

Diversification and Refinement in Collaborative Filtering Recommender (Full Version)

Rubi Boim, Tova Milo, Slava Novgorodov

School of Computer Science

Tel-Aviv University

{boim,milo,slavanov}@post.tau.ac.il

Abstract—This paper considers a popular class of recommender systems that are based on Collaborative Filtering (CF) and proposes a novel technique for diversifying the recommendations that they give to users. Items are clustered based on a unique notion of *priority-medoids* that provides a natural balance between the need to present highly ranked items vs. highly diverse ones. Our solution estimates items diversity by comparing the rankings that different users gave to the items, thereby enabling diversification even in common scenarios where no semantic information on the items is available. It also provides a natural zoom-in mechanism to focus on items (clusters) of interest and recommending diversified similar items. We present **DiRec**, a plug-in that implements the above concepts and allows CF Recommender systems to diversify their recommendations. We illustrate the operation of **DiRec** in the context of a movie recommendation system and present a thorough experimental study that demonstrates the effectiveness of our recommendation diversification technique and its superiority over previous solutions.

I. INTRODUCTION

Online shopping has grown rapidly over the past few years. Besides the convenience of shopping directly from one's home, an important advantage of e-commerce is the great variety of items that online stores offer. However, with such a large number of items, it becomes harder for vendors to determine which items are more relevant for a given user and, given the limited size of the screen, which of these possibly relevant items should be presented first.

Much research has been devoted recently to the development of *Recommender systems*[1]. These systems predict the rating (e.g., a grade on a scale of 1 to 5) that a user would assign to an unseen item, and consider items with a high predicted rating to be relevant. But, which of these highly rated items should be presented first to the user? A naive solution would be to simply sort the items by their estimated rating and present the top-k that fit onto the screen. This however may result in an over-specialized items list. For example, suppose that a user is interested in movie recommendations. Assume that only 5 movies may fit onto the screen and that the top-5 ranked movies, for this user, all happen to be Star Wars sequels. While the given user may indeed like this series, a more *diverse* and wider view of the highly ranked movies may be desirable. For instance one that includes a Star Wars movie, but also other movies like Star Trek or E.T., with the access to more Star Wars sequels enabled via a “more of that” zoom-in button.

This papers aims to provide precisely such diversification and zoom-in facilities. Specifically, we focus here on a popular class of recommender systems that is based of Collaborative Filtering (CF), in which user ratings to items based on previous ratings of (similar) items by (similar) users[2]. A first question that needs to be addressed when designing such a diversification mechanism is how to measure the similarity/diversity of two given items. Previous proposals are often based on the assumption that some semantic information on items (e.g. the genre of the movie, the director, the actors) is given. CF recommender systems, however, typically *do not carry such semantic information* [1]. But even if they had, a problem is that it is not always clear how to define item diversity based on a given semantic information [3]. For example, some movies of the same director/leading actor may indeed be similar, whereas others may not. To overcome this difficulty, we follow the CF approach [2], [4] and instead of relying on semantic information, determine items similarity (and correspondingly diversity) based solely on ratings that previous users gave to the items. Intuitively, each item here is viewed as a vector of ratings, with vector distance (measured, e.g., by cosine, L^i distance, or Pearson correlation coefficient) used as measure of similarity/diversity.

A second important challenge is the need to balance, when choosing items, between two possibly conflicting objectives: presenting highest ranked items vs. choosing highly diverse ones. Some previous works attempted to resolve this by assigning a weight to each objective and selecting an items set that maximizes the weighted sum[5]; others used thresholds to bound the allowed similarity between items and the drop in rank [6]. But the difficult question always is *which weights or thresholds to choose?*. Indeed, a manual tuning of weights/thresholds (e.g. by experimentation) for a given data set is not only time consuming but is also no longer effective when the data changes [5]. To solve this problem we propose here a novel approach that *avoids the use of weights/thresholds altogether*. We introduce the notion of *priority-medoids*, an adaptation of the classical notion of *medoids*[7] to a context where items have priorities (ratings). *Priority-medoids* (to be defined formally in the sequel) allow for natural clustering of items and the selection of cluster representatives that balance rank and diversity. The clustering further allows the realization of an intuitive “zoom-in” mechanism, where users can focus on specific items on the screen and view similar

recommended items. Priority-medoids sub-clustering is then used, recursively, to diversify their presentation (and to allow further zooming-in).

