
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Zhang, Guangyi; Gionis, Aristides
Maximizing diversity over clustered data*

Published in:
Proceedings of the 2020 SIAM International Conference on Data Mining, SDM 2020

DOI:
[10.1137/1.9781611976236.73](https://doi.org/10.1137/1.9781611976236.73)

Published: 01/01/2020

Document Version
Publisher's PDF, also known as Version of record

Please cite the original version:
Zhang, G., & Gionis, A. (2020). Maximizing diversity over clustered data*. In C. Demeniconi, & N. Chawla (Eds.), *Proceedings of the 2020 SIAM International Conference on Data Mining, SDM 2020* (pp. 649-657). Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9781611976236.73>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Maximizing diversity over clustered data*

Guangyi Zhang[†]

Aristides Gionis[‡]

Abstract

Maximum diversity aims at selecting a diverse set of high-quality objects from a collection, which is a fundamental problem and has a wide range of applications, e.g., in Web search. Diversity under a uniform or partition matroid constraint naturally describes useful cardinality or budget requirements, and admits simple approximation algorithms [5]. When applied to clustered data, however, popular algorithms such as picking objects iteratively and performing local search lose their approximation guarantees towards maximum intra-cluster diversity because they fail to optimize the objective in a global manner. We propose an algorithm that greedily adds a pair of objects instead of a singleton, and which attains a constant approximation factor over clustered data. We further extend the algorithm to the case of monotone and submodular quality function, and under a partition matroid constraint. We also devise a technique to make our algorithm scalable, and on the way we obtain a modification that gives better solutions in practice while maintaining the approximation guarantee in theory. Our algorithm achieves excellent performance, compared to strong baselines in a mix of synthetic and real-world datasets.

1 Introduction

The *dispersion problem* aims at selecting a diverse set of objects from a collection. The problem has been studied extensively for different notions of dispersion and different settings. In their seminal work, Ravi et al. [17] studied the problem of selecting a subset of objects in a metric space so as to maximize some dispersion criterion, such as the minimum distance among all pairs of selected objects (MAXMIN or *remote edge*), and the sum of all pairwise distances (MAXSUM or *remote clique*),

and provided algorithms with constant-factor approximation guarantees. Stronger theoretical guarantees have been established later for the metric case [4, 11], while a polynomial-time approximation scheme is possible for negative-type distances [14].

The *diversification problem* is an extension of the dispersion problem that aims to maximize not only the pairwise distances of the selected objects, but also quality (or relevance in information retrieval) over those selected objects, via a weighted sum of both. The problem has important applications in databases, operations research, and information retrieval. Gollapudi and Sharma [8] showed that if the quality function is modular, the diversification problem can be re-casted as a dispersion problem. Borodin et al. [5] further extend the diversification problem to permit a submodular quality function, which is common in real-life applications. Furthermore, they generalized the diversification setting to allow for an arbitrary matroid constraint. The algorithms proposed by Borodin et al. rely on a greedy strategy that selects one object at a time, as well as a local-search strategy.

The aforementioned works, however, assume that the objects to be selected come from a single collection. On the other hand, in many cases, the collection of objects is split into a number of potentially-overlapping clusters. Consider the following example on analyzing communities of different political leaning in the Twitter network. Users from different communities discussing a topic may use the same hashtags, thus, inducing an overlapping clustering on the set of hashtags. We want to design a *description* functionality for each community, which creates descriptors by selecting a subset of hashtags used by users in each community. We are naturally interested in selecting a high-quality and diverse subset of hashtags, and at the same time, increasing exposure of different hashtags as much as possible, i.e., a hashtag will not be repeatedly selected.

In this paper we focus on the problem of maximizing diversity over data that form overlapping clusters, which models the setting discussed in the previous example. In more detail, we maximize the sum of dispersion within each selected set. The setting can be seen as applying MAXSUM diversification to each cluster separately under a global partition matroid constraint. Sim-

*This work was supported by the Academy of Finland project “Active knowledge discovery in graphs (AGRA)” (313927), the EC H2020 RIA project “SoBigData++” (871042), and the Wallenberg AI, Autonomous Systems and Software Program (WASP).

[†]Department of Computer Science, KTH Royal Institute of Technology, Sweden. This work was done while the author was with Aalto University.

[‡]Department of Computer Science, KTH Royal Institute of Technology, Sweden, and Department of Computer Science, Aalto University, Finland.

ilar to Borodin et al., we assume that distances between our objects are measured by a metric, while the quality function is monotone non-decreasing and submodular. To the best of our knowledge, Abbassi et al. [1] are the only work that covers diversity maximization over clustered data, but they focus on global dispersion, which is a sum of intra- and inter-cluster dispersion and may cause undesirable consequences in many applications. For example, in the above example, the presence of inter-cluster dispersion prevents one community from selecting hashtags similar to the one taken by the other community, and leads to incomplete summarization of communities.

Our first observation is that in the presence of overlapping clusters, the approximation algorithms of Borodin et al. (greedy and local search) have unbounded approximation ratios. We then propose algorithms with provable approximation guarantees. Our key finding is that it is beneficial to greedily add *pairs of objects* in the solution, instead of singleton objects. This idea is inspired by the work of Hassin et al. [11].

Additionally, we propose a modification of our algorithm, which improves scalability, and also enables additional flexibility. The modified algorithm allows to obtain better solutions in practice while maintaining the approximation guarantee in theory.

We summarize our contributions as follows.

- To the best of our knowledge, this is the first work to study the problem of intra-cluster diversity maximization in the presence of overlapping clusters.
- We propose a greedy algorithm with constant-factor approximation guarantee, while other popular heuristics are shown to have unbounded approximation ratio. We further extend the analysis to the case of monotone and submodular quality function, and under a partition matroid constraint.
- We show how to make our algorithm scalable, and on the way we obtain a modification that gives better solutions in practice while maintaining the approximation guarantee in theory.
- Our algorithm achieves excellent performance, compared to strong baselines in a mix of synthetic and real-world datasets.

