}(\sigma_1, \sigma_2) =$
454 $\phi(\sigma_1)^T \phi(\sigma_2)$ for the product-symmetric weight kernel as given in Equation 5. Recalling from
455 Equation 2, we have $k^{wk}(\sigma_1, \sigma_2)$ as follows:

$$\begin{aligned} k^{wk}(\sigma_1, \sigma_2) &= \frac{1}{\binom{n}{2}} \cdot \sum_{i < j} w((\mathbf{p}_i^{\sigma_1}, \mathbf{p}_j^{\sigma_1}), (\mathbf{p}_i^{\sigma_2}, \mathbf{p}_j^{\sigma_2})) \cdot \eta_{i,j}(\sigma_1, \sigma_2), \\ &= \frac{1}{\binom{n}{2}} \cdot \sum_{i < j} w_s(\mathbf{p}_i^{\sigma_1}, \mathbf{p}_j^{\sigma_1}) \cdot w_s(\mathbf{p}_i^{\sigma_2}, \mathbf{p}_j^{\sigma_2}) \cdot \eta_{i,j}(\sigma_1, \sigma_2), \end{aligned} \quad (7)$$

456 where the second line incorporates the use of the product-symmetric weight kernel. Next, our focus
457 shifts to the simplification of $\eta_{i,j}(\sigma_1, \sigma_2)$, which is elaborated as follows:

$$\begin{aligned} \eta_{i,j}(\sigma_1, \sigma_2) &= \mathbf{p}_{i < j}^{\sigma_1} \cdot \mathbf{p}_{i < j}^{\sigma_2} + \mathbf{p}_{i > j}^{\sigma_1} \cdot \mathbf{p}_{i > j}^{\sigma_2} - \mathbf{p}_{i < j}^{\sigma_1} \cdot \mathbf{p}_{i > j}^{\sigma_2} - \mathbf{p}_{i > j}^{\sigma_1} \cdot \mathbf{p}_{i < j}^{\sigma_2}, \\ &= \mathbf{p}_{i < j}^{\sigma_1} \cdot (\mathbf{p}_{i < j}^{\sigma_2} - \mathbf{p}_{i > j}^{\sigma_2}) + \mathbf{p}_{i > j}^{\sigma_1} \cdot (\mathbf{p}_{i > j}^{\sigma_2} - \mathbf{p}_{i < j}^{\sigma_2}), \\ &= (\mathbf{p}_{i < j}^{\sigma_1} - \mathbf{p}_{i > j}^{\sigma_1}) \cdot (\mathbf{p}_{i < j}^{\sigma_2} - \mathbf{p}_{i > j}^{\sigma_2}). \end{aligned}$$

458 Combining the above factorization of $\eta_{i,j}$ with Equation 7, we get:

$$k^{wk}(\sigma_1, \sigma_2) = \frac{1}{\binom{n}{2}} \cdot \sum_{i < j} w_s(\mathbf{p}_i^{\sigma_1}, \mathbf{p}_j^{\sigma_1}) \cdot w_s(\mathbf{p}_i^{\sigma_2}, \mathbf{p}_j^{\sigma_2}) \cdot (\mathbf{p}_{i < j}^{\sigma_1} - \mathbf{p}_{i > j}^{\sigma_1}) \cdot (\mathbf{p}_{i < j}^{\sigma_2} - \mathbf{p}_{i > j}^{\sigma_2})$$

$$\begin{aligned}
&= \frac{1}{\binom{n}{2}} \cdot \sum_{i < j} \phi_{i,j}^{wk}(\sigma_1) \cdot \phi_{i,j}^{wk}(\sigma_2) \\
&= \phi(\sigma_1)^T \phi(\sigma_2).
\end{aligned}$$

459

□

Claim 3. Let $\phi^{wck}(\pi) : \Pi^k \mapsto \mathbb{R}^{\binom{n}{2}}$ be a vector indexed by unique item pairs (i, j) given as: $\phi_{i,j}^{wck}(\pi) := \frac{1}{\sqrt{\binom{n}{2}}} \cdot \mathbf{w}_{i,j}^{wck}(\pi) \cdot (\mathbf{p}_{i < j}^\pi - \mathbf{p}_{i > j}^\pi)$, where $\mathbf{w}_{i,j}^{wck}(\pi)$ is determined as follows:

$$\mathbf{w}_{i,j}^{wck}(\pi) = \begin{cases} w_s(\mathbf{p}_i^\pi, \mathbf{p}_j^\pi) & \text{if } \mathbf{p}_i^\pi \in [k] \& \mathbf{p}_j^\pi \in [k] \\ w_s(\mathbf{p}_i^\pi, \cdot) & \text{else if } \mathbf{p}_i^\pi \in [k] \& \mathbf{p}_j^\pi \notin [k], \\ w_s(\mathbf{p}_j^\pi, \cdot) & \text{else if } \mathbf{p}_i^\pi \notin [k] \& \mathbf{p}_j^\pi \in [k], \\ 0 & \text{otherwise,} \end{cases}$$

where w_s denotes symmetric weights and $w_s(\ell, \cdot) = \frac{1}{n-k} \sum_{j=k+1}^n w_s(\ell, j)$. Then, the vector ϕ^{wck} is a linear feature vector for the WCK kernel k^{wck} . By uniformly setting $w_s(\cdot, \cdot) \equiv 1$ in the definitions above, $\phi_{i,j}^{wck}(\pi)$ specializes to a linear feature vector for the CK kernel.

460 *Proof.* The main idea revolves around leveraging the feature representation of the Weighted Kendall
461 kernel for a full ranking and the linearity of the convolution operation. It is already established
462 that $k^{wck}(\sigma_1, \sigma_2) = \phi^{wk}(\sigma_1)^T \phi^{wk}(\sigma_2)$, as demonstrated in Claim 2. Recall that the WCK kernel
463 requires a double summation over pairs of rankings from B_{π_1} and B_{π_2} , which represent the sets of
464 full rankings consistent with their respective top-k rankings, as described in Equation 4. We simplify
465 the WCK kernel as follows:

$$\begin{aligned}
k^{wck}(\pi_1, \pi_2) &= \frac{1}{|B_{\pi_1}|} \cdot \frac{1}{|B_{\pi_2}|} \cdot \sum_{\sigma_1 \in B_{\pi_1}} \sum_{\sigma_2 \in B_{\pi_2}} \phi^{wk}(\sigma_1)^T \phi^{wk}(\sigma_2) \\
&= \left(\frac{1}{|B_{\pi_1}|} \cdot \sum_{\sigma_1 \in B_{\pi_1}} \phi^{wk}(\sigma_1)^T \right) \cdot \underbrace{\left(\frac{1}{|B_{\pi_2}|} \cdot \sum_{\sigma_2 \in B_{\pi_2}} \phi^{wk}(\sigma_2) \right)}_{:= \phi^{wck}(\pi_2)} \\
&= \phi^{wck}(\pi_1)^T \phi^{wck}(\pi_2).
\end{aligned}$$

466 The simplification above reveals that the feature representation, ϕ^{wck} , for the WCK kernel, is a $\binom{n}{2}$
467 dimensional vector and can be indexed by unique pairs of items (i, j) , much like the ϕ^{wk} . However,
468 the double summation is over an exponentially large number of pairs of rankings. Moving forward,
469 we shift our focus to the individual entries of this representation involving this summation, elucidating
470 the analytical values within the summation by exploring four unique cases, each dependent on whether
471 these specific items fall within the top-k rankings.

472 In Case 1, we examine the scenario where items i and j are within the top-k ranking π . Here, the
473 focus is on the feature representation of the pair, specifically when both elements are ranked among
474 the top-k positions.

475

476 Case 1: $\mathbf{p}_i^\pi \in [k]$ and $\mathbf{p}_j^\pi \in [k]$.

$$\begin{aligned}
\phi_{i,j}^{wck}(\pi) &= \frac{1}{|B_\pi|} \cdot \sum_{\sigma \in B_\pi} \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_i^\sigma, \mathbf{p}_j^\sigma) \cdot (\mathbf{p}_{i < j}^\sigma - \mathbf{p}_{i > j}^\sigma) \\
&= \frac{1}{|B_\pi|} \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_i^\pi, \mathbf{p}_j^\pi) \cdot \left(\sum_{\sigma \in B_\pi} \mathbf{p}_{i < j}^\sigma - \sum_{\sigma \in B_\pi} \mathbf{p}_{i > j}^\sigma \right)
\end{aligned}$$

$$\begin{aligned}
&= \frac{1}{|B_\pi|} \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_i^\pi, \mathbf{p}_j^\pi) \cdot (|B_\pi| \cdot \mathbf{p}_{i < j}^\pi - |B_\pi| \cdot \mathbf{p}_{i > j}^\pi) \\
&= \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_i^\pi, \mathbf{p}_j^\pi) \cdot (\mathbf{p}_{i < j}^\pi - \mathbf{p}_{i > j}^\pi). \tag{8}
\end{aligned}$$

477 The simplification in lines 3rd and 4th follows from the fact that any full ranking $\sigma \in B_\pi$, consistent
478 with the top-k ranking π , the relative ranks and weights of items i and j remains unchanged, given
479 $\mathbf{p}_i^\pi \in [k]$ and $\mathbf{p}_j^\pi \in [k]$. Concretely, this implies $\mathbf{p}_{i < j}^\sigma = \mathbf{p}_{i < j}^\pi$ for all $\sigma \in B_\pi$ and similar with the
480 other term.

481 In Case 2, we analyze when item i is in the top-k ranking while item j is not.

482

483 Case 2: $\mathbf{p}_i^\pi \in [k]$ and $\mathbf{p}_j^\pi \notin [k]$.

$$\begin{aligned}
\phi_{i,j}^{wck}(\pi) &= \frac{1}{|B_\pi|} \cdot \sum_{\sigma \in B_\pi} \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_i^\sigma, \mathbf{p}_j^\sigma) \cdot (\mathbf{p}_{i < j}^\sigma - \mathbf{p}_{i > j}^\sigma) \\
&= \frac{1}{|B_\pi|} \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot \sum_{\sigma \in B_\pi} w_s(\mathbf{p}_i^\sigma, \mathbf{p}_j^\sigma) \cdot (1 - 0) \quad (\text{since } \mathbf{p}_i^\pi \in [k] \text{ and } \mathbf{p}_j^\pi \notin [k]) \\
&= \frac{1}{|B_\pi|} \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot \sum_{\sigma \in B_\pi} w_s(\mathbf{p}_i^\pi, \mathbf{p}_j^\sigma).
\end{aligned}$$

484 Next, every possible consistent ranking is considered jointly while fixating on a specific rank outside
485 top-k elements, leading to $(n - k - 1)!$ different rankings. Given that $|B_\pi| = (n - k)!$, we can refine
486 the above expression as follows:

$$\begin{aligned}
\phi_{i,j}^{wck}(\pi) &= \frac{1}{|B_\pi|} \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot \sum_{l=k+1}^n w_s(\mathbf{p}_i^\pi, l) \cdot (n - k - 1)! \\
&= \frac{(n - k - 1)!}{|B_\pi|} \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot \sum_{l=k+1}^n w_s(\mathbf{p}_i^\pi, l) \\
&= \frac{1}{\sqrt{\binom{n}{2}}} \cdot \frac{1}{n - k} \cdot \sum_{l=k+1}^n w_s(\mathbf{p}_i^\pi, l) \\
&= \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_i^\pi, \cdot). \tag{9}
\end{aligned}$$