To best of our knowledge, the only other previous algorithm without weights/threshold is Algorithm Greedy of [6], which does not support zoom-in. The tradeoff between ranking and diversity is hard-coded in the algorithm and no declarative notion of optimality is given. As always, an advantage of a declarative definition is that it is not tied to a particular algorithm and thus allows for formal analysis and optimization. While we show that identifying the optimal priority-medoids is NP-hard, we present an efficient (ptime) heuristic based on *priority cover-trees*, a particular sub-class of cover-trees [8] that proves to be extremely effective in this context. Our experiments show that the representatives chosen by our algorithms, with no need for weights tuning whatsoever, are as good and sometimes even superior to those obtained even with optimally-tuned weights of previous algorithms. (The comparison measures and the experiments are detailed in Section V). We further present an optimization technique that exploits the properties of our algorithm for an efficient realization of the above mentioned “zoom-in” mechanism.

DiRec : To demonstrate the effectiveness of our approach, we implemented the above solution in the DiRec prototype system. DiRec is designed as a plug-in that can be deployed on CF-based recommender systems, by implementing a simple API, and was demonstrated in [9]. [9] provides a high-level description of the system, while the current paper presents the underlying model and algorithms. Our experimental study examines the operation of DiRec in the context of a movie recommendations system using real data from Netflix [10]. This data set provides only raw user ratings to movies (such as 1 to 5 “stars” given by individual users) and does not hold semantic information on the movies. The results thus illustrate the effectiveness of our recommendation diversification technique even in the absence of semantic information.

Related Work: We next give a brief summary of previous work, highlighting the contributions of this paper.

Modern recommender systems aim to generate a personalized set of recommendations to each user. The difficulty of gathering semantic information on items and user preferences has triggered the development of recommender systems that are based on Collaborative Filtering (CF) rather than on semantics [1]. Such systems are very common nowadays and are the focus of the present work.

As explained above, a recommendations list that consists of the items with top-k predicted ratings may suffer from over specialization. Indeed, [11] evaluated the diversity of top-k items generated by traditional CF algorithms and showed it to be fairly low. Several algorithms that attempt to diversify the recommendation were proposed in the literature (see [5] for a survey). They fall generally into two classes: *greedy heuristics*, where the recommendation list is constructed “one-by-one” by maximizing a given distance function at each step (e.g. [12], [3], [6]), and *interchange (Swap) heuristics*, where an initial

list is first constructed and then refined by a series of actions that forms the final list (e.g. [4], [6]). But common to most is the use of *predefined weights or thresholds* to determine the balances between ranking and diversity or to bound the allowed similarity between items and the drop in rank [12], [3], [4], [13]. While the selected weights/thresholds clearly affects the performance of the algorithms, their precise choice is left open in all the works we are aware of. Such use of weights/thresholds is problematic since their manual tuning (e.g. by experimentation) for a given data set is not only time consuming but is also no longer effective when the data changes [5]. As explained in the Introduction, our solution employs, instead, priority medoids, to declaratively capture the desired balance. This is in contrast to [6] where the tradeoff is hard-coded in the algorithm. It further has the advantage of allowing for a natural (and optimizable) zoom-in mechanism.

The importance of results diversification has been recognized also in the context of database queries [14], [15], [16], [17], [18]. For example, [14] proposes a notion of diversity over structured query results which are post-processed and organized in a decision tree to help users navigate them; [15] introduces a hierarchical notion of diversity in databases and develops efficient top-k processing algorithms; [16] proposes to diversify query results to avoid over personalization; [17] uses attributes content to group tuples in a meaningful way that allows for convenient data exploration. This line of works however relies heavily on the *structured data content*. But such structured (semantic) information is *not available* in CF Recommender systems. Our work alleviates this problem by adopting the CF approach and relying on CF (dis)similarity measures rather than semantic ones.

The most relevant to our work, although also targeted to structured databases, is [18], where the authors used the notion of classical medoids (approximated by classical cover-trees) to select representatives for the query results and to zoom in on similar answers. While our work was inspired by [18], a key difference is that [18] *completely ignores tuples rating/priority*. We will see that our use of *priority* medoids and, resp., *priority* cover-trees, over the classical ones, prove to be extremely effective and greatly improves that generated recommendations.