The rest of the paper is organized as follows. We first discuss related work in Section 2. Then a formal definition for the problem is formulated in Section 3, followed by theoretical analysis in Sections 4 and 5, with and without a quality function, respectively. A modification is introduced in Section 6 to cope with large-scale data. We evaluate our algorithms on both synthetic and real-world datasets in Section 7. Finally, we conclude in Section 8.

For lack of space, the proofs of most of our technical results, together with the case of partition matroid constraint, and some additional experiments are given in the full version of this paper [23].

2 Related work

Maximum dispersion was first studied by Ravi et al. [17]. Gollapudi and Sharma [8] incorporate a quality objective into the dispersion framework. Borodin et al. [5] extend the dispersion problem to submodular quality functions. Abbassi et al. [1] study global dispersion under a partition matroid constraint over clustered data, whereas we study a sum of intra-cluster dispersion.

Most of the existing work on maximum dispersion, including the approaches mentioned above, adopts a greedy strategy of adding iteratively a single element or local-search strategy. One exception is the work of Cevallos et al. [14], who apply convex quadratic programming with a negative-type distance function. Bhaskara et al. [3] solve a similar SUM-MIN problem using an Linear Programming relaxation. Our algorithm is inspired by the work of Hassin et al. [11], who utilize maximum matching for multi-subset selections and a more efficient greedy algorithm on disjoint edges for the case of a single subset.

In web search, the importance of providing users a list of documents that are not only relevant to their queries but also diverse has been noticed early on. Carbonell et al. [6] propose a greedy algorithm with respect to *Maximal Marginal Relevance* (MMR) to reduce the redundancy among returned documents. Zhai et al. [21] invent two novel distances for documents, and evaluate diversity using coverage on subtopics. Agrawal et al. [2] incorporate an existing taxonomy of topics and develop a probabilistic model whose goal is to cover every intention of a query, where the probabilities of documents covering an intention are considered independent, determined by their relevance to the query and their topic distributions.

Zhang et al. [22] form an affinity graph of documents, where the most informative document given by PageRank is selected at each iteration, and then its neighbors are immediately removed to avoid redundancy. Radlinski et al. [16] assume that users with different interpretations of the same query click diverse documents, leading to the idea of optimizing a ranking of documents using only user clickthrough data such that the probability of users not clicking any document on the list is minimized. Compared to these approaches, diversification provides a solid theoretical framework, which can be applied to many different settings, once an appropriate object representation and distance metric are defined.

Other related work includes the team-formation problem [13], where it is asked to select team members who bear a small communication cost, formulated as a Steiner tree or diameter. Esfandiari et al. [7] have a different focus, partitioning students into different groups to encourage peer learning, so as to optimize some special affinity structure such as a star in each group.

The scalability of dispersion algorithms becomes an essential concern as the data size increases. Distributed algorithms have been invented to tackle the computational complexity for dispersion problems [20]. However, little work focuses on the scalability of dispersion algorithms in a clustering setting.

3 Notation and problem definition

In this section we formally define the *intra-cluster dispersion* problem, which is the focus of our paper.

We start with a *ground set* $U = \{u_1, \dots, u_n\}$, and a *distance metric* $d : U \times U \rightarrow \mathbb{R}_+$ defined over U . In some cases a *quality function* $Q : 2^U \rightarrow \mathbb{R}$ is available, which evaluates the quality of subsets of U according to some desirable property.

In addition we assume that we are given as input a clustering $\mathcal{C} = \{C_1, \dots, C_k\}$ over the ground set U , which may contain overlapping clusters. Finally, we are given cardinality upper bounds b_1, \dots, b_k , one for each cluster of \mathcal{C} , i.e., b_j is a cardinality upper bound for cluster C_j , for $j = 1, \dots, k$.

The objective of the intra-cluster dispersion problem is to pick pairwise disjoint sets $S = \{S_1, \dots, S_k\}$, with $S_j \subseteq C_j$ and $|S_j| \leq b_j$, for $j = 1, \dots, k$, so as to maximize the dispersion function

$$(3.1) \quad D(S) = \sum_{j=1}^k d(S_j),$$

where $d(S_j) = \sum_{u,v \in S_j} d(u,v)$ is the sum of all pairwise distances of the elements S_j picked from cluster C_j .

When a quality function Q is available, it can be incorporated in the intra-cluster dispersion problem by asking to maximize the objective function

$$(3.2) \quad D_Q(S) = Q(S) + \lambda D(S),$$

where λ is a user-defined parameter that is used to provide a trade-off between quality and dispersion. Equation (3.2) is often used to formulate the *diversification problems* [8], where we typically want to select high-quality sets (as measured by Q) which are also diverse (as measured by D).

We refer to the problems of maximizing Equations (3.1) and (3.2) by k -MAXSUM and k -MAXSUM-Q, respectively. In the case of a single cluster, we refer to them as 1-MAXSUM and 1-MAXSUM-Q, respectively.

An interesting generalization of the intra-cluster dispersion problem is when in addition to the clustering $\mathcal{C} = \{C_1, \dots, C_k\}$ we are also given a partition $\mathcal{P} = \{P_1, \dots, P_m\}$, the goal remains unchanged with an additional requirement of selecting at most one element from each set in the partition \mathcal{P} . Without loss of generality, we can assume that each element forms its own partition, resulting in the constraint of the sets in S being disjoint. However, our methods generalize to a general partition matroid constraint by removing a whole set of the partition when we select a single element from it. A detailed elaboration can be found in Section S.2.

Finally, recall that a set function $Q : 2^U \rightarrow \mathbb{R}$ is *monotone non-decreasing* if $Q(B) \geq Q(A)$, for all $A \subseteq B \subseteq U$. The function Q is *submodular* if it satisfies the “*diminishing returns*” property, that is, for all $A, B \subseteq U$ with $A \subseteq B$ and for all $u \in U \setminus B$, it holds that $Q(A \cup \{u\}) - Q(A) \geq Q(B \cup \{u\}) - Q(B)$. A set function Q is *supermodular* if $-Q$ is submodular, and it is *modular* if it is both submodular and supermodular.