487 In Case 3, we analyze when item i is not in the top-k ranking while item j is.

488

489 Case 3: $\mathbf{p}_i^\pi \notin [k]$ and $\mathbf{p}_j^\pi \in [k]$. Similar to case 2, the simplification follows analogously, with the
490 only change being $\mathbf{1}_{\mathbf{p}_{i < j}^\sigma} - \mathbf{1}_{\mathbf{p}_{i > j}^\sigma} = -1$ instead of 1. Thus, by symmetry between i and j , we have
491 the following:

$$\phi_{i,j}^{wck}(\pi) = \frac{-1}{\sqrt{\binom{n}{2}}} \cdot w_s(\cdot, \mathbf{p}_j^\pi) = \frac{-1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_j^\pi, \cdot) \quad (\text{using symmetry of } w_s). \tag{10}$$

492 Lastly, in Case 4, we analyze when items i and j are not in the top-k ranking.

493 Case 4: $\mathbf{p}_i^\sigma \notin [k]$ and $\mathbf{p}_j^\sigma \notin [k]$.

$$\phi_{i,j}^{wck}(\pi) = \frac{1}{|B_\pi|} \cdot \sum_{\sigma \in B_\pi} \phi_{i,j}^{wck}(\sigma)$$

$$\begin{aligned}
&= \frac{1}{|B_\pi|} \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot \sum_{\sigma \in B_\pi} w_s(\mathbf{p}_i^\sigma, \mathbf{p}_j^\sigma) \cdot (\mathbf{p}_{i < j}^\sigma - \mathbf{p}_{i > j}^\sigma) \\
&= 0 \quad (\text{by symmetry}). \tag{11}
\end{aligned}$$

494 The result of zero arises from symmetry. Since \mathbf{p}_i^σ and \mathbf{p}_j^σ are not in the top- k ranking, they are
495 treated symmetrically in the summation overall rankings in B_π . For any ranking σ , suppose $\mathbf{p}_i^\sigma = l$
496 and $\mathbf{p}_j^\sigma = m$, there exists a corresponding ranking σ' such that only the items i and j are swapped.
497 Therefore, jointly, these two rankings yield $w_s(l, m)$ and $-w_s(l, m)$. Since w_s is symmetric, the
498 overall contribution from each pair of such rankings is zero. Hence, the entire summation nets to
499 zero.

500 Thus, with the explanation provided for each case and combining results from Equations 8, 9, 10 and
501 11, it's trivial to validate the Claim 3, i.e., $\phi_{i,j}^{wck}(\pi) = \frac{1}{\sqrt{\binom{n}{2}}} \cdot \mathbf{w}_{i,j}^{wck}(\pi) \cdot (\mathbf{p}_{i < j}^\pi - \mathbf{p}_{i > j}^\pi)$ for all unique
502 pair of items. From Case 4, we have $\mathcal{O}((n - k)^2)$ entries leaving at max only $\mathcal{O}(k^2 + 2nk)$ non-zero
503 entries. \square

504 A.2 Algorithms for Computing Kendall Kernels for top-k Rankings

505 In this section, we provide and delve into the proofs of Algorithms 2 and 3 for the weighted
506 convolutional Kendall kernel and the convolutional Kendall kernel, as previously discussed in Section
507 2. Section A.2.1 for valid both the correctness and computational complexity of Algorithm 2 as given
508 earlier in Claim 1. Following this, Section A.2.2 revisits Algorithm 3, initially introduced by Jiao
509 et al. [10]. The original publication presented the algorithm without formal proof of its correctness,
510 which we rectify and offer in Section A.2.2.

511 A.2.1 Efficiently Computing the Weighted Convolutional Kendall Kernel

512 This section provides a proof to Claim 1 to establish the efficiency and accuracy of Algorithm 2 in
513 computing the weighted convolutional Kendall kernel, as specified in Equation 4, with a focus on its
514 computational complexity.

Claim 1. *The weighted convolutional Kendall kernel (Equation 4) with product-symmetric rank
weights (Equation 5) can be computed in $\mathcal{O}(k^2)$ time.*

515 *Proof.* The claim is proven through Algorithm 2, where we establish its correctness and demonstrate
516 its computation requirement is $\mathcal{O}(k^2)$. The essence of our proof centers on analyzing the feature
517 representation of the WCK kernel, ϕ^{wck} , as outlined in Claim 3. The feature vectors of ϕ^{wck}
518 reside in a $\binom{n}{2}$ dimensional space, indexed by pairs of items. Our approach is to demonstrate
519 that Algorithm 2 accurately computes the right-hand side (RHS) of the equation $k^{wck}(\pi_1, \pi_2) =$
520 $\phi^{wck}(\pi_1)^T \phi^{wck}(\pi_2)$. This involves a summation over item pairs, expressed as $k^{wck}(\pi_1, \pi_2) =$
521 $\sum_{l < m} \phi_{l,m}^{wck}(\pi_1)^T \phi_{l,m}^{wck}(\pi_2)$.

522 Our proof analyzes various scenarios: cases where pairs of items, namely l and m , fall within the
523 top- k , scenarios with one item within the top- k and the other outside, and situations where neither
524 item is within the top- k . Each of these cases contributes distinctively to the computation of the
525 overall kernel, resulting in different terms in the algorithmic computation. This is encapsulated in
526 Algorithm 2, where $k^{wck}(\pi_1, \pi_2) = \sum_{i=1}^5 s_i(\pi_1, \pi_2)$, and each s_i corresponds to the terms given
527 earlier in Algorithm 2 from Section 2.

528 Before proceeding with the cases of this summation as given in Table 4, we recall the notations utilized
529 by Algorithm 3 in Definition 1. Also, remember that we will be proving for product-symmetric
530 weights as given in Equation 5, where, $w_s : [n] \times [n] \mapsto \mathbb{R}^n$ and its one-dimensional marginals are
531 $w_s(\ell, \cdot) = \frac{1}{n-k} \sum_{j=k+1}^n w_s(\ell, j)$. Table 4 shows how these cases are organized and relate to different
532 s_i terms required for computing the WCK kernel. The key strategy involves breaking down the
533 kernel's computation into cases based on the positioning of item pairs within the top- k rankings.
534 In case 1, we consider all the scenarios when both indices are within the set of items in both top- k
535 rankings, i.e., all items in the set $I_1 \cup I_2$.

Case	Description
1	Both items (l, m) in $I_1 \cup I_2$. Branches into the following three sub-cases based on the presence of items in $I_1 \cap I_2$: 1-a: Both items in $I_1 \cap I_2$. The concerned term is s_1 . 1-b: One item in $I_1 \cap I_2$. Subdivided into 1-b-i (other in $I_1 \setminus I_2$) and 1-b-ii (other in $I_2 \setminus I_1$); concerned terms are s_2 and s_3 . 1-c: No item in $I_1 \cap I_2$. Addresses cases where l and m are in different sets ($I_1 \setminus I_2$ and $I_2 \setminus I_1$); concerned term is s_4 .
2	One item in $I_1 \cup I_2$. I.e., either l is $I_1 \cup I_2$ or m is in $I_1 \cup I_2$, leading to sub-cases 2-a and 2-b; concerned term is s_5 .
3	No item in $I_1 \cup I_2$. Addresses the scenario where neither l nor m is in $I_1 \cup I_2$; value trivially zero.

Table 4: Case categorization for the proof of Algorithms 2 and 3 based on item pair ranks, where I_1 and I_2 are the sets of items for top-k rankings π_1 and π_2 , respectively.

Definition 1. Algorithm 2 and 3 and utilize following notations.

- I_1 and I_2 are the sets of items in rankings π_1 and π_2 , respectively.
- $\sigma_1 \in \Pi^{|I_1|}$ and $\tau_1 \in \Pi^{|I_1 \cap I_2|}$ are the full rankings of I_1 and $I_1 \cap I_2$, both consistent with the input top-k ranking π_1 . I.e., relative ranks of items is same yielding $\forall l, m \in I_1 \cap I_2, \mathfrak{p}_{i>j}^{\pi_1} = \mathfrak{p}_{i>j}^{\tau_1}$.
- Analogously, σ_2 and τ_2 are constructed utilizing the set I_2 and ranking π_2 .

Algorithm 2 Computing Weighted Convolutional Kendall Kernel

Input: Two permutations $\pi_1, \pi_2 \in \Pi^k$. Ranking weighting function $w_s : [n] \times [n] \mapsto \mathbb{R}^n$ and its one dimensional marginals are $w_s(\ell, \cdot) = \frac{1}{n-k} \sum_{j=k+1}^n w_s(\ell, j)$.

Output: Convolutional Weighted Kendall kernel $k^{wck}(\pi_1, \pi_2)$.

- Let I_1 and I_2 be the sets of items in rankings π_1 and π_2 , respectively.

- 1: **if** $|I_1 \cap I_2| \geq 2$ **then**
- 2: $s_1(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \sum_{1 \leq l < m \leq n} \sum_{l, m \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2}) \cdot \eta_{l, m}(\pi_1, \pi_2)$
- 3: **end if**
- 4: **if** $|I_1 \cap I_2| \geq 1$ and $|I_1 \setminus I_2| \geq 1$ **then**
- 5: $s_2(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot \sum_{l \in I_1 \cap I_2 \setminus m \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) (\mathfrak{p}_{l < m}^{\pi_1} - \mathfrak{p}_{l > m}^{\pi_1})$
- 6: **end if**
- 7: **if** $|I_1 \cap I_2| \geq 1$ and $|I_2 \setminus I_1| \geq 1$ **then**
- 8: $s_3(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot \sum_{l \in I_1 \cap I_2 \setminus m \in I_2 \setminus I_1} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2}) \cdot (\mathfrak{p}_{l < m}^{\pi_2} - \mathfrak{p}_{l > m}^{\pi_2})$
- 9: **end if**
- 10: **if** $|I_1 \setminus I_2| \geq 1$ and $|I_2 \setminus I_1| \geq 1$ **then**
- 11: $s_4(\pi_1, \pi_2) = -\frac{1}{\binom{n}{2}} \cdot \sum_{l \in I_1 \setminus I_2 \setminus m \in I_2 \setminus I_1} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_m^{\pi_2}, \cdot)$
- 12: **end if**
- 13: **if** $|I_1 \cap I_2| \geq 1$ and $|(n) \setminus (I_1 \cup I_2)| \geq 1$ **then**
- 14: $s_5(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot (n - |I_1 \cup I_2|) \cdot \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot)$
- 15: **end if**
- 16: $k^{wck}(\pi_1, \pi_2) = s_1(\pi_1, \pi_2) + s_2(\pi_1, \pi_2) + s_3(\pi_1, \pi_2) + s_4(\pi_1, \pi_2) + s_5(\pi_1, \pi_2)$

536 **Case 1:** The pair $(l, m) \in I_1 \cup I_2$ falls within the top-k, leading to three distinct cases. Below, we
537 provide s_i terms for each case as given in Table 4.