Contributions: The technical contributions of this paper can be summarized as follows:

- We introduce the novel notion of priority-medoids as a tool for selecting item representatives. Our approach naturally balances the rating and the diversity of the recommended items and is applicable even in the absence of semantic information (as often is the case in CF recommender systems).
- We show that finding optimal priority-medoids is NP-hard and provide an alternative effective heuristic based on priority cover-trees.
- We exploit the properties of our algorithm to design an efficient, incremental zoom-in mechanism that allows to focus on individual items, identify their neighborhood (similar) items and select appropriate representatives for

them.

- We discuss the implementation of the DiRec plug-in that implements the aforementioned algorithms and present an experimental study that demonstrates the superiority of our solution relative to previous algorithms as well as the efficiency of the algorithms.

The paper is organized as follows. Section II presents priority-medoids. Section III then explains how they are approximated and how item representatives are selected. Section IV describes the zoom-in mechanism. The system implementation and our experiments are described in Section V. Finally, we conclude in Section VI.

II. PRIORITY-MEDOIDS

We start by providing the needed background and notation for Collaborative Filtering (CF). We then consider the problem of balancing the ratings and the diversity of the recommended items and define priority-medoids.

A. Collaborative Filtering

Common CF algorithms are *item-based*, consisting of two main steps: (1) choosing for each item a *neighborhood* of similar items, and (2) predicting the rating that a user u will give to an item i using some aggregation function on the actual ratings, gave by u , to the items within the neighborhood of i [2]. Symmetric *user-based* variants also exist but are less frequently used [2]. A key ingredient in the algorithm is thus the estimation of similarity between two items. Intuitively, each item is viewed as a vector of ratings in a multi-dimensional space, where each dimension corresponds to a given user, (recording her rating of that item). The distance between item vectors is then used as a measure for the items similarity. In principle, any distance measure can be employed here (e.g. Cosine or L^i distance) but Pearson's correlation coefficient [19] seems to be the preferred choice in most major systems. The basic intuition behind Pearson's measure is to give a high similarity score for two items that tend to be rated the same by many users. In the reminder of this paper we use $rate(u, i)$ to denote the predicted rating of a user u to item i . When u is known from the context we omit it and simply write $rate(i)$. We use $dist(i, j)$ to denote the distance between items i and j . W.l.o.g. we assume below that distance values are in the range of $[0, 1]$. (When this is not the case one may naturally map the values to this range). The smaller the distance is, the more similar (and less diverse) are the items.

B. Balancing Rating and Distance

We can now describe our main problem: given a set I of items and a size k , where the former holds the most relevant items for a given user (as decided by the given CF recommender system) and the latter denotes the number of items fitted onto the screen, we need to choose the subset $I_k \subseteq I$ (of size k) that will be presented to the user. A naive solution simply selects the top- k items with the highest $rate()$ values, namely the ones who maximize the sum $\sum_{i \in I_k} rate(i)$. This however may result in an over-specialized subset, as described




Fig. 1. Priority-medoid vs. standard medoid

in the Introduction. Thus, we should consider the diversity of the items as well, which may be analogically measured by the sum $\sum_{i,j \in I_k} dist(i, j)$. As these are two opposing measures, one needs to provide a good balance between the two. Additionally, this subset should also provide a good coverage for the entire set I , which intuitively can be viewed as a classic clustering problem, where we need to minimize the distances between the items in I and the “center” of the cluster they belong to.

Our solution, as mentioned in the Introduction, is based on the notion of *priority-medoids*, an adaptation of the classical notion of medoids to this context. To explain this, let us first briefly (and informally) recall what standard medoids are. Consider a set I of items split into k disjoint subsets, referred to as clusters. The *medoid* of a given cluster (also called the cluster's representative) is an element in the cluster s.t. the sum of the distances from it to the other items in the cluster is minimal. Other variants that consider e.g. the average, min or max distance, also exist [7]. This sum is called the cluster's weight. The classical goal is to find a clustering that minimize the overall sum of cluster weights. Note that, given a set $I_k \subseteq I$ of k items in I , the minimal-weight clustering for which the I_k items serve as representatives (medoids), is one where each item $i \in I$ is associated (clustered) with the element in I_k that is closest to it. Thus to find the best clustering one essentially needs to identify the best I_k set.