Complexity. The k -MAXSUM and k -MAXSUM-Q problems studied in this paper extend the *max-sum diversification* problem [5], which aims at maximizing the objective (3.2) in a single cluster. The max-sum diversification is proven to be **NP**-hard [9]. As a consequence, the problems k -MAXSUM and k -MAXSUM-Q are **NP**-hard, as they generalize the max-sum diversification problem in the case that clusters are available. Thus, our objective is to develop approximation algorithms with provable guarantees.

4 Maximizing dispersion

We start our discussion about limitations of existing algorithms for the k -MAXSUM and k -MAXSUM-Q problems (i.e., maximizing Equation (3.1) and Equation (3.2), respectively), and proceed to the proposed algorithm for the k -MAXSUM problem. The case of k -MAXSUM-Q is discussed in Section 5.

For the 1-MAXSUM-Q problem, which is a special case of the k -MAXSUM-Q problem when no clustering is available, the *greedy algorithm* is a natural choice. The greedy adds elements into S iteratively, each time selecting the element that yields the maximal gain in a slightly-modified version of the objective function (3.2). Borodin et al. [5] showed that if the quality function Q is submodular, the greedy algorithm achieves a factor-2 approximation guarantee.

A natural extension of the greedy algorithm for the k -MAXSUM-Q problem is the following: iterate over the clusters C_1, \dots, C_k , and for each cluster C_j , form the set S_j by iteratively adding b_j elements into S_j from C_j ; each time add the element that yields the maximal

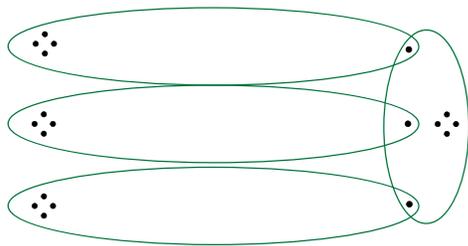


Figure 1: A bad instance for algorithms GREEDY and LS-INTRA for the intra-cluster dispersion problems (k -MAXSUM and k -MAXSUM-Q).

gain for the objective function (3.2). We refer to this simple algorithm as GREEDY.

Another powerful technique in combinatorial optimization is *local search*: start by a feasible solution and perform incremental changes until the objective cannot be improved. For the k -MAXSUM-Q problem, an intuitive way to apply the local-search strategy is by *element swapping*, which is finding a cluster for which it is possible to swap an element in the current solution and improve the objective. We call this algorithm LS-INTRA. When the objective is replaced with global dispersion, we call it LS-GLOBAL.

Our first result is to show that for the k -MAXSUM and k -MAXSUM-Q problems the algorithms GREEDY and LS-INTRA do not have a bounded approximation factor. This is in contrast to the approximation guarantee of GREEDY for the 1-MAXSUM problem.

LEMMA 1. *The approximation factor of algorithms GREEDY and LS-INTRA is unbounded.*

Proof. We prove this argument by the example shown in Figure 1, which shows 4 overlapping clusters. The horizontal distance of the three clusters on the left can be arbitrarily large. If GREEDY considers the cluster on the right before any of the clusters of the left, the approximation ratio can be arbitrarily large. For the LS-INTRA algorithm, a selection containing any overlapping element can be a feasible initial selection for the right cluster, leading again to an arbitrarily large approximation ratio. \square

From the proof of Lemma 1, it is worth noting that the performance of the two approaches, GREEDY and LS-INTRA, depends on the order of processing the data clusters, and the initial solution, respectively. Notice that the dependence of GREEDY on the order of processing the clusters is unavoidable, as the first element of each cluster that gets added to the solution contributes nothing to the objective function.

Proposed method. We now introduce our proposed method for the k -MAXSUM problem. The main idea is to overcome the drawback of GREEDY and its dependence on cluster ordering. We can achieve this by finding a *pair of elements*, which both belong to the same cluster, and add them both to the solution. After adding a pair of elements to a set S_j , the corresponding budget b_j is reduced by 2. A desirable feature of our algorithm, is that at each step we can select the pair of elements with the maximal contribution to some intuitive measure; the order of processing clusters is now determined completely by the algorithm, and it is not arbitrary. Pseudocode of our algorithm, which we call GREEDYPAIRS, is shown in Algorithm 1.

ALGORITHM 1. GREEDYPAIRS

Input: Ground set U , clustering $\mathcal{C} = \{C_1, \dots, C_k\}$, upper bounds b_1, \dots, b_k .

Output: Intra-cluster dispersion set $S = \{S_1, \dots, S_k\}$.

```

1: for  $j = 1, \dots, k$  do
2:    $S_j \leftarrow \emptyset$ ;  $b'_j \leftarrow 2 \lfloor \frac{b_j}{2} \rfloor$ 
3:   while (exists  $j$  such that  $|S_j| < b'_j$ ) do
4:      $R \leftarrow \{j \in [k] \text{ s.t. } |S_j| < b'_j\}$ 
5:     Select a pair  $p = \{u, v\}$  over clusters in  $R$  whose
       both elements  $u$  and  $v$  belong in some cluster  $C_j$  and
        $(b_j - 1) d(p)$  is maximized
6:      $S_j \leftarrow S_j \cup \{u, v\}$ 
7:      $U \leftarrow U \setminus \{u, v\}$ 
8:   for  $j = 1, \dots, k$  do
9:     if ( $b_j$  is odd) then
10:      Add an arbitrary element  $v \in C_j$  into  $S_j$ 
11: Return  $S_1, \dots, S_k$ 

```

Our first result is to demonstrate the relative power of greedily adding pairs of elements, instead of individual elements. In particular, in contrast to Lemma 1 where we showed that GREEDY has an unbounded approximation ratio, GREEDYPAIRS offers a provable constant-factor approximation.