538 **Case 1-a:** Two items in $I_1 \cap I_2$, meaning both l and m belong to $I_1 \cap I_2$. Using Claim 3 regarding
539 the feature vector ϕ^{wck} , we simplify s_1 as follows:

$$\begin{aligned}
s_1(\pi_1, \pi_2) &= \sum_{1 \leq l < m \leq n | l, m \in I_1 \cap I_2} \phi_{l,m}^{wck}(\pi_1) \cdot \phi_{l,m}^{wck}(\pi_2) \\
&= \sum_{1 \leq l < m \leq n | l, m \in I_1 \cap I_2} \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \\
&\quad \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2}) \cdot (\mathbf{p}_{l < m}^{\pi_2} - \mathbf{p}_{l > m}^{\pi_2}) \\
&= \frac{1}{\binom{n}{2}} \sum_{1 \leq l < m \leq n | l, m \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2}) \cdot \eta_{l,m}(\pi_1, \pi_2). \quad (12)
\end{aligned}$$

540 **Case 1-b:** When one item is in $I_1 \cap I_2$, the other must reside either in $I_1 \setminus I_2$ or $I_2 \setminus I_1$, thus leading
541 to two distinct sub-cases. This is specified in Table 4. Concretely, if the other item is in $I_1 \setminus I_2$, it
542 contributes to the s_2 terms, whereas if it's in $I_2 \setminus I_1$, it contributes to the s_3 terms.

543 Corresponding to Case 1-b-i, when the other item is in $I_1 \cap I_2$, i.e., s_2 is the term corresponding to
544 indices where l is in $I_1 \cap I_2$ and m in $I_1 \setminus I_2$, or the reverse, represented by partial sums u and v . For
545 the partial sum u , with l in $I_1 \cap I_2$ and m in $I_1 \setminus I_2$, we find that $\mathbf{p}_l^{\pi_2}$ is in $[k]$, while $\mathbf{p}_m^{\pi_2}$ is not. The
546 simplification of u proceeds using Claim 3 as follows:

$$\begin{aligned}
u &= \sum_{1 \leq l < m \leq n | l \in I_1 \cap I_2 | m \in I_1 \setminus I_2} \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \\
&\quad \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) (\mathbf{p}_{l < m}^{\pi_2} - \mathbf{p}_{l > m}^{\pi_2}) \\
&= \frac{1}{\binom{n}{2}} \sum_{1 \leq l < m \leq n | l \in I_1 \cap I_2 | m \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}).
\end{aligned}$$

547 Similarly, the partial sum v can be simplified as follows:

$$\begin{aligned}
v &= \sum_{1 \leq l < m \leq n | m \in I_1 \cap I_2 | l \in I_1 \setminus I_2} \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \\
&\quad \cdot \frac{-1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) \cdot (\mathbf{p}_{l < m}^{\pi_2} - \mathbf{p}_{l > m}^{\pi_2}) \\
&= \frac{-1}{\binom{n}{2}} \sum_{1 \leq l < m \leq n | m \in I_1 \cap I_2 | l \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) \cdot (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \\
&= \frac{-1}{\binom{n}{2}} \sum_{1 \leq m < l \leq n | l \in I_1 \cap I_2 | m \in I_1 \setminus I_2} w_s(\mathbf{p}_m^{\pi_1}, \mathbf{p}_l^{\pi_1}) \cdot w_s(\mathbf{p}_m^{\pi_2}, \cdot) \cdot (\mathbf{p}_{m < l}^{\pi_1} - \mathbf{p}_{m > l}^{\pi_1}) \\
&= \frac{1}{\binom{n}{2}} \sum_{1 \leq m < l \leq n | l \in I_1 \cap I_2 | m \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}).
\end{aligned}$$

548 In the above, the first two lines use results from Claim 3 and use similarity of w_s . In the following
549 line, l and m are exchanged. Lastly, the negative sign is pushed into the indicator functions to make
550 the summand function of this partial sum v similar to the partial sum u , and the similarity of the w_s
551 is utilized. The above partial sums simplify s_2 as follows:

$$s_2(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot \sum_{l \in I_1 \cap I_2 | m \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}). \quad (13)$$

552 Analogously, in Case 1-b-ii, we deduce the corresponding term s_3 for the pair of indices as described
553 in Table 4 through symmetry. Specifically, the term s_3 can be outlined as follows:

$$s_3(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot \sum_{l \in I_1 \cap I_2 | m \in I_2 \setminus I_1} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2}) \cdot (\mathbf{p}_{l < m}^{\pi_2} - \mathbf{p}_{l > m}^{\pi_2}). \quad (14)$$

554 **Case 1-c:** Both items are outside $I_1 \cap I_2$, specifically, $l \in I_1 \setminus I_2$ and $m \in I_2 \setminus I_1$ or the reverse.
555 Like Case 1-b-i, we divide s_4 into partial summations u and v . Now, we calculate u under the
556 condition that $l \in I_1 \setminus I_2$ and $m \in I_2 \setminus I_1$.

$$\begin{aligned} u &= \sum_{1 \leq l < m \leq n | l \in I_1 \setminus I_2 | m \in I_2 \setminus I_1} \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \\ &\quad \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_m^{\pi_2}, \cdot) \cdot (\mathbf{p}_{l < m}^{\pi_2} - \mathbf{p}_{l > m}^{\pi_2}), \\ &= \frac{1}{\binom{n}{2}} \cdot \sum_{1 \leq l < m \leq n | l \in I_1 \setminus I_2 | m \in I_2 \setminus I_1} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot (1 - 0) \cdot w_s(\mathbf{p}_m^{\pi_2}, \cdot) \cdot (0 - 1), \\ &= \frac{-1}{\binom{n}{2}} \cdot \sum_{1 \leq l < m \leq n | l \in I_1 \setminus I_2 | m \in I_2 \setminus I_1} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_m^{\pi_2}, \cdot). \end{aligned}$$

557 Similarly, we can estimate partial sum v for the set $l \in I_2 \setminus I_1$ & $m \in I_1 \setminus I_2$. Using calculations
558 similar to Case-1-b-i for summing u and v , we have:

$$s_4(\pi_1, \pi_2) = \frac{-1}{\binom{n}{2}} \cdot \sum_{l \in I_1 \setminus I_2 | m \in I_2 \setminus I_1} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_m^{\pi_2}, \cdot). \quad (15)$$

559 **Case 2:** One item exists in $I_1 \cap I_2$, the other in $[n] \setminus (I_1 \cap I_2)$. It branches into two sub-cases: Case
560 2-a with one item in $I_1 \cup I_2$, and Case 2-b, where one item outside $I_1 \cap I_2$ but is in $I_1 \cup I_2$. Focusing
561 on Case 2-a, represented by s_5 , we simplify as follows. This involves two index scenarios, either
562 $l \in I_1 \cap I_2$ and $m \notin I_1 \cup I_2$ or vice versa, represented by partial sums u and v . We now simplify u
563 below:

$$\begin{aligned} u &= \frac{1}{\binom{n}{2}} \sum_{1 \leq l < m \leq n | l \in I_1 \cap I_2 | m \notin I_1 \cup I_2} \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \\ &\quad \cdot \frac{1}{\sqrt{\binom{n}{2}}} \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) \cdot (\mathbf{p}_{l < m}^{\pi_2} - \mathbf{p}_{l > m}^{\pi_2}), \\ &= \frac{1}{\binom{n}{2}} \sum_{1 \leq l < m \leq n | l \in I_1 \cap I_2 | m \notin I_1 \cup I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot), \\ &= \frac{1}{\binom{n}{2}} \sum_{1 \leq l < m \leq n | l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) \cdot (n - |I_1 \cup I_2|). \end{aligned}$$

564 Using steps similar to the previous case, we get the following value for s_5 :

$$s_5(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot (n - |I_1 \cup I_2|) \cdot \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot). \quad (16)$$

565 For Case 2-b, l or m are absent from I_1 or I_2 , leading to two sub-scenarios. Consequently, either
566 $\phi_{l,m}^{wck}(\pi_1)$ is zero or $\phi_{l,m}^{wck}(\pi_2)$ is zero. Therefore, these terms don't contribute to the overall WCK
567 kernel value.

568 **Case 3:** No item is in the top-k, i.e., both $l, m \notin I_1 \cup I_2$. As both items are absent from the top-k
569 in either ranking, the value trivially reduces to zero.

570 After covering all configurations of l and m , we incorporate results from Equations 12, 13, 14, 15,
571 and 16. This integration yields the expression $k^{wck}(\pi_1, \pi_2) = \sum_{i=1}^5 s_i(\pi_1, \pi_2)$, where, each term s_i
572 matches precisely with its corresponding expression in Algorithm 2. The proof for the correctness of
573 Algorithm 2 is complete, as each term s_i corresponds to its respective expression in the algorithm.
574 Regarding the time complexity of Algorithm 2, each term s_i sums at most k^2 quantities, and each
575 quantity summed can be computed in $\mathcal{O}(1)$ time. Therefore, the computation time required for
576 Algorithm 2 is $\mathcal{O}(k^2)$. \square

577 A.2.2 Efficiently Computing the Convolutional Kendall Kernel

578 This section provides Algorithm 3 for computing the convolutional Kendall kernel, as specified in
579 Equation 3. Later, its efficiency and accuracy are proved in Claim 4.

Algorithm 3 Computing Convolutional Kendall Kernel [10]

Input: Two top-k rankings $\pi_1, \pi_2 \in \Pi^k$.

Output: Convolutional Kendall kernel $k^{ck}(\pi_1, \pi_2)$.

- Let I_1 and I_2 be the sets of items in rankings π_1 and π_2 , respectively.
- Let $\sigma_1 \in \Pi^{|I_1|}$ and $\tau_1 \in \Pi^{|I_1 \cap I_2|}$ be the full rankings of I_1 and $I_1 \cap I_2$, both consistent
with the input top-k ranking π_1 .
- Analogously, construct σ_2 and τ_2 utilizing the set I_2 and ranking π_2 .

- 1: **if** $|I_1 \cap I_2| \geq 2$ **then**
- 2: $s_1(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot \binom{|I_1 \cap I_2|}{2} \cdot k^{sk}(\tau_1, \tau_2)$
- 3: **end if**
- 4: **if** $|I_1 \cap I_2| \geq 1$ and $|I_1 \setminus I_2| \geq 1$ **then**
- 5: $s_2(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot \sum_{l \in I_1 \cap I_2} 2 \cdot (\sigma_1(l) - \tau_1(l)) - k + |I_1 \cap I_2|$
- 6: **end if**
- 7: **if** $|I_1 \cap I_2| \geq 1$ and $|I_2 \setminus I_1| \geq 1$ **then**
- 8: $s_3(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot \sum_{l \in I_1 \cap I_2} 2 \cdot (\sigma_2(l) - \tau_2(l)) - k + |I_1 \cap I_2|$
- 9: **end if**
- 10: $s_4(\pi_1, \pi_2) = -\frac{1}{\binom{n}{2}} \cdot |I_1 \setminus I_2| \cdot |I_1 \setminus I_2|$
- 11: $s_5(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot |I_1 \cap I_2| \cdot |[n] \setminus (I_1 \cup I_2)|$
- 12: $k^{ck}(\pi_1, \pi_2) = s_1(\pi_1, \pi_2) + s_2(\pi_1, \pi_2) + s_3(\pi_1, \pi_2) + s_4(\pi_1, \pi_2) + s_5(\pi_1, \pi_2)$

Claim 4. Algorithm 3 computes the convolutional Kendall kernel (as given in the Equation 3)
with a computational complexity of $\mathcal{O}(k^2)$.