In our context we are interested in representatives with high rating. Priority-medoids therefore add the requirement that the representatives are the ones having *highest rating* in their corresponding clusters. Thus, when considering a set I_k of priority-medoids, the clusters it forms are different than the ones formed when considering standard medoids.

More formally, consider a subset $I_k \subseteq I$ of size k of items, s.t. I_k contains, among others, an item having the highest rating in I . We will explain below why having such an item in I_k is important. For an element $i \in I$, we denote by $rep(i)$ the item within I_k satisfying the following two constraints:

- the rating of $rep(i)$ is greater or equal to that of i
- among all items in I_k satisfying the above, $rep(i)$ is the closest to i , namely there is no other $j \in I_k$ with $dist(i, j) < dist(i, rep(i))$.

The items with the same representative $rep(i)$ form a cluster, and thus I_k yields a clustering formation for the items of I . We refer to the items in I_k as the *priority-medoids* of their corresponding clusters. Note that the fact that the highest rated

element in I is a member of I_k guarantees that all elements in I indeed have a cluster to which they may belong.

Example 2.1: The example in Figure 1 illustrates the difference between the cluster formation in regular medoids and that of priority-medoids. Assume that $k = 2$ and that the set I_k consists of the two items i' and i'' . Also assume that the euclidian distance between items describes their similarity. The number in parenthesis, next to each item, describes its rating. On the right (resp. left) hand side we see the clusters formed when the items in I_k are treated as regular (resp. priority-) medoids. The arrows outgoing the elements $i_1 - i_4$ point to their cluster representative. On the right hand side (regular medoids) items i_2, i_3 and i_4 are clustered with i'' , as they are closer to it than to i' . i_1 is clustered with i' . In contrast, on the left hand side (priority-medoids), the clustering formation is different as the item ratings are now also being considered. Specifically, since the rating of i'' is lower than that of i_2 , i_2 is clustered with (and represented by) i' , that has higher rating, even though it is slightly further.

The quality of the obtained clustering (and thereby the quality of the set I_k of priority-medoids that yielded the clustering), is measured, as before, by the distance of the items to the corresponding cluster representatives, namely by $\sum_{i \in I} \text{dist}(i, \text{rep}(i))$. When I is known from the context, we use a simplified notation and denote this sum by $\text{weight}(I_k)$. As for standard medoids, the lower the weight, the better the clustering (and the set I_k of priority-medoids) is. We are thus interested in a set I_k with minimal $\text{weight}(I_k)$. However, we point out that there may be several sets with the same minimal weight, in which case we break the tie by choosing the one where rating values are lexicographically higher.

For example, assume $k = 3$ and we have two sets $I_3 = \{i_1, i_2, i_3\}$, $I'_3 = \{i'_1, i'_2, i'_3\}$, where $\text{weight}(I_3) = \text{weight}(I'_3)$. If the rating values of the three elements in I_3 (resp. I'_3), sorted in decreasing order are 5, 4, 1, (resp. 5, 3, 2) then we choose I_3 over I'_3 . (An alternative could be to prefer, e.g., item sets with higher average/sum of rating). Ties may still occur when distinct items have the same rating, in which case we break it arbitrarily.

We can show that identifying the best priority-medoids is NP-hard (proof is omitted). We thus use a heuristic, based on *priority cover-trees* - an adaptation of the classical cover-trees [8] to our context. We explain this next.

III. PRIORITY COVER-TREE

A cover-tree is a data structure originally designed to speed up a nearest neighbor search [8]. The use of cover-trees as a tool for selecting (regular) medoids was recently proposed in [18], in the context of diversification of query results. A key difference from the present work is that *item ratings were not taken into consideration*. We next show that a fairly simple modification to the algorithm of [18] allows to account for such ratings and thereby gradely improve the quality of the generated recommendations.

A conventional cover-tree can be thought of as a hierarchy of levels, where each node corresponds to a specific item, and




Fig. 2. Priority Cover-Tree

each level is a “cover” for the level beneath it. (The root is at level zero, its children at level one, and so on). Each node in the tree is associated with an item in I . An item can be associated with multiple nodes but can appear at most once in every level l . A conventional cover-tree obeys, for all levels, the first three invariants below. In our algorithms we use a special sub-class of these trees, which we call *priority cover-trees*, that further obey the *forth invariant*.