THEOREM 2. *Algorithm GREEDYPAIRS is 6-approximation for k -MAXSUM problem.*

Proof. Let S^* be an optimal solution to an instance of the k -MAXSUM problem, and let S be the solution computed by GREEDYPAIRS. Let S_j^* be the subset of the optimal solution in cluster C_j , so that $S^* = \bigcup_{j=1}^k S_j^*$, and similarly $S = \bigcup_{j=1}^k S_j$. Let \bar{S}_j^* be the intersection of S_j^* with the elements of the solution S other than S_j , i.e., $\bar{S}_j^* = S_j^* \cap (S \setminus S_j)$.

If $\bar{S}_j^* = \emptyset$, then GREEDYPAIRS can be analyzed separately for each cluster C_j , as by Hassin et al. [11], and get $d(S_j) \geq \frac{1}{2} d(S_j^*)$, obtaining a 2-approximation.

If $\bar{S}_j^* \neq \emptyset$, then $d(S_j) \geq \frac{1}{2}d(S_j^* \setminus \bar{S}_j^*)$ still holds. Additionally, we have to bound the remaining part $d(S_j^*) - d(S_j^* \setminus \bar{S}_j^*)$. By the definition of distance d we have $d(S_j^*) - d(S_j^* \setminus \bar{S}_j^*) \leq 2 \sum_{v \in \bar{S}_j^*} d(v, S_j^*)$. Consider now an element $v \in \bar{S}_j^*$ and let $\{u, v\}$ be the selected pair of elements in S . Then $d(v, S_j^*) \leq (b_j - 1)d(\{u, v\})$ and further $d(v, S_j^*) \leq d(\{u, v\}) + \sum_{w \in (S_j \setminus \{u, v\})} d(u, w) + d(v, w)$. Summing over all elements $v \in \bar{S}_j^*$ and all clusters we get $\sum_{j=1}^k \sum_{v \in \bar{S}_j^*} d(v, S_j^*) \leq 2 \sum_{j=1}^k d(S_j) = 2D(S)$, because each element v is unique in S^* , and can be mapped to some pair in S . On the other hand, each pair of elements is mapped at most twice. We can conclude that

$$\begin{aligned} 6D(S) &\geq \sum_{j=1}^k d(S_j^* \setminus \bar{S}_j^*) + 2 \sum_{j=1}^k \sum_{v \in \bar{S}_j^*} d(v, S_j^*) \\ &\geq D(S^*). \end{aligned}$$

□

Furthermore, we can show that our analysis of GREEDYPAIRS is tight.

REMARK 3. ([23]) *The 6-approximation bound obtained in Theorem 2 for GREEDYPAIRS is tight, even if the distance function d is a metric.*

5 Dispersion with submodular quality function

We now turn our attention to the k -MAXSUM-Q problem, with the quality function Q being monotone non-decreasing and submodular. Notice that non-negative modular functions serve as a special case of such submodular functions, so we will devote our discussion only to the submodular case. In this case, the GREEDYPAIRS algorithm gives an equally strong guarantee with the following intuitive modification: In the i -th iteration we select the i -th pair of elements $\{u, v\}$ in some cluster j , among all unsaturated clusters, so as to maximize

$$(5.3) \quad \phi_{\{u,v\}}(S^{i-1}) = Q_{\{u,v\}}(S^{i-1}) + \lambda 2(b'_j - 1)d(u, v),$$

where $b'_j = 2\lceil b_j/2 \rceil$ and $Q_{\{u,v\}}(S^{i-1}) = Q(S^{i-1} \cup \{u, v\}) - Q(S^{i-1})$ is the marginal gain to the quality function Q from adding the pair $\{u, v\}$ into the current solution S^{i-1} . In our analysis we will make the mild assumption that the budgets b_j are even. We propose a simple heuristic to deal with the case that some budget b_j is odd, but the approximation guarantee becomes slightly weaker.

LEMMA 4. ([23]) *The GREEDYPAIRS algorithm with greedy rule (5.3) gives a 4-approximation (or $(4 \min\{\frac{b+1}{b-1}, 2\})$ -approximation if b is odd) for the 1-MAXSUM-Q problem.*

Notice that in the case of a single cluster, the GREEDY algorithm is superior giving a 2-approximation [5]. However, when many clusters are present, the approximation ratio of GREEDY becomes unbounded (Lemma 1), while GREEDYPAIRS maintains a constant-factor approximation guarantee, as we show next.

THEOREM 5. *The GREEDYPAIRS algorithm with greedy rule (5.3) gives a 6-approximation for the k -MAXSUM-Q problem, when all b_j are even.*

Proof. We denote the optimal solution as $S^* = \cup_{j=1}^k S_j^*$, and the solution found by GREEDYPAIRS as $S = \cup_{j=1}^k S_j$. Also, let $\{p_i\}$ be the set of pairs of elements selected by GREEDYPAIRS. As in the proof of Theorem 2 we define $\bar{S}_j^* = S_j^* \cap (S \setminus S_j)$. If $\bar{S}_j^* = \emptyset$ for every cluster C_j , then by Lemma 4, each set S_j is 4-approximation with respect to S_j^* . Hence S is 4-approximation.

Thus, we proceed to analyze the case that $\bar{S}_j^* \neq \emptyset$ for some j . For each set S_j^* we consider a *maximum matching* $\{p_{j i'}^*\}$.¹ We define a mapping π from $\{p_{j i'}^*\}$ to $\{p_i\}$, as follows:

- if $p_{j i'}^* \cap p_i \neq \emptyset$ for some i and cluster j is not saturated before the i -th iteration, then $\pi(p_{j i'}^*) = p_i$ where the smallest i is chosen;
- otherwise $p_{j i'}^*$ is mapped to some infrequently mapped pair p_i of S_j .