580 *Proof.* To establish the correctness of Algorithm 3, we will adopt the same proof approach as the one
581 used for Claim 1 concerning Algorithm 2. Specifically, we will adhere to the earlier categorization in
582 Table 4 and notations given in Definition 1. Since the CK kernel can be derived by uniformly setting
583 the weight function $w_s(i, j) = 1$, we will insert them in s_i terms as given in Algorithm 2. These
584 cases will be revisited and simplified by applying the condition $w_s(i, j) = 1$. Note that this also
585 implies its one-direction marginal weights to be 1, i.e., $w_s(i, \cdot) = 1$

586 **Simplifying the s_1 Term:** For the WCK kernel, Case 1-a leads to the expression of s_1 as stated in
587 Equation 12. In this case, when two items, specifically l and m , are both in the intersection $I_1 \cap I_2$,
588 it implies that $p_l^{\pi_1}, p_m^{\pi_1}, p_l^{\pi_2}$, and $p_m^{\pi_2}$ all rank within the top-k, denoted as $[k]$. We simplify the s_1
589 term for CK kernel as follows:

$$\begin{aligned}
s_1(\pi_1, \pi_2) &= \frac{1}{\binom{n}{2}} \sum_{1 \leq l < m \leq n | l, m \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2}) \cdot \eta_{l,m}(\pi_1, \pi_2) \\
&= \frac{1}{\binom{n}{2}} \sum_{1 \leq l < m \leq n | l, m \in I_1 \cap I_2} \eta_{l,m}(\pi_1, \pi_2) \\
&= \frac{1}{\binom{n}{2}} \sum_{1 < l' < m' \leq |I_1 \cap I_2|} \eta_{l',m'}(\pi_1, \pi_2) = \frac{\binom{|I_1 \cap I_2|}{2}}{\binom{n}{2}} k^{sk}(\pi_1, \pi_2). \tag{17}
\end{aligned}$$

590 The simplification process begins by assigning unit rank weights in the first line, i.e., $\mathbf{w}_i = 1$.
591 Following this, by relabeling the items in $I_1 \cap I_2$ and using τ_1 and τ_2 , which are the rankings of
592 π_1 and π_2 limited to the set $I_1 \cap I_2$ as defined in Definition 1, it is established that $\eta_{l',m'}(\tau_1, \tau_2) =$
593 $\eta_{l,m}(\pi_1, \pi_2)$. This is because the relative order of any pair of items is maintained in τ_1 and τ_2 .
594 Consequently, this leads to the final simplification to a scaled value of the standard Kendall kernel
595 k^{sk} , as given in Equation 1.

596 **Simplifying the s_2 and s_3 Terms:** The s_2 and s_3 terms are obtained for Case 1-b, which is for
597 case when one item is in $I_1 \cap I_2$ and the other item is either in $I_1 \setminus I_2$ or $I_2 \setminus I_1$. We divide this into
598 two sub-cases. Case 1-b-i: The other item is in $I_1 \setminus I_2$, with s_2 representing the summation terms
599 derived from the CK's inner product. Case 1-b-ii: The other item is $I_2 \setminus I_1$, where s_3 denotes the
600 summation terms. We simplify the s_2 term for the CK kernel as follows:

$$\begin{aligned}
s_2(\pi_1, \pi_2) &= \frac{1}{\binom{n}{2}} \sum_{l \in I_1 \cap I_2 | m \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_m^{\pi_2}, \cdot) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \\
&= \frac{1}{\binom{n}{2}} \sum_{l \in I_1 \cap I_2 | m \in I_1 \setminus I_2} (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \\
&= \frac{1}{\binom{n}{2}} \underbrace{\sum_{l \in I_1 \cap I_2 | m \in I_1 \setminus I_2} \mathbf{p}_{l < m}^{\pi_1}}_{:=u} - \frac{1}{\binom{n}{2}} \underbrace{\sum_{l \in I_1 \cap I_2 | m \in I_1 \setminus I_2} \mathbf{p}_{l > m}^{\pi_1}}_{:=v}.
\end{aligned}$$

601 Next, we examine the terms u and v in detail, starting with u . The term u , which corresponds to
602 $\mathbf{p}_{l < m}^{\pi_1}$, signifies instances where item l is ranked before item m in the top-k ranking π_1 . This can be
603 derived from the observation that $\sigma_1(l) - 1$ items are positioned before item l in the set I_1 . Out of
604 these items, $\tau_1(l) - 1$ also belong to the intersection $I_1 \cap I_2$. This follows from the definition of
605 the full rankings σ_1 and τ_1 on the set I_1 and the intersection $I_1 \cap I_2$, respectively. Consequently, it
606 can be concluded that $\sigma_1(l) - \tau_1(l)$ items from the set difference $I_1 \setminus I_2$ are ranked before item l .
607 The second term, v , corresponds to $\mathbf{p}_{l > m}^{\pi_1}$ and involves a calculation that takes into account the items
608 ranked after the l -th item in the set I . Specifically, there are $k - \sigma_1(l)$ items following the l -th item.
609 Within the intersection $I_1 \cap I_2$, the number of items before l is given by $|I_1 \cap I_2| - \tau_1(l)$. Therefore,
610 the expression $(k - \sigma_1(l)) - (|I_1 \cap I_2| - \tau_1(l))$ represents the count of elements that are positioned
611 after l in the set difference $I_1 \setminus I_2$.

612 Combining the above calculations for both terms u and v , the s_2 term for the CK kernel can be
613 simplified as follows:

$$s_2(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \sum_{l \in I_1 \cap I_2} 2 \cdot (\sigma_1(l) - \tau_1(l)) - k + |I_1 \cap I_2|. \tag{18}$$

614 Using the symmetry between Case 1-b-i and Case 1-b-ii, we can simplify s_3 for the CK kernel as
615 follows:

$$s_3(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \sum_{l \in I_1 \cap I_2} 2 \cdot (\sigma_2(l) - \tau_2(l)) - k + |I_1 \cap I_2|. \tag{19}$$

616 **Simplifying the s_4 and s_5 Terms:** We simplify the s_4 and s_5 terms for the CK kernel starting from
617 Equation 15 and Equation 16, respectively, as follows:

$$s_4(\pi_1, \pi_2) = \frac{-1}{\binom{n}{2}} \cdot \sum_{l \in I_1 \setminus I_2 \mid m \in I_2 \setminus I_1} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_m^{\pi_2}, \cdot) = \frac{-|I_1 \setminus I_2| \cdot |I_2 \setminus I_1|}{\binom{n}{2}} \quad (20)$$

$$s_5(\pi_1, \pi_2) = \frac{1}{\binom{n}{2}} \cdot (n - |I_1 \cup I_2|) \cdot \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) = \frac{|I_1 \cap I_2| \cdot |[n] \setminus (I_1 \cup I_2)|}{\binom{n}{2}}. \quad (21)$$

618 We have obtained the values of all the simplified s_i terms for the CK kernel in Equations 17, 18,
619 19, 20, and 21. By combining these terms, we get $k^{ck}(\pi_1, \pi_2) = \sum_{i=1}^5 s_i(\pi_1, \pi_2)$, where each term
620 s_i precisely matches its corresponding expression in Algorithm 3. This completes the proof of the
621 correctness of Algorithm 3. Regarding its time complexity, each term s_i sums at most k^2 quantities,
622 and each quantity can be computed in $\mathcal{O}(1)$ time. Therefore, the time required for Algorithm 3 to
623 compute the CK kernel is $\mathcal{O}(k^2)$. \square

624 A.3 Fast Matrix-Vector Multiplication with Kendall Kernel Matrix on Top-k Rankings

625 This section revisits Theorem 1 about the fact matrix-vector multiplication time for the Kendall kernel
626 matrix for top-k rankings. Specifically, we aim to eliminate the $\text{mvm}(K_X)$'s dependence on the
627 number of items, i.e., n on and linear dependence in the number of rounds, i.e., T , as claimed in
628 Theorem 1.

Theorem 1. *For the WCK kernel with product-symmetric weights w_{ps} , the computational complexity of multiplying the kernel matrix K_{X_t} with any admissible vector is $\mathcal{O}(k^2 t)$, i.e., $\text{mvm}(K_{X_t}) = \mathcal{O}(k^2 t)$, where X_t is any arbitrary set of t top-k rankings.*

629 *Proof.* The cornerstone of this proof lies in the computation of the WCK kernel, as delineated in
630 Algorithm 2. This algorithm requires only $\mathcal{O}(k^2)$ computation. For brevity, we write X to represent
631 X_T , and the proof follows for any X_t , i.e., any value of t , not just T .

632 As also suggested previously, we will demonstrate through the equation $K_X = (\Phi_X^a)^T \Phi_X^b$, where
633 both matrices Φ_X^a and Φ_X^b have columns with only $\mathcal{O}(k^2)$ non-zero entries. Consequently, this
634 leads to the computational complexity of matrix-vector multiplication, denoted as $\text{mvm}(K_X)$, being
635 $\mathcal{O}(k^2 \cdot T)$.

636 From Algorithm 2, we know that each entry of the kernel matrix $k(\pi_1, \pi_2)$, can be expressed as a sum
637 $\sum_{i=1}^5 s_i(\pi_1, \pi_2)$. Assuming each $s_i(\pi_1, \pi_2)$ equals $\phi^{a_i}(\pi_1)^T \phi^{b_i}(\pi_2)$, and considering that all vectors
638 ϕ^{a_i} and ϕ^{b_i} exhibit this property, we can express K_X as $(\Phi_X^a)^T \Phi_X^b$. Here, the i^{th} row of $(\Phi_X^a)^T$
639 and the j^{th} column of Φ_X^b are represented by $[\phi^{a_1}(\pi_i)^T, \dots, \phi^{a_5}(\pi_i)^T]$ and $[\phi^{b_1}(\pi_j), \dots, \phi^{b_5}(\pi_j)]$,
640 respectively. Therefore, the overall mvm complexity can be characterized by the sparsity of the
641 vectors ϕ^{a_i} and ϕ^{b_i} , as is formalized in the claim presented below.