- 1) (Nesting) If a node is associated with an item i , then one of its children must also be associated with i .
- 2) (Separation) All nodes at level l are at least $\frac{1}{2^l}$ far from one another.
- 3) (Covering) Each node at level l is within distance $\frac{1}{2^l}$ to its children in level $l + 1$.
- 4) (Priority) Each node has a rating higher or equal to that of any of its children.

Example 3.1: To illustrate how a priority cover tree looks like, consider for example the items $i_1 - i_7$ on the right hand side of Figure 2. While in general we do not assume that the pairwise distances between items (the $\text{dist}(i, j)$ measure) forms a metric, the distances between the items in this example are presented in the euclidian space for the readers convenience. The number in the parentheses next to each item denotes its rating. The tree on the left hand side of the figure is an example for a priority cover-tree for these elements. To maintain invariant 4, the item with the highest priority i_6 is assigned to the root. Analogously, all other nodes are assigned to items with higher priorities than their descendants. For instance, in this specific tree i_4 is a descendant of i_5 . We note that without invariant 4, as e.g. in the case of standard cover-trees, an alternative tree where the roles of the two items is reversed would also be legal. Namely a tree where i_4 is now direct descendant of the root, and correspondingly i_5 serve as a descendant of i_4 .

We next explain the role that invariant 4 plays in our algorithms for constructing priority-cover trees and for selecting items representatives, based on the constructed tree.

A. Construction

Our construction algorithm for priority cover-trees resembles the one of the original cover-tree [8]. Nevertheless, we

next detail the algorithm and the specific changes as they are important for understanding our refinement (zoom-in) algorithm.

The construction consists of two main components: the first (Algorithm 1) starts by sorting the set I of items according to their rating, in a descending order. This will allow preserving the invariants of the priority cover-tree (see Theorem 3.3 below). The algorithm then constructs a new tree whose root is the item with the highest rating. It then iterates orderly over the rest of elements and adds them into the tree one by one, by calling the second component: the `InsertSingleItem` function (depicted in Algorithm 2). This function finds the first level into which the given item can be inserted. It works in a recursive fashion: it is given as input the item i to be inserted, the current level of tree l under which the item may potentially be inserted (initially level 0 - the level of the root), and the set Q_l of the nodes in this level - the node that are considered to serve as i 's parent. Specifically, it contains all the items whose $dist()$ measure with the added item is below $\frac{1}{2^l}$ (due to invariant 3). On lines 1-3 the algorithm checks if the separation (invariant 2) between the added item and the descendants of items in Q_l holds. The function $Ldist(Q, i)$ returns the item $i' \in Q$ whose $dist(i, i')$ value is the smallest among all members of Q . If the condition holds, the algorithm returns *true* and a previous instance of the recursion will add the item. On line 5 the algorithm advances to the next level $l + 1$ by choosing the item who is the most similar to i from all candidates that preserves invariant 3. Intuitively, this will set the closest item as the item's parent. This item is saved in the set Q_{l+1} , and on line 6 the call for the recursion is done. If the recursion returned true, and the condition holds, the algorithm chooses, on line 8, the parent (out of the legal candidates) and adds the new item underneath it, on line 9. We point however that because of line 5, the selection on line 8 is unique as the set Q_l holds a single item at any time.

We illustrate the operation of the algorithm with the following example.

Example 3.2: To continue with our running example, the priority cover-tree in Figure 2 is obtained by running the construction algorithm on the items in the figure. The algorithm starts by sorting the items by their ratings. Item i_6 is first set as the root of the tree, followed by the insertion of items i_3 and i_5 . Note that when two or more items share the same rating value, their relative order is arbitrary; in this example we used the lexicographic ordering of their corresponding item number (thus item i_3 was inserted prior to item i_5). Then, item i_7 is inserted and added at level 2, as its distance from with i_6 , $dist(i_6, i_7)$, is too small, and invariant 2 prevents it from residing at level 1. Note that it is placed underneath i_6 due to invariant 3. Analogously, items i_1 and i_4 are later added at level 3 underneath items i_2 and i_5 correspondingly. Both are not added at level 2 because their $dist(i_1, i_2)$ and $dist(i_4, i_5)$ values are too high (that is, they are too similar) for preserving invariant 2. We finally note that the tree presented in Figure 2 is "fully presented", where every node is shown. It is not necessary to repeat a node when it has no (different)

descendants (for example, i_3 , i_6 and i_7 in Figure 2).