The key point of mapping π is that each optimal pair $p_{j i'}^*$ that falls into the second requirement has to be mapped to some selected pair p_i within cluster j , because pairs in other clusters can not be guaranteed to be better than $p_{j i'}^*$, as they may be selected simply because cluster j is saturated, i.e., a selected pair is greedy only among unsaturated clusters in its iteration. Notice that due to such a requirement, each pair p_i can be the image of at most three pairs $p_{j i'}^*$: two from the two pairs that intersect with p_i and one from some pair that does not intersect with p_i but is in the same cluster as p_i . As each pair p_i achieves the best marginal gain at the time it is selected, we obtain the following inequality: $\phi_{p_i}(S^{i-1}) \geq \phi_{p_{j i'}^*}(S^{i-1}) \geq \phi_{p_{j i'}^*}(S)$. The second inequality above is due to submodularity.

Assume that a total of B elements are selected from all clusters. We have $3 \sum_{i=1}^{B/2} \phi_{p_i}(S^{i-1}) \geq \sum_{j=1}^k \sum_{i'=1}^{b_j/2} \phi_{p_{j i'}^*}(S) = \sum_{j=1}^k \sum_{i'=1}^{b_j/2} Q(S \cup p_{j i'}^*) - Q(S) + \lambda 2(b_j - 1)d(p_{j i'}^*) \geq Q(S \cup S^*) - Q(S) + \sum_{j=1}^k \sum_{i'=1}^{b_j/2} \lambda 2(b_j - 1)d(p_{j i'}^*) \geq Q(S^*) - Q(S) + \sum_{j=1}^k \sum_{i'=1}^{b_j/2} \lambda 2(b_j - 1)d(p_{j i'}^*)$, where the last inequality uses the monotonicity of Q . Putting together that

¹These maximum matchings are only to be considered conceptually, for the purposes of the proof, not computed, as S^* is not known.

$3 \sum_{i=1}^{B/2} \phi_{p_i}(S^{i-1}) = 3Q(S) + \lambda \sum_{i=1}^{B/2} 6(b_{\xi(p_i)} - 1)d(p_i)$ where ξ maps a pair p_i to its cluster number, the inequalities $2(b_j - 1) \sum_{i'=1}^{b_j/2} d(p_{j i'}^*) \geq d(S_j^*)$ (by the virtue of $\{p_{j i'}^*\}$ being a maximum matching) and $d(S_j) \geq b_j \sum_{i=1}^{b_j/2} d(\tilde{p}_{ji})$ (by triangular inequality), where $\{\tilde{p}_{ji}\}$ can be an arbitrary set of disjoint pairs in S_j , we eventually have $6D_Q(S) \geq 4Q(S) + \lambda \sum_{i=1}^{B/2} 6(b_{\xi(p_i)} - 1)d(p_i) \geq Q(S^*) + \sum_{j=1}^k \sum_{i'=1}^{b_j/2} \lambda 2(b_j - 1)d(p_{j i'}^*) \geq D_Q(S^*)$. As modular functions is a subset of submodular functions, Remark 3 applies, and thus, our bound for GREEDYPAIRS is tight. \square

When some b_j is odd, we run the algorithm to select $2\lceil b_j/2 \rceil$ elements for each cluster, and then we remove an element in each cluster with an odd b_j that contributes the least with respect to some measure. A guarantee of $(6 \min\{\frac{b_m+1}{b_m-1}, 2\})$ -approximation can be proved, where b_m is the smallest odd b_j [23].

6 Scalable algorithms

As the data size increases, a vanilla implementation of our methods may be impractical. In this section we examine the scalability of our algorithm, and discuss extensions to improve its efficiency. We focus on scaling up the algorithm on a single machine; we are not after parallel or distributed implementations, which are of independent interest, but beyond the scope of this paper. We first inspect the time complexity of the three methods we have presented.

Method	Time complexity
GREEDY	$\mathcal{O}(kb_M^2 n)$
LS-INTRA	$\mathcal{O}(\frac{1}{\epsilon} kb_M^2 n)$
GREEDYPAIRS	$\mathcal{O}(kb_M n^2)$

Recall that n is the number of elements in U , k is the number of clusters in \mathcal{C} , and b_M is the maximum budget over all b_j . The accuracy parameter ϵ controls the termination of the LS-INTRA algorithm.

In a practical scenario we expect n to be much larger than k and b_M , thus, we are interested in optimizing time complexity with respect to n . From the above table we can see that GREEDY and LS-INTRA² are both linear in n , though the latter is slower in practice due to the need for convergence. The quadratic complexity of GREEDYPAIRS is due to the need of finding the *furthest pair* of elements, which we also refer to as *diameter*. To cope with the quadratic complexity, we borrow a well-known 2-approximation algorithm for finding the

²The optimal objective value is assumed to be a constant, and a swap is performed only when it increases the objective by more than ϵ

diameter in a metric space. Hereinafter, we use the terms “pair” and “edge” interchangeably.

LEMMA 6. *Picking an element x arbitrarily, and the element y that is the furthest from x , gives a 2-approximation to the problem of finding the diameter of a set of elements in a metric space.*

Proof. If optimal diameter is (u, v) , then $d(u, v) \leq d(u, x) + d(x, v) \leq 2d(x, y)$. \square

The idea behind Lemma 6 can be used not only to improve the running time but also to provide additional flexibility in adding elements to our solution, and obtaining solutions of higher quality in practice.

In particular, for the case of maximizing dispersion in a single cluster (problem 1-MAXSUM), if S^{i-1} is the set of currently-selected elements, in the next iteration we can select the first element x to be the one that maximizes $d(x, S^{i-1})$. For the second element y , instead of selecting the one that is the furthest from x (as Lemma 6 dictates), we have observed that, in practice, it is better to find an element y that balances its distance between x and S^{i-1} , via a hyper-parameter $\alpha \in [0, 1]$. In particular, if y^* is the furthest element from x , we find y so that $d(x, y) \geq \alpha d(x, y^*)$ and $d(y, S^{i-1})$ is maximized. We refer to this version of GREEDYPAIRS as GREEDYPAIRS $_{\alpha}$.