Claim 5. *Consider a kernel matrix K_X corresponding to any set X of cardinality T . Each entry of K_X , denoted as $k(x_1, x_2)$, is defined by the sum $\sum_{i=1}^5 s_i(x_1, x_2)$, where each $s_i(x_1, x_2)$ is the result of the dot product $\phi^{a_i}(x_1)^T \phi^{b_i}(x_2)$, where, ϕ^{a_i} and ϕ^{b_i} are vectors characterized by having $\mathcal{O}(z)$ non-zero entries. Given this structure, the matrix-vector multiplication complexity for K_X is $O(nnz \cdot T)$, i.e., $\text{mvm}(K_X) = O(z \cdot T)$.*

642 *Proof.* We will demonstrate this in the following discussion by concentrating on the k^{th} entry of the
643 output vector, specifically $K_X \mathbf{v}$, for any arbitrary vector \mathbf{v} :

$$(K_X \mathbf{v})_k = \sum_j K_X(k, j) v_j = \sum_j \left(\sum_{i=1}^5 s_i(\pi_k, \pi_j) \right) v_j = \sum_j \left(\sum_{i=1}^5 \phi^{a_i}(\pi_k)^T \phi^{b_i}(\pi_j) \right) v_j,$$

$$= \sum_{i=1}^5 \left(\sum_j \phi^{a_i}(\pi_k)^T \phi^{b_i}(\pi_j) v_j \right) = \sum_{i=1}^5 \phi^{a_i}(\pi_k)^T \left(\sum_j \phi^{b_i}(\pi_j) v_j \right).$$

644 Given that for all i , ϕ^{b_i} possesses only $\mathcal{O}(z)$ non-zero entries for any π_j , the computation of
645 $\sum_j \phi^{b_i}(\pi_j) v_j$ requires $\mathcal{O}(z)$ operations. This implies that the expression $\sum_j \phi^{b_i}(\pi_j) v_j$ also necess-
646 itates $\mathcal{O}(z)$ computation. Applying a similar rationale to ϕ^{a_i} , it follows that computing $(K_X v)_k$
647 demands only $\mathcal{O}(z)$ operations. Extending this argument to all entries of the output vector, it is
648 evident that computing $K_X \mathbf{v}$ requires only $\mathcal{O}(z \cdot T)$ computation \square

649 Utilizing Claim 5, it suffices to complete the proof by showcasing that these exist vectors ϕ^{a_i} and
650 ϕ^{b_i} , each with only $\mathcal{O}(k^2)$ non-zero elements, corresponding to each s_i as specified in Algorithm 2.
651 Additionally, these vectors ensure that $s_i(\pi_1, \pi_2) = \phi^{a_i}(\pi_1)^T \phi^{b_i}(\pi_2)$. We will next establish such
652 vectors for all s_i terms. Starting with the s_1 term below.

653 **Showcasing** $s_1(\pi_1, \pi_2) = \phi^{a_1}(\pi_1)^T \phi^{b_1}(\pi_2)$ for sparse $\phi^{a_1}(\pi_1)$ and $\phi^{b_1}(\pi_2)$ vectors. We begin
654 by manipulating s_1 , as defined in Equation 12. For the sake of brevity, their scalar factors will be
655 omitted in the following explanation.

$$\begin{aligned} s_1(\pi_1, \pi_2) &= \sum_{1 \leq l < m \leq n | l, m \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2}) \cdot \eta_{l,m}(\pi_1, \pi_2), \\ &= \sum_{1 \leq l < m \leq n | l, m \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2}) \cdot (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \cdot (\mathbf{p}_{l < m}^{\pi_2} - \mathbf{p}_{l > m}^{\pi_2}), \\ &= \sum_{1 \leq l < m \leq n | l, m \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot (\mathbf{p}_{i < j}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2}) \cdot (\mathbf{p}_{l < m}^{\pi_2} - \mathbf{p}_{l > m}^{\pi_2}), \\ &= \sum_{1 \leq l < m \leq n} \underbrace{w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot (\mathbf{p}_{i < j}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \cdot \mathbf{1}_{\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1} \in [k]}}_{:= \phi_{l,m}^{a_1}(\pi_1)} \\ &\quad \cdot \underbrace{w_s(\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2}) \cdot (\mathbf{p}_{l < m}^{\pi_2} - \mathbf{p}_{l > m}^{\pi_2}) \cdot \mathbf{1}_{\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2} \in [k]}}_{:= \phi_{l,m}^{b_1}(\pi_2)}, \\ &= (\phi^{a_1}(\pi_1)^T \phi^{b_1}(\pi_2)). \end{aligned} \tag{22}$$

656 Both ϕ^{a_1} and ϕ^{b_1} are sparse by design, taking non-zero values only when l and m appear in the top- k
657 rankings. This demonstrates the existence of sparse vectors for the s_1 term. Next, we will establish
658 the same for the s_2 and s_3 terms.

659 **Showcasing sparse vectors for s_2 and s_3 .** We begin by manipulating s_2 , as defined in Equation 13,
660 while ignoring its scalar factor. We will exploit symmetry between s_2 and s_3 terms.

$$\begin{aligned} s_2(\pi_1, \pi_2) &= \sum_{l \in I_1 \cap I_2 | m \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}), \\ &= \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_2}, \cdot) \sum_{m \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}), \\ &= \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_2}, \cdot) \left(\sum_{m \in I_1} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) - \sum_{m \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \right), \\ &= \sum_{l \in [n]} \underbrace{\mathbf{1}_{\mathbf{p}_l^{\pi_2} \in [k]} w_s(\mathbf{p}_l^{\pi_2}, \cdot)}_{:= \phi_l^{b_21}(\pi_2)} \underbrace{\mathbf{1}_{\mathbf{p}_l^{\pi_1} \in [k]} \sum_{m \in I_1} w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1})}_{:= \phi_l^{a_21}(\pi_1)} \end{aligned}$$

$$- \sum_{l,m \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_2}, \cdot) w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}), \quad (23)$$

$$\begin{aligned} &= \phi^{a_{21}}(\pi_1)^T \phi^{b_{21}}(\pi_2) - \sum_{l,m \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_2}, \cdot) w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}), \\ &= \phi^{a_{21}}(\pi_1)^T \phi^{b_{21}}(\pi_2) + \sum_{l,m \in [n]} \underbrace{-w_s(\mathbf{p}_l^{\pi_2}, \cdot) \mathbf{1}_{\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2} \in [k]}}_{:= \phi_{l,m}^{b_{22}}} \\ &\quad \cdot \underbrace{w_s(\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1}) (\mathbf{p}_{l < m}^{\pi_1} - \mathbf{p}_{l > m}^{\pi_1}) \mathbf{1}_{\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1} \in [k]}}_{:= \phi_{l,m}^{a_{22}}}, \end{aligned} \quad (24)$$

$$\begin{aligned} &= \phi^{a_{21}}(\pi_1)^T \phi^{b_{21}}(\pi_2) + \phi^{a_{22}}(\pi_1)^T \phi^{b_{22}}(\pi_2), \\ &= \underbrace{[\phi^{a_{21}}(\pi_1); \phi^{a_{22}}(\pi_2)]^T}_{:= \phi^{a_2}(\pi_1)^T} \underbrace{[\phi^{a_{21}}((\pi_2)); \phi^{b_{22}}((\pi_2))]}_{:= \phi^{b_2}(\pi_2)} = \phi^{a_2}(\pi_1)^T \phi^{b_2}(\pi_2). \end{aligned} \quad (25)$$

661 Equation 25 demonstrates the existence of vectors ϕ^{a_2} and ϕ^{b_2} for the s_2 term. The vectors $\phi^{a_{21}}$ and
662 $\phi^{a_{22}}$, possessing $\mathcal{O}(k)$ and $\mathcal{O}(k^2)$ non-zero entries respectively, are defined in Equations 23 and 24.
663 Consequently, the ϕ^{a_2} vector has $\mathcal{O}(k^2)$ non-zero entries. Similarly, it can be shown that ϕ^{b_2} contains
664 $\mathcal{O}(k^2)$ non-zero entries, thus fulfilling the proof requirements for proving the s_2 term. For the s_3
665 term, we observe a symmetry between s_2 and s_3 , namely $s_3(\pi_1, \pi_2) = s_2(\pi_2, \pi_1)$. This symmetry
666 makes it trivial to satisfy the requirements, as further highlighted by the following equation:

$$s_3(\pi_1, \pi_2) = s_2(\pi_2, \pi_1) = \phi^{a_2}(\pi_2)^T \phi^{b_2}(\pi_1) = \underbrace{\phi^{b_2}(\pi_1)^T}_{:= \phi^{a_3}(\pi_1)} \underbrace{\phi^{a_2}(\pi_2)}_{:= \phi^{b_3}(\pi_2)} = \phi^{a_3}(\pi_1)^T \phi^{b_3}(\pi_2). \quad (26)$$

667 **Showcasing sparse vectors** $s_4(\pi_1, \pi_2) = \phi^{4a}(\pi_1)^T \phi^{4b}(\pi_2)$. We begin by manipulating the s_4
668 term without scalar, as defined in Equation 15.

$$\begin{aligned} s_4(\pi_1, \pi_2) &= - \sum_{l \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_m^{\pi_2}, \cdot), \\ &= - \sum_{l \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot \left(\sum_{m \in I_2} w_s(\mathbf{p}_m^{\pi_2}, \cdot) - \sum_{m \in I_1 \cap I_2} w_s(\mathbf{p}_m^{\pi_2}, \cdot) \right). \end{aligned}$$

669 Observing that $\bar{w} := \sum_{m \in I_2} w_s(\mathbf{p}_m^{\pi_2}, \cdot)$ represents a constant value that does not depend on I_2 , we can
670 further simplify the above expression for s_4 as follows:

$$\begin{aligned} s_4(\pi_1, \pi_2) &= - \sum_{l \in I_1 \setminus I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot \left(\bar{w} - \sum_{m \in I_1 \cap I_2} w_s(\mathbf{p}_m^{\pi_2}, \cdot) \right), \\ &= - \left(\bar{w} - \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \right) \cdot \left(\bar{w} - \sum_{m \in I_1 \cap I_2} w_s(\mathbf{p}_m^{\pi_2}, \cdot) \right), \\ &= -\bar{w}^2 + \bar{w} \left(\sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) + \sum_{m \in I_1 \cap I_2} w_s(\mathbf{p}_m^{\pi_2}, \cdot) \right) \\ &\quad - \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \sum_{m \in I_1 \cap I_2} w_s(\mathbf{p}_m^{\pi_2}, \cdot). \end{aligned} \quad (27)$$

671 Next, to simplify the above equation, we first focus on the second term and have the following:

$$\bar{w} \left(\sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) + \sum_{m \in I_1 \cap I_2} w_s(\mathbf{p}_m^{\pi_2}, \cdot) \right)$$

$$\begin{aligned}
&= \sum_{l \in [n]} \underbrace{\mathbf{1}_{\mathbf{p}_l^{\pi_1} \in [k]} w_s(\mathbf{p}_l^{\pi_1}, \cdot)}_{:= \phi_l^{4a_1}(\pi_1)} \underbrace{\mathbf{1}_{\mathbf{p}_l^{\pi_2} \in [k]} \bar{w}}_{:= \phi_l^{4b_1}(\pi_2)} + \sum_{m \in I_1 \cap I_2} \bar{w} \cdot w_s(\mathbf{p}_m^{\pi_2}, \cdot), \\
&= \phi^{4a_1}(\pi_1)^T \phi^{4b_1}(\pi_2) + \sum_{m \in [k]} \underbrace{\mathbf{1}_{\mathbf{p}_m^{\pi_1} \in [k]} \bar{w}}_{:= \phi_m^{4a_2}(\pi_1)} \underbrace{w_s(\mathbf{p}_m^{\pi_1}, \cdot)}_{:= \phi_m^{4b_2}(\pi_2)} \\
&= \phi^{4a_1}(\pi_1)^T \phi^{4b_1}(\pi_2) + \phi^{4a_2}(\pi_1)^T \phi^{4b_2}(\pi_2). \tag{28}
\end{aligned}$$

672 Next, we simplify the third and last term in the Equation 27 as follows:

$$\begin{aligned}
\sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \sum_{m \in I_1 \cap I_2} w_s(\mathbf{p}_m^{\pi_2}, \cdot) &= \sum_{l \in [n], m \in [n]} \underbrace{w_s(\mathbf{p}_l^{\pi_1}, \cdot) \mathbf{1}_{\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1} \in [k]}}_{:= \phi_{l,m}^{4a_3}(\pi_1)} \underbrace{w_s(\mathbf{p}_m^{\pi_2}, \cdot) \mathbf{1}_{\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2} \in [k]}}_{:= \phi_{l,m}^{4b_3}(\pi_2)}, \\
&= \phi^{4a_3}(\pi_1)^T \phi^{4b_3}(\pi_2). \tag{29}
\end{aligned}$$

673 Next, combining the results from Equations 27, 28, and 29, we obtain the following:

$$\begin{aligned}
s_4(\pi_1, \pi_2) &= \underbrace{[\bar{w}, \phi^{4a_1}(\pi_1); \phi^{4a_1}(\pi_1); \phi^{4a_3}(\pi_1)]^T}_{:= \phi^{4a}(\pi_1)^T} \underbrace{[-\bar{w}; \phi^{4b_1}(\pi_2); \phi^{4b_2}(\pi_2); -\phi^{4b_3}(\pi_2)]}_{:= \phi^{4b}(\pi_2)} \\
&= \phi^{4a}(\pi_1)^T \phi^{4b}(\pi_2). \tag{30}
\end{aligned}$$

674 Equation 30 showcases both ϕ^{4a} and ϕ^{4b} has three components with having only $\mathcal{O}(k^2)$ non-zero entries, thus fulfilling the requirements for the s_4 term. Next, we focus on the s_5 term.

676 **Showcasing sparse vectors** $s_5(\pi_1, \pi_2) = \phi^{5a}(\pi_1)^T \phi^{5b}(\pi_2)$. We begin by examining the s_5 term, 677 excluding its scalar component, as outlined in Equation 16.

$$\begin{aligned}
s_5(\pi_1, \pi_2) &= (n - |I_1 \cup I_2|) \cdot \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot), \\
&= (n - (2k - |I_1 \cap I_2|)) \cdot \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot), \\
&= (n - 2k) \cdot \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) + |I_1 \cap I_2| \cdot \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot), \\
&= \sum_{l \in I_1 \cap I_2} \sqrt{n - 2k} \cdot w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot \sqrt{n - 2k} \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) \\
&\quad + \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) \cdot |I_1 \cap I_2|, \tag{31} \\
&= \sum_{l \in [n]} \underbrace{\sqrt{n - 2k} \cdot w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot \mathbf{1}_{\mathbf{p}_l^{\pi_1} \in [k]}}_{:= \phi_l^{5a_1}(\pi_1)} \underbrace{\sqrt{n - 2k} \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) \cdot \mathbf{1}_{\mathbf{p}_l^{\pi_2} \in [k]}}_{:= \phi_l^{5b_1}(\pi_2)} \\
&\quad + \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) \cdot |I_1 \cap I_2|, \\
&= \phi^{5a_1}(\pi_1)^T \phi^{5b_1}(\pi_2) + \sum_{l \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot) \cdot \sum_{m \in I_1 \cap I_2} 1, \\
&= (\phi^{5a_1}(\pi_1)^T \phi^{5b_1}(\pi_2) + \sum_{l \in I_1 \cap I_2, m \in I_1 \cap I_2} w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot w_s(\mathbf{p}_l^{\pi_2}, \cdot)), \\
&= \phi^{5a_1}(\pi_1)^T \phi^{5b_1}(\pi_2) + \sum_{l \in [n], m \in [n]} \underbrace{w_s(\mathbf{p}_l^{\pi_1}, \cdot) \cdot \mathbf{1}_{\mathbf{p}_l^{\pi_1}, \mathbf{p}_m^{\pi_1} \in [k]}}_{:= \phi_{l,m}^{5a_2}(\pi_1)} \underbrace{w_s(\mathbf{p}_l^{\pi_2}, \cdot) \cdot \mathbf{1}_{\mathbf{p}_l^{\pi_2}, \mathbf{p}_m^{\pi_2} \in [k]}}_{:= \phi_{l,m}^{5b_2}(\pi_2)}, \tag{32}
\end{aligned}$$

$$\begin{aligned}
&= \phi^{5a_1}(\pi_1)^T \phi^{5b_1}(\pi_2) + \phi^{5a_2}(\pi_1)^T \phi^{5b_2}(\pi_2), \\
&= \underbrace{[\phi^{5a_1}(\pi_1); \phi^{5a_2}(\pi_1)]^T}_{:=\phi^{5a}(\pi_1)^T} \underbrace{[\phi^{5b_1}(\pi_2) + \phi^{5b_2}(\pi_2)]}_{:=\phi^{5b}(\pi_2)} = \phi^{5a}(\pi_1)^T \phi^{5b}(\pi_2).
\end{aligned} \tag{33}$$

678 The equation shows that $s_5(\pi_1, \pi_2) = \phi^{5a}(\pi_1)^T \phi^{5b}(\pi_2)$, where both ϕ^{5a} and ϕ^{5b} possess components with a maximum number of non-zero entries, as indicated in Equations 31 and 32. This completes the proof requirements for the s_5 term.

681 By combining the results from Equations 22, 25, 26, 30, and 33, we have demonstrated the existence
682 of vectors ϕ^{a_i} and ϕ^{b_i} , each containing only $\mathcal{O}(k^2)$ non-zero elements, and have established that
683 $s_i(\pi_1, \pi_2) = \phi^{a_i}(\pi_1)^T \phi^{b_i}(\pi_2)$ for each $i \in 1, 2, 3, 4, 5$. In conjunction with Claim 5, this completes
684 the proof. \square

685 B Proposed GP-TopK Bandit Algorithm- Omitted Details

686 This section includes the proofs that were omitted from Section 4, presented in the following order:

- 687 • Section B.1 outlines a brief of Gaussian process regression for any domain.
- 688 • Section B.2 summarizes the committed details about the local search utilized for optimizing
689 the UCB function.
- 690 • Section B.3 provides the removed proof for the Theorem 2 concerning the overall time for
691 the bandit algorithm.
- 692 • Section B.4 provides the proof for Theorem 3 concerning regret analysis of the proposed
693 bandit algorithm.

694 B.1 Gaussian Process Regression

695 In GP regression [22], the training data are modeled as noisy measurements of a random function
696 f drawn from a GP prior, denoted $f \sim \mathcal{N}(0, k(\cdot, \cdot))$, where $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a kernel function
697 over any domain \mathcal{X} . The observed training pairs (\mathbf{x}_i, y_i) are collected as $X = [\mathbf{x}_1, \dots, \mathbf{x}_T]$ and
698 $\mathbf{y} = [y_1, \dots, y_T] \in \mathbb{R}^T$, where, for an input \mathbf{x}_i , the observed value is modeled as $\mathbf{y}_i = f(\mathbf{x}_i) + \epsilon$,
699 with $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$. The kernel matrix on data is $K_X = [k(\mathbf{x}_i, \mathbf{x}_j)]_{i,j=1}^T \in \mathbb{R}^{T \times T}$. The posterior
700 mean $\mu_{f|\mathcal{D}}$ and variance $\sigma_{f|\mathcal{D}}$ functions for GPs are:

$$\mu_{f|\mathcal{D}}(\mathbf{x}) := \mathbf{k}_x^T \mathbf{z} \tag{34}$$

$$\sigma_{f|\mathcal{D}}(\mathbf{x}) := k(\mathbf{x}, \mathbf{x}) - \mathbf{k}_x^T (K_X + \sigma^2 I)^{-1} \mathbf{k}_x \tag{35}$$

701 where $\mathbf{k}_x \in \mathbb{R}^T$ has as its i^{th} entry $k(\mathbf{x}, \mathbf{x}_i)$, $\mathbf{z} = (K_X + \sigma^2 I)^{-1} \mathbf{y}$, and I is an identity matrix. For
702 GP regression on an arbitrary domain \mathcal{X} , the kernel function must be a p.d. kernel [23].

703 Naive approaches rely on the Cholesky decomposition of the matrix $K_X + \sigma^2 I$, which takes $\Theta(T^3)$
704 time [23]. To circumvent the $\Theta(T^3)$ runtime, recent works use iterative algorithms such as the
705 conjugate gradient algorithm, which facilitate GP inference by exploiting fast kernel matrix-vector
706 multiplication (MVM) algorithms, i.e., $\mathbf{v} \mapsto \mathbf{K}_X \mathbf{v}$ [3]. In practice, these methods yield highly
707 accurate approximations for GP posterior functions with a complexity of $\Theta(p \cdot T^2)$ for p CG
708 iterations, as $\text{mvm}(K_X) = T^2$, and $\text{mvm}(M)$ is the operation count for multiplying matrix M by
709 a vector. $p \ll T$ proves to be efficient in practical application [3].

710 B.2 Contextual GP Reward Model

711 Optimizing the \mathcal{AF} , i.e., UCB function, poses a significant challenge due to its enormous size of Π^k .
712 Drawing inspiration from prior research on Bayesian optimization within combinatorial spaces, we
713 employ a breadth-first local search (BFLS) to optimize the UCB acquisition function [2, 19]. The
714 BFLS begins with the selection of several random top-k rankings. Subsequently, each specific top-k
715 ranking is compared with the UCB values of its neighboring rankings, proceeding to the one with the
716 highest UCB value.

717 The neighbors of a top-k ranking include all its permutations and the permutations of modified top-k
 718 rankings obtained by swapping one item with any of the remaining items. For any top-k ranking,
 719 there are $(n - k) \cdot k! + k!$ neighbors, which is often not huge as k is often ≤ 6 . This search continues
 720 until no neighboring top-k ranking with a higher value is discovered. Although BFLS is a local
 721 search, the initial random selection and multiple restart points help it evade local minima, a strategy
 722 that previous studies have corroborated [19].

723 **B.3 Assessing GP-TopK Compute Requirements**

Theorem 2. *Assuming a fixed number of iterations required by the iterative algorithms, the total computational time for running the GP-TopK bandit algorithm for T rounds of top-k recommendations, using the contextual product kernel (Equation 6), is $\mathcal{O}(k^2 c \ell T^2)$. This applies to WK, CK, and WCK top-k ranking kernels, where ℓ is the number of local search evaluations for selecting the next arm in every round.*

724 *Proof.* The proof can be straightforwardly derived by combining the results presented in Table 1,
 725 which succinctly summarizes the time complexities for each step of computing the UCB using both
 726 feature and kernel approaches. It is important to emphasize that iterative algorithms enhance results
 727 from $\mathcal{O}(T^4)$ to $\mathcal{O}(T^3)$ in computational complexity. Furthermore, these algorithms can further
 728 reduce complexity to $\mathcal{O}(T^2)$ when used with the feature approach.