GREEDYPAIRS $_{\alpha=1}$ is the approximate version of GREEDYPAIRS, which selects the diameter using the 2-approximation method of Lemma 6. GREEDYPAIRS $_{\alpha=0}$ degenerates into GREEDY. For intermediate values of α , GREEDYPAIRS $_{\alpha}$ combines the best of both worlds: the theoretical properties of GREEDYPAIRS and the empirical effectiveness of GREEDY.

LEMMA 7. ([23]) *The GREEDYPAIRS $_{\alpha}$ algorithm with $0 < \alpha \leq 1$, is a $(4/\alpha)$ -approximation for the 1-MAXSUM problem.*

The case of maximizing dispersion for multiple clusters (problem k -MAXSUM) is similar. In each iteration, we select a furthest pair in each unsaturated cluster, and add to the solution the best pair out of these candidates.

THEOREM 8. ([23]) *The GREEDYPAIRS $_{\alpha}$ algorithm with $0 < \alpha \leq 1$, is a $(12/\alpha)$ -approximation for the k -MAXSUM problem.*

The situation is more complicated when a quality function Q is involved. As before we focus on monotone non-decreasing and submodular functions. First notice that evaluating Equation (5.3) has quadratic complexity, as we need to optimize quality as well, so we

need another fast heuristic. We use \hat{p} to represent a 2-approximation diameter, and p for an exact one. In this case, we select the first endpoint x that maximizes $Q_x(\hat{S}^{i-1})$ in the i -th iteration, and then search for the best pair \hat{p}_i^x with respect to $\phi_{\hat{p}_i^x}(\hat{S}^{i-1})$. Note that this pair is not necessarily the maximum-distance pair $\hat{p}_i^{\max,x}$ that includes x . As a result: $2\phi_{\hat{p}_i^x}(\hat{S}^{i-1}) \geq 2\phi_{\hat{p}_i^{\max,x}}(\hat{S}^{i-1}) \geq \phi_{p_i}(\hat{S}^{i-1})$. Similarly, we can still apply the same trick to look for a first endpoint x' of quality at least $\alpha Q_x(\hat{S}^{i-1})$ and a second endpoint y' that gains an additional value at least $\alpha(\phi_{\hat{p}_i^{x'}}(\hat{S}^{i-1}) - Q_{x'}(\hat{S}^{i-1}))$. Hence, $\frac{2}{\alpha}\phi_{(x',y')}(\hat{S}^{i-1}) \geq \frac{2}{\alpha}\phi_{\hat{p}_i^{\max,x'}}(\hat{S}^{i-1}) \geq \phi_{p_i}(\hat{S}^{i-1})$. We can show the following.

LEMMA 9. ([23]) *The GREEDYPAIRS $_{\alpha}$ algorithm with $0 < \alpha \leq 1$ and a greedy rule in Equation (5.3) is $(6/\alpha)$ -approximation for the 1-MAXSUM-Q problem when b is even.*

With multiple clusters, the selection scheme is the same as its modular counterpart, but with a different objective, that is, we select a best pair as in Lemma 9 within each unsaturated cluster and pick the final best one among them.

THEOREM 10. ([23]) *The GREEDYPAIRS $_{\alpha}$ algorithm with $0 < \alpha \leq 1$ and a greedy rule in Equation (5.3) is $(12/\alpha)$ -approximation for the k -MAXSUM-Q problem when b_j is all even.*

When some b_j is odd, we can apply the same technique as in Section 5, and the approximation ratio increases only by a multiplicative factor of $\min\{\frac{b_m+1}{b_m-1}, 2\}$, where b_m is the smallest odd b_j .

7 Experiments

We evaluate our methods using six datasets, two synthetic, and four real-world data, to examine their ability to obtain a selection of both good quality and diversity. We use two of the real-world datasets to evaluate dispersion alone, and the other two to evaluate dispersion combined with a submodular quality function. All baselines are used in the comparison with small datasets, and only the scalable ones are applied to large datasets.

The list of all the algorithms we evaluated, and their abbreviations, is the following: GREEDYPAIRS $_{\alpha}$ (GP) — the proposed algorithm; GREEDY (GV), LS-GLOBAL (LSG), LS-INTRA (LSI), MAXCOVER (MC), and RANDOM (RN). Unlike our algorithm GREEDYPAIRS $_{\alpha}$, all baselines optimize the objective locally, so we will run them multiple times with input clusters in different order. We run our algorithm with several different α as well. min/avg/max of multiple runs are reported. We also investigate the sensitivity of GREEDYPAIRS $_{\alpha}$ to its hyperparameter α .

The implementation of our algorithm and all baselines is available online.³

Results on synthetic datasets. The first dataset is a set of random vectors, each assigned to a fixed number of clusters randomly. The second one tries to mimic structure of coherent clusters in real data by creating Normally-distributed random vectors around predefined center vectors, one for each cluster (aka prototypes). In order to generate overlapping clusters, we also assign a vector to its closest prototype.

We first investigate sensitivity of GREEDYPAIRS $_{\alpha}$ to its hyperparameter. In both of the synthetic datasets, random and prototype, a larger α , i.e., trying to find a diameter almost as good as the approximated one, is likely to have better performance than smaller ones. Since our algorithm GREEDYPAIRS $_{\alpha}$ draws closer to GREEDY with a smaller α , this phenomenon also reflects the effectiveness of GREEDYPAIRS $_{\alpha}$. For more detailed results we refer to the Supplementary material, Table S.1.

In the small synthetic datasets shown in Table 1, GREEDYPAIRS $_{\alpha}$ outperforms all baselines in the random dataset, and unsurprisingly, RANDOM has the worst performance. LS-GLOBAL also performs poorly as it is optimizing a different objective, global dispersion. More interesting results are observed in the prototype dataset. When the number of selected elements is small, GREEDYPAIRS $_{\alpha}$ surpasses all others. However, when the number of selected points becomes extremely large, a random initialization for LS-INTRA and LS-GLOBAL obtains the best result as it does not favor any elements, while other algorithms have a preference for overlapping elements, causing some clusters to not be able to select enough elements. Interestingly, RANDOM performs equally well in this case, indicating the lack of structure in these problem instances.

It is worth noting that GREEDYPAIRS $_{\alpha}$ is quite robust to such extreme cases as it select best pairs globally. The performance remains similar with data of a larger scale. Detailed results are deferred to the supplementary material, Table S.2. Similarly, validation on the linear-time scalability of GREEDYPAIRS $_{\alpha}$ is shown in the supplementary material, Figure S.1.

Results on real-world data: Topical document aggregation. In order to create overlapping topics, we employ LDA [12] on two document datasets, 2 797 documents from five topics in *20NewsGroups* dataset and 18 713 blogs, one for each blogger [18], and assign each document to the topics to which estimated probability is larger than a threshold. Note that Cosine distance is computed between Word2Vec vectors of doc-

³<https://github.com/Guangyi-Zhang/clustering-max-diversity>

Table 1: Performance for small synthetic datasets ($n = 1000$, 10 clusters) where dim stands for dimension.

Dataset	b	dim	GPmin	GPavg	GPmax	GVmin	GVavg	GVmax	LSImin	LSIavg	LSImax	LSG	RN
random	10	2	0.987	0.993	1.0	0.975	0.984	0.988	0.970	0.975	0.979	0.942	0.656
random	10	10	0.993	0.996	1.0	0.983	0.987	0.991	0.987	0.993	0.996	0.962	0.792
random	100	2	0.975	0.987	1.0	0.923	0.931	0.935	0.944	0.955	0.972	0.949	0.918
random	100	10	0.957	0.981	1.0	0.935	0.941	0.944	0.925	0.935	0.942	0.937	0.919
proto	10	2	0.975	0.989	1.000	0.963	0.978	0.984	0.950	0.956	0.964	0.579	0.405
proto	10	10	0.993	0.998	1.000	0.978	0.988	0.993	0.992	0.994	0.997	0.804	0.751
proto	100	2	0.957	0.972	0.979	0.798	0.810	0.826	0.876	0.934	1.000	1.000	1.000
proto	100	10	0.987	0.992	0.997	0.985	0.986	0.987	0.994	0.998	1.000	1.000	1.000

uments, which is a metric for normalized vectors. We observe that $\text{GREEDYPAIRS}_\alpha$ surpasses the strong baseline GREEDY in various data settings, showing a superior ability to select diverse documents. Due to space limitation, the results are shown in the supplementary material, Table S.3.

Results on real-world data: Descriptors for scholar communities and movie categories. Upon the scholar collaboration network provided by *Aminer* [19], we run K-clique community detection algorithm [15] to discover dense and overlapping scholar communities. In this dataset, each scholar is associated with a set of keywords, representing their main research interests. To be more specific, there are 544 communities formed by 6 529 scholars and 23 827 keywords. Our goal is to select a small number of keywords for each community in the hope of fully characterizing it, in a way that the keywords cover as many scholars in its community as possible and they each represent distinct concepts. In more detail, the quality of a set of keywords is defined by the cardinality of their coverage over all scholars, and Jaccard distance is applied between different keywords. The reason we quantify the quality by global coverage instead of intra-cluster coverage is that the quality function is not submodular with the latter formulation. A simple technique to avoid selecting a keyword that has large global coverage from a cluster on which it has small coverage is by preprocessing, i.e., removing a keyword from clusters it covers weakly.

Another dataset we adopt is *Movielens* [10], consisting 27 278 movies, each falling into several of 20 categories and described by many short comments from different users. We extract adjectives from comments to represent user emotions or opinions. Our goal is to describe each category by diverse user opinions. We formulate quality of an adjective as the number of movies it describes, and distance between two adjectives as their Cosine distance in Word2Vec embedding space.

In order to evaluate our methods both for quality and dispersion, we compare $\text{GREEDYPAIRS}_\alpha$ with

two baselines, GREEDY and MAXCOVER, the latter of which is specialized for quality. We vary two model parameters, the number of selected elements (b) and the bi-objectives trade-off (λ), with the increase of which, dispersion plays a more important role in the combined objective. Besides, due to the large number of clusters, we will not run GREEDY multiple times with input clusters of different order. Accordingly, we also fix the hyperparameter of $\text{GREEDYPAIRS}_\alpha$ at $\alpha = 0.95$.

The result is shown in Table 2, where $\text{GREEDYPAIRS}_\alpha$ consistently achieves the best performance with respect to a combined objective under different model parameters. Though it can be expected that $\text{GREEDYPAIRS}_\alpha$ will achieve a good compromise between these two objectives, the result turns out to be even better, as $\text{GREEDYPAIRS}_\alpha$ obtains the best quality and the best dispersion simultaneously until model parameters distort the balance between the two objectives too much. The reason for this is that, MAXCOVER fails to select enough elements in highly overlapping clusters, while $\text{GREEDYPAIRS}_\alpha$ manages to do so, in other words, maintaining a balance among different clusters.

A use case for scholar communities is shown in the supplementary material in Table S.4.

8 Conclusion

We propose a provable algorithm for intra-cluster dispersion and diversification problems, for which popular heuristics fail to achieve any guarantee. For the diversification version, the theoretical guarantee is achieved when the quality function is modular and submodular, thus, accommodating a wide range of real-life applications. Our algorithm selects pairs of elements greedily. To achieve a desirable linear-time complexity, the algorithm incorporates an approximate diameter computation, which also permits additional flexibility in choosing diverse pairs of endpoints. In the empirical study on both synthetic and real-world datasets, our algorithm exhibits a superior ability to maximize dispersion and

Table 2: Performance on clustering descriptors selection, where columns starting with Q means quality, one with D means dispersion and one with QD means a combination of both.