729 The results presented in Table 1 can be validated through straightforward observations and by
 730 leveraging findings from previous Sections 2. Specifically, Section 2 offers proof for the **mvm**(K_X)
 731 row explicitly. For the *compute* K_{X_t} row, the complexity of kernel approaches is deduced from
 732 Algorithms 2 and 3. For feature approaches, the *compute* K_{X_t} row is inferred from the sparsity of the
 733 feature representations as stated in Claim 3. Lastly, the *memory* row is straightforwardly deduced for
 734 the kernel approach by counting its entries. For the feature approach, it is derived from the sparsity of
 735 the feature representations. \square

736 **B.4 Regret Analysis**

737 In this section, we revisit Theorem 3 and provide its proof. The proofs build on the work by Krause
 738 et al. [14], delivering results for bounding the contextual regret in the context of the top-k ranking
 739 problem. To set the stage for our regret analysis, let's first define the critical term *maximum mutual
 740 information*, denoted by γ_t , is given below:

$$\gamma_t := \max_{X \subseteq \mathcal{X}: |X|=t} I(y_X; f), \quad I(y_X; f) = H(y_X) - H(y_X|f),$$

741 where $I(y_X; f)$ quantifies the reduction in uncertainty (measured in terms of differential Shannon
 742 entropy) about f achieved by revealing y_A [28]. In Gaussian observation case, the entropy can be
 743 computed in closed form: $H(N(\mu, \Sigma)) = \frac{1}{2} \log |2\pi e \Sigma|$, so that $I(y_X; f) = \frac{1}{2} \log |I + \sigma^{-2} K_X|$,
 744 where $K_X = [k(x, x')]_{x, x' \in X}$ is the Gram matrix of k evaluated on set $X \subseteq \mathcal{X}$. For the contextual
 745 bandit algorithm, X represents contexts and arms considered until round t .

746 Before proving Theorem 3, we align the Krause et al. [14] results with our notation for consistency.
 747 Furthermore, we modify β_t to accommodate embeddings encompassing negative values, aligning
 748 with the fact that contextual embeddings may exhibit negative dimensions.

Proposition 1 (Theorem 1, [14]). *Let $\delta \in (0, 1)$, and the unknown reward function \hat{f} be sampled*

749

from the known GP prior with known noise variance σ^2 . Suppose one of the following holds:

1. Assumption 1 holds and set $\beta_t = 2 \log(|\mathcal{X}|t^2\pi^2/6\delta)$.
2. Assumption 2 holds and set $\beta_t = 2B^2 + 300\gamma_t \ln^3(t/\delta)$.

Then the cumulative regret \mathcal{R}_T of the contextual GP bandit algorithm with the UCB acquisition function is bounded by $\tilde{\mathcal{O}}(\sqrt{C_1 T \gamma_T \beta_T})$ w.h.p. Precisely, $\Pr\{R_T \leq \sqrt{C_1 T \gamma_T \beta_T} + 2 \forall T \geq 1\} \geq 1 - \delta$, where, $C_1 = 8/\log(1 + \sigma^{-2})$ and the notation $\tilde{\mathcal{O}}$ hides logarithmic factors in $n, \frac{1}{\delta}$ and T .

750

751 Proposition 1 shows that the regret \mathcal{R}_T for the contextual GP bandit algorithm, utilizing the UCB
752 acquisition function is bounded with high probability within $\tilde{\mathcal{O}}(\sqrt{C_1 T \gamma_T \beta_T})$, where the notation
753 $\tilde{\mathcal{O}}$ hides logarithmic factors in $n, \frac{1}{\delta}$ and T . To ascertain the $\tilde{\mathcal{O}}$ order for \mathcal{R}_T , it is imperative to first
754 bound the $\tilde{\mathcal{O}}$ order of $\gamma_T \beta_t$. We begin by examining the γ_T term in the subsequent proposition.

Proposition 2. Under the assumptions of Theorem 3, γ_T can be succinctly characterized as $\gamma_T = \mathcal{O}(n^2 c \log(n^2 T) + c \log T)$, which also simplifies to $\tilde{\mathcal{O}}(n^2 c)$, where the $\tilde{\mathcal{O}}$ notation omits logarithmic factors in n and T .

755 *Proof.* For the GP bandit algorithm with the UCB acquisition function, $\gamma_T = C \cdot$
756 $\log(|I + \sigma^{-2} K_{X_T}|)$, where C equals $(1/2) \cdot (1 - 1/e)^{-1}$ and K_{X_T} represents the kernel ma-
757 trix computed over contexts and arms across T rounds [28, 14]. Precisely, K_{X_T} is calculated using
758 the contextual kernel defined in Equation 6. It is applied to contexts and top-k ratings from the
759 feedback data \mathcal{D}_t , corresponding to Line 6 of the generic contextual bandit Algorithm 1.

760 Next, we leverage the characteristic of the contextual kernel being a product kernel. Consequently, the
761 maximum mutual information term for the joint kernel, γ_T , can be upper bounded by $c \cdot (\gamma_T^\pi + \log T)$,
762 where c denotes the dimensionality of contexts and γ_T^π represents the maximum information gain in a
763 non-contextual setting [14]. Specifically, γ_T^π is computed similarly but is confined to top-k rankings.
764 That is, $\gamma_T^\pi = C \cdot \log(|I + \sigma^{-2} K_{X^\pi}|)$, with $K_{X_T^\pi}$ being calculated exclusively using the top-k
765 kernels on the top-k rankings as selected by the bandit algorithm. X_T^π represents the top-k rankings
766 selected by the bandit algorithm, i.e., excluding the contexts from the collected feedback.

767 Recalling the formulation for top-k rankings kernels, we have $K_{X_T} = \Phi_{X_T^\pi}^T \Phi_{X_T^\pi}$, where $\Phi_{X^\pi} \in$
768 $\mathbb{R}^{\binom{n}{2} \times T}$ comprises feature columns pertinent to the top- k ranking kernels, as elucidated in Section A.
769 Utilizing the Weinstein–Aronszajn identity, γ_T^π is expressed as $C \cdot \log(|I + \sigma^{-2} \Phi_{X_T^\pi} \Phi_{X_T^\pi}^T|)$. Further,
770 we deduce that $\gamma_T^\pi \leq C \cdot \sum_{i=1}^{\binom{n}{2}} \log(|1 + \sigma^{-2} \lambda_i|)$, where λ_i is an eigenvalue of $\Phi_{X_T^\pi} \Phi_{X_T^\pi}^T$. Given
771 the Gershgorin circle theorem, which bounds all eigenvalues of a matrix by the maximum absolute
772 sum of its rows, therefore we can conclude that $\gamma_T^\pi = \mathcal{O}(n^2 \log(n^2 T))$, as for all the columns of the
773 Φ_{X^π} have bounded normed as given in Claims 2 and 3, i.e., $\|\phi(\pi)\|_2^2 \leq 1$ [30].
774 By combining $\gamma_T^\pi = \mathcal{O}(n^2 \log(n^2 T))$ with the contextual product kernel, we obtain $\gamma_T =$
775 $\mathcal{O}(n^2 c \log(n^2 T) + c \log T)$, thereby providing the claimed bound in the proposition. \square

776 Next, we build on Propositions 1 and 2 to prove the main theorem regarding the regret of the proposed
777 GP-TopK bandit algorithm for top-k recommendations.

Theorem 3. If either Assumptions 1 or 2 hold, setting β_t as $2 \log\left(\frac{|\mathcal{C}| \cdot |\Pi^k| \cdot t^2 \cdot \pi^2}{6\delta}\right)$
and $300\gamma_t \ln^3\left(\frac{t}{\delta}\right)$ respectively, the cumulative regret \mathcal{R}_T of the GP-TopK bandit
algorithm for top-k recommendations can, with at least $1 - \delta$ probability, be bounded by $\tilde{\mathcal{O}}(n \sqrt{C_1 T c (\log |\mathcal{C}| + k + \log(T^2 \pi^2 / 6\delta))})$ under Assumption 1, and
 $\tilde{\mathcal{O}}(n \sqrt{C_1 (2B^2 c + 300n^2 c^2 \ln^3(T/\delta)) T})$ under Assumption 2. Here, $C_1 = \frac{8}{\log(1 + \xi^{-2})}$, and $\tilde{\mathcal{O}}$
excludes logarithmic factors related to n, k , and T .

778 *Proof.* We will prove the above theorem for both cases separately.

779 **For Assumption-1.** Given $|\mathcal{C}|$ is finite and $\beta_T = 2 \log(|\mathcal{D}|T^2\pi^2/6\delta)$. First, we focus on bounding
780 β_T as follows:

$$\begin{aligned}\beta_T &= 2 \log(|\mathcal{D}|T^2\pi^2/6\delta) \\ &= \mathcal{O}(\log|\mathcal{C}| + \log|\Pi^k| + \log(T^2\pi^2/6\delta))\end{aligned}$$

781 As $\binom{n}{k} \leq n^k$ and $k! \leq k^k$, we also have $\log|\Pi^k| = \log(\binom{n}{k}k!) \leq \log(n^k k^k) = \mathcal{O}(k \log(nk))$,
782 which implies that $\beta_T = \mathcal{O}(\log|\mathcal{C}| + k \log(nk) + \log(T^2\pi^2/6\delta))$. Combining this with Proposition
783 2, we have following:

$$\begin{aligned}\mathcal{O}(\gamma_T \beta_T) &= \mathcal{O}((n^2 c \log(n^2 T) + c \log T)(\log|\mathcal{C}| + k \log(nk) + \log(T^2\pi^2/6\delta))) \\ &= \mathcal{O}(n^2 c \log(n^2 T)(\log|\mathcal{C}| + k \log(nk) + \log(T^2\pi^2/6\delta))) \quad (\text{Ignoring } c \log T \text{ term}) \\ &= \tilde{\mathcal{O}}(n^2 c (\log|\mathcal{C}| + k + \log(T^2\pi^2/6\delta))).\end{aligned}$$

784 Thus, we showcase the asserted bound for the regret \mathcal{R}_T as $\tilde{\mathcal{O}}(\sqrt{C_1 T \gamma_T \beta_T}) =$
785 $\tilde{\mathcal{O}}(n \sqrt{C_1 T c (\log|\mathcal{C}| + k + \log(T^2\pi^2/6\delta))})$.

786 **For Assumption-2.** Given $\|f\|_k \leq B$ and $\beta_t = 2B^2 + 300\gamma_t \ln^3(t/\delta)$. First, we bound the β_T term
787 using Proposition 2 as follows:

$$\begin{aligned}\beta_T &= 2B^2 + 300 \cdot \gamma_T \cdot \ln^3(T/\delta), \\ &= 2B^2 + 300 \cdot (n^2 c \log(n^2 T) + c \log T) \cdot \ln^3(T/\delta).\end{aligned}$$

788 Using the above result, we have the following:

$$\begin{aligned}\mathcal{O}(\sqrt{C_1 T \gamma_T \beta_T}) &= \mathcal{O}\left(\sqrt{C_1 T \gamma_T \cdot (2B^2 + 300 \cdot \gamma_T \cdot \ln^3(T/\delta))}\right), \\ &= \mathcal{O}\left(\sqrt{C_1 T n^2 c \log(n^2 T) \cdot (2B^2 + 300 \cdot n^2 c \log(n^2 T) \cdot \ln^3(T/\delta))}\right), \\ &= \tilde{\mathcal{O}}\left(n \sqrt{C_1 T c (2B^2 + 300 n^2 c \ln^3(T/\delta))}\right).\end{aligned}$$

789 □

790 **Comparison with Srinivas et al. (2010).** Using the identity kernel for top-k rankings, we can
791 develop a finite-dimensional feature for the contextual kernel and apply Theorem 5 by Srinivas et al.
792 (2010). Given that $\gamma_T = \mathcal{O}(n^k c \log T)$, the regret bounds are as follows under both assumptions. For
793 instance, the calculations for the $\mathcal{O}(\sqrt{C_1 T \gamma_T \beta_T})$ under the Assumption 2 are as follows:

$$\begin{aligned}\mathcal{O}(\sqrt{C_1 T \gamma_T \beta_T}) &= \mathcal{O}\left(\sqrt{C_1 T \gamma_T \cdot (2B^2 + 300 \cdot \gamma_T \cdot \ln^3(T/\delta))}\right), \\ &= \mathcal{O}\left(\sqrt{C_1 T (n^k c \log T) \cdot (2B^2 + 300 \cdot (n^k c \log T) \cdot \ln^3(T/\delta))}\right), \\ &= \tilde{\mathcal{O}}\left(n^{\frac{k}{2}} \sqrt{C_1 T c (2B^2 + 300 n^k c \ln^3(T/\delta))}\right).\end{aligned}$$

794 Similarly, we can analogously perform the analysis for Assumption 1 and combine it with Proposi-
795 tion 1 to obtain the regret bounds mentioned in the Table 3.