Dataset	b	λ	QDGP	QDGV	QDMC	DGP	DGV	DMC	QGP	QGV	QMC
scholar	6	1	22258.25	22053.14	14949.15	15734.25	15631.14	8452.15	6524.0	6422.0	6497.0
scholar	6	5	85159.28	85010.09	48757.75	15728.06	15717.02	8452.15	6519.0	6425.0	6497.0
scholar	6	10	163833.55	163942.39	91018.50	15731.06	15751.24	8452.15	6523.0	6430.0	6497.0
scholar	10	1	52925.29	52548.52	32461.70	46396.29	46038.52	25939.70	6529.0	6510.0	6522.0
scholar	10	5	238424.28	237782.79	136220.49	46379.06	46254.56	25939.70	6529.0	6510.0	6522.0
scholar	10	10	470291.14	469702.83	265918.98	46376.21	46319.38	25939.70	6529.0	6509.0	6522.0
movielens	10	1	10145.64	10017.05	9736.03	1546.64	1490.05	1201.03	8599.0	8527.0	8535.0
movielens	10	5	16399.05	16073.75	14540.16	1572.41	1510.15	1201.03	8537.0	8523.0	8535.0
movielens	10	10	24340.60	23799.07	20545.32	1584.56	1527.31	1201.03	8495.0	8526.0	8535.0
movielens	20	1	15429.97	15079.33	14566.02	6565.97	6254.33	5702.02	8864.0	8825.0	8864.0
movielens	20	5	41906.11	40759.44	37374.10	6614.82	6390.29	5702.02	8832.0	8808.0	8864.0
movielens	20	10	75123.33	72992.49	65884.19	6629.73	6420.25	5702.02	8826.0	8790.0	8864.0

quality while maintaining balance among clusters.

Potential future directions include extensions to other quality functions and other notions of dispersion.

References

- [1] Z. ABBASSI, V. S. MIRROKNI, AND M. THAKUR, *Diversity maximization under matroid constraints*, in KDD, 2013.
- [2] R. AGRAWAL, S. GOLLAPUDI, A. HALVERSON, AND S. IEONG, *Diversifying search results*, in WSDM, 2009.
- [3] A. BHASKARA, M. GHADIRI, V. MIRROKNI, AND O. SVENSSON, *Linear relaxations for finding diverse elements in metric spaces*, in NeurIPS, 2016.
- [4] B. BIRNBAUM AND K. J. GOLDMAN, *An improved analysis for a greedy remote-clique algorithm using factor-revealing lps*, Algorithmica, (2009).
- [5] A. BORODIN, H. C. LEE, AND Y. YE, *Max-sum diversification, monotone submodular functions and dynamic updates*, in PODS, 2012.
- [6] J. G. CARBONELL AND J. GOLDSTEIN, *The use of mmr, diversity-based reranking for reordering documents and producing summaries.*, in SIGIR, 1998.
- [7] M. ESFANDIARI, D. WEI, S. AMER-YAHIA, AND S. BASU ROY, *Optimizing peer learning in online groups with affinities*, in KDD, 2019.
- [8] S. GOLLAPUDI AND A. SHARMA, *An axiomatic approach for result diversification*, in WWW, 2009.
- [9] P. HANSEN AND I. D. MOON, *Dispersing facilities on a network*, Cahiers du GERAD, (1995).
- [10] F. M. HARPER AND J. A. KONSTAN, *The movielens datasets: History and context*, Acm TIIS, (2016).
- [11] R. HASSIN, S. RUBINSTEIN, AND A. TAMIR, *Approximation algorithms for maximum dispersion*, Operations research letters, (1997).
- [12] M. HOFFMAN, F. R. BACH, AND D. M. BLEI, *Online learning for latent dirichlet allocation*, in NeurIPS, 2010.
- [13] T. LAPPAS, K. LIU, AND E. TERZI, *Finding a team of experts in social networks*, in KDD, 2009.
- [14] C. MANZANO AND A. BOLÍVAR, *Approximation algorithms for geometric dispersion*, tech. rep., 2016.
- [15] G. PALLA, I. DERÉNYI, I. FARKAS, AND T. VICSEK, *Uncovering the overlapping community structure of complex networks in nature and society*, nature, (2005).
- [16] F. RADLINSKI, R. KLEINBERG, AND T. JOACHIMS, *Learning diverse rankings with multi-armed bandits*, in ICML, 2008.
- [17] S. S. RAVI, D. J. ROSENKRANTZ, AND G. K. TAYI, *Heuristic and special case algorithms for dispersion problems*, Operations Research, (1994).
- [18] J. SCHLER, M. KOPPEL, S. ARGAMON, AND J. W. PENNEBAKER, *Effects of age and gender on blogging.*, in AAAI spring symposium: Computational approaches to analyzing weblogs, 2006.
- [19] J. TANG, J. ZHANG, L. YAO, J. LI, L. ZHANG, AND Z. SU, *Arnetminer: extraction and mining of academic social networks*, in KDD, 2008.
- [20] S. A. ZADEH, M. GHADIRI, V. MIRROKNI, AND M. ZADIMOGHADDAM, *Scalable feature selection via distributed diversity maximization*, in AAAI, 2017.
- [21] C. ZHAI, W. W. COHEN, AND J. LAFFERTY, *Beyond independent relevance: methods and evaluation metrics for subtopic retrieval*, in SIGIR, 2015.
- [22] B. ZHANG, H. LI, Y. LIU, L. JI, W. XI, W. FAN, Z. CHEN, AND W.-Y. MA, *Improving web search results using affinity graph*, in SIGIR, 2005.
- [23] G. ZHANG AND A. GIONIS, *Maximizing diversity over clustered data*, arXiv preprint arXiv:2001.03050, (2020).