796 C Experiments – Omitted Details

797 This section presents omitted details from the main body of the text.

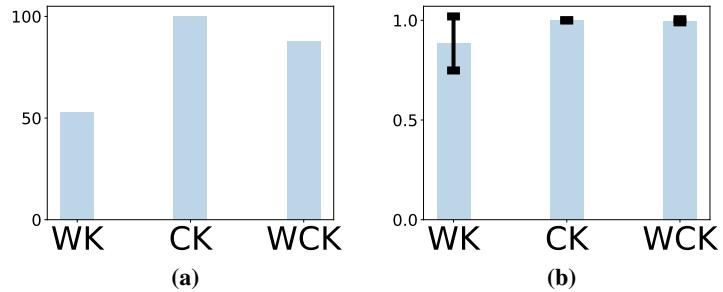


Figure 4: Local search results for optimizing combinatorial objectives in Π^k for $n = 50$ and $k = 6$. For details, see the textual description. Left (a) shows how many times out of 100 trials the local search recovers the exact maximizer, i.e., π^* , and right plot (b) shows the average value of the objective for the returned maximizer. These results indicate that the local search utilized in this work is effective.

798 C.1 Compute resources

799 We utilized multiple NVIDIA Tesla M40 GPUs with 40 GB RAM on our in-house cluster for our
800 experiments. The experiments in Section 5 required approximately 5 GPU-hours for small arm
801 space and 24 GPU-hours per iteration for large arm space. We conducted about 50 to 100 iterations
802 throughout the project. The results reported in Section C.3 required the same computational resources
803 as the large arm space experiments.

804 C.2 Bandit Simulation and Hyper-parameter Configurations – Omitted Details

805 To set up the simulation, we utilized embeddings trained on the MovieLens dataset using a collabora-
806 tive filtering approach [6]. We consider a $1M$ variant of the MovieLens dataset, which contains 1
807 million ratings from 6040 users for 3677 items. Specifically, we train user embeddings \mathbf{c}_u and item
808 embeddings θ_i such that the user’s attraction to the items are captured by the inner product of the user
809 embedding with the item embeddings, respectively. Both context and item embeddings, i.e., \mathbf{c}_u and
810 θ_i , are 5-dimensional, optimized by considering the 5-fold performance on this dataset. The reward
811 provided in our experiments is contaminated with zero mean and standard deviation equals 0.05.

812 For the ϵ -greedy baselines, we considered various values of ϵ are considered, specifically $\epsilon =$
813 $\{0.01, 0.05, 0.1\}$. The outcomes are presented for the configuration that demonstrates optimal
814 performance. For *MAB-UCB* baseline, the algorithm has an upper confidence score $ucb(i) =$
815 $\bar{\mu}_i + \beta_{mab} \sqrt{\frac{2 \ln(t+1)}{n_i}}$ [11]. Here, $\bar{\mu}_i$ represents the average reward, n denotes the total number of
816 rounds, and n_i signifies the frequency of arm i being played. β_{mab} is a hyper-parameter. We evaluate
817 β_{mab} values within the set $\{0.1, 0.25, 0.5\}$ and disclose results for the best-performing configuration.
818 For the parameters of proposed GP-TopK bandit algorithms, we set $\beta_t = \beta_{gp} \cdot \log(|\mathcal{X}| \cdot t^2 \cdot \pi^2)$ with
819 $\beta_{gp} \in \{0.05, 0.1, 0.5\}$, reporting results the value that yields the best performance. The choice of β_t
820 is informed by prior work in GP bandits [28]. The selection of σ for all variants is determined by
821 optimizing the log-likelihood of the observed after very 10 rounds by considering values in the set
822 $\{0.01, 0.05, 0.1\}$.

823 C.3 Additional results

824 **Local search** results for optimizing combinatorial objectives in Π^k for $n = 50$ and $k = 6$. Specif-
825 ically, $\pi^* = \max_{\pi} \phi^r(\pi)^T \phi^r(\pi^*)$, where $\phi^r(\pi^*)$ represents the feature vector for Kendall kernels
826 on top-k rankings. Notably, for this optimization problem, it is known that the optimal value is 1
827 obtained by only π^* . Figure 4 shows results for this optimization problem when applied to WK, CK,
828 and WCK kernels.

829 **Reward results** for large arm space for the nDCG + diversity reward. Similar to Figure 3, a large
830 setup with $n = 50$ for $k = 3$ and $k = 6$, is considered. For $k = 6$, the possible arms are over
831 1.1×10^{10} possible top-k rankings. Given the vastness of this arm space, computing the optimal
832 arm for the diversity reward is not straightforward. Therefore, we focus on reporting the cumulative

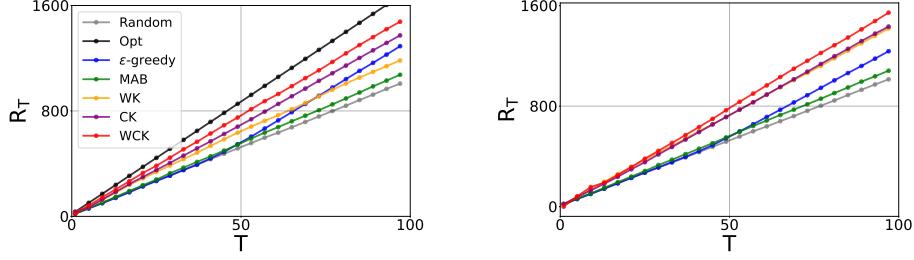


Figure 5: Comparative evaluation of bandit algorithms for large arm spaces for the nDCG + diversity reward, with $> 1.1 \times 10^5$ for the left plot and $> 1.1 \times 10^{10}$ for the right plot, respectively. Cumulative reward with respect to the rounds of the bandit algorithm is depicted. Results are averaged over 6 trials. In both settings, the WCK approach outperforms other baselines. For more details, see the textual description.

833 reward in Figure 5. We implement this setup using a Local search in batch mode, updating every 5
 834 round and considering a substantial horizon of $T = 100$ rounds. Specifically, we use 5 restarts, 5
 835 steps in every search direction, and start with 1000 initial candidates. Figure 5 shows that the WCK
 836 approach demonstrates superior performance, continuing to learn effectively even after extensive
 837 rounds.

838 **NeurIPS Paper Checklist**

839 **1. Claims**

840 Question: Do the main claims made in the abstract and introduction accurately reflect the
841 paper's contributions and scope?

842 Answer: [\[Yes\]](#)

843 Justification: Section 1 briefs both contributions and scope of this work.

844 **2. Limitations**

845 Question: Does the paper discuss the limitations of the work performed by the authors?

846 Answer: [\[Yes\]](#)

847 Justification: Section 6 reflects on the limitations of this work.

848 **3. Theory Assumptions and Proofs**

849 Question: For each theoretical result, does the paper provide the full set of assumptions and
850 a complete (and correct) proof?

851 Answer: [\[Yes\]](#)

852 Justification: Proofs of all Claims and Theorems are provided in the Appendix.

853 **4. Experimental Result Reproducibility**

854 Question: Does the paper fully disclose all the information needed to reproduce the main ex-
855 perimental results of the paper to the extent that it affects the main claims and/or conclusions
856 of the paper (regardless of whether the code and data are provided or not)?

857 Answer: [\[Yes\]](#)

858 Justification: Section 5 provides necessary details of bandit simulator and experimental
859 setups considered in this work.

860 **5. Open access to data and code**

861 Question: Does the paper provide open access to the data and code, with sufficient instruc-
862 tions to faithfully reproduce the main experimental results, as described in supplemental
863 material?

864 Answer: [\[Yes\]](#)

865 Justification: Our code can be accessed using this hyper-link.

866 **6. Experimental Setting/Details**

867 Question: Does the paper specify all the training and test details (e.g., data splits, hyper-
868 parameters, how they were chosen, type of optimizer, etc.) necessary to understand the
869 results?

870 Answer: [\[Yes\]](#)

871 Justification: Section 5 provides experimental details and a few remaining details are given
872 in the Appendix C.

873 **7. Experiment Statistical Significance**

874 Question: Does the paper report error bars suitably and correctly defined or other appropriate
875 information about the statistical significance of the experiments?

876 Answer: [\[Yes\]](#)

877 Justification: All figures reported in this work have errorbars with them.

878 **8. Experiments Compute Resources**

879 Question: For each experiment, does the paper provide sufficient information on the com-
880 puter resources (type of compute workers, memory, time of execution) needed to reproduce
881 the experiments?

882 Answer: [\[Yes\]](#)

883 Justification: Section C provides relevant information about compute resources.

884 **9. Code Of Ethics**

885 Question: Does the research conducted in the paper conform, in every respect, with the
886 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

887 Answer: [Yes]

888 10. Broader Impacts

889 Question: Does the paper discuss both potential positive societal impacts and negative
890 societal impacts of the work performed?

891 Answer: [Yes]

892 Justification: Section 6 reflects on the impact of this work.

893 11. Safeguards

894 Question: Does the paper describe safeguards that have been put in place for responsible
895 release of data or models that have a high risk for misuse (e.g., pretrained language models,
896 image generators, or scraped datasets)?

897 Answer: [NA]

898 Justification: this work does not release any such resource or asset.

899 12. Licenses for existing assets

900 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
901 the paper, properly credited and are the license and terms of use explicitly mentioned and
902 properly respected?

903 Answer: [Yes]

904 13. New Assets

905 Question: Are new assets introduced in the paper well documented and is the documentation
906 provided alongside the assets?

907 Answer: [Yes]

908 Justification: this work release only code with instructions for its usage.

909 14. Crowdsourcing and Research with Human Subjects

910 Question: For crowdsourcing experiments and research with human subjects, does the paper
911 include the full text of instructions given to participants and screenshots, if applicable, as
912 well as details about compensation (if any)?

913 Answer: [NA]

914 15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human 915 Subjects

916 Question: Does the paper describe potential risks incurred by study participants, whether
917 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)
918 approvals (or an equivalent approval/review based on the requirements of your country or
919 institution) were obtained?

920 Answer: [NA]