Algorithm 1 Construct Priority Cover-tree

Require: Set of items I

- 1: Sort I by relevance
 - 2: $i_{root} = \text{pop first } I$
 - 3: set i_{root} as root tree
 - 4: **for all** $i \in I \setminus \{i_{root}\}$ **do**
 - 5: $InsertSingleItem(i, 0, \{i_{root}\})$
 - 6: **end for**
-

Algorithm 2 Insert Single Item

Require: Item i , Cover set Q_l , level l

- 1: $Q = \{Children(i') | i' \in Q_l\}$
 - 2: **if** $dist(Ldist(Q, i), i) > \frac{1}{2^l}$ **then**
 - 3: return *true* (parent found)
 - 4: **else**
 - 5: $Q_{l+1} = \{Ldist(Q, i)\}$
 - 6: $found = InsertSingleItem(i, Q_{l+1}, l + 1)$
 - 7: **if** $found == \text{true AND } dist(Ldist(Q_l, i), i) \leq \frac{1}{2^l}$ **then**
 - 8: $q = \text{the node in } Q_l$
 - 9: insert i into $children(q)$
 - 10: return *false* (finish)
 - 11: **else**
 - 12: return *false*
 - 13: **end if**
 - 14: **end if**
-

Comparison with the construction of regular cover-trees

The construction algorithm for priority cover-tree can be viewed as a refinement of the basic algorithm for (regular) cover-tree construction from [8]. Specifically, there are two particular actions which our prioritized variant takes:

- *Ordered Insertion.* In the conventional cover-tree construction algorithm, items are inserted in an arbitrary order. Instead, our algorithm first sorts the items w.r.t their ratings, then inserts them in descending order.
- *Tight Insertion.* In the original algorithm, at the point when the algorithm finally inserts the given item into the tree, it may have several candidate nodes that may act as the item's parent. As any node in this set is a 'legal' parent (in the sense that the invariants will be preserved), the original algorithm chooses among them arbitrary. In contrast, our refined algorithm always prefers the node with the smallest $dist()$ measure to the inserted items.

We note here that ordered insertion suffices to guarantee Invariant 4 (see Theorem 3.3). Tight insertion, on the other hand, is not mandatory for the correctness of the construction. However it will prove useful later, when considering recommendations refinement (zoom-in). We discuss this issue in more detail in Section IV.

Theorem 3.3: At any point of the construction, the tree built by Algorithm 1 preserves all four invariants of the priority cover-tree.

Proof: As our algorithm is a refinement of the regular cover-tree construction algorithm, that preserves the first three invariants, so does Algorithm 1. For the fourth invariant, we need to show that for any item i and any of its descended i' , the rating of the former is higher than the one of the latter. Indeed, as our tree is constructed in a top-down fashion, and i is higher in the tree than i' , we know that i was inserted prior to i' . As the order of the insertion respects the rating (in descending order), it follows that i is of higher rating. ■

B. Representative Selection

We next explain how the constructed priority cover-tree is used to select item representatives. We use below the following notations. Given a node n in the tree and an integer $l \geq 0$, we say that n is at level l in the tree, if the distance between n and the root of the tree (counted by the number of edges on the path between them) is l . We use C_l to denote the set of nodes at level l .

The pseudo-code of the algorithm is given in Algorithm 3 and we sketch below the key ideas. Consider a set I of items which the CF recommender determined as most relevant for a given user. As explained above, the k items that we want to present to users are the priority-medoids of I . To find such a subset, with low weight, we use the priority cover-tree of I as follows: Starting from the tree root, we search for the first level l within the tree that includes at least k nodes, namely where $|C_{l-1}| < k$ and $|C_l| \geq k$ (lines 1-7). We then choose our k representatives from within C_l (lines 8-14). Recall that due to the structure of the priority cover-tree the nodes in C_i are at least $\frac{1}{2^i}$ far from one another and have rating higher than their descendants. If C_l contains exactly k nodes, then we simply choose these nodes for representatives and we are done. Otherwise, we need to select a subset of size k . To select good representatives, we use the tree structure: Recall the *nesting* property of the tree, which implies that $C_{l-1} \subseteq C_l$. Moreover, because of the hierarchical structure of the tree, the nodes (items) within C_{l-1} have rating \geq than those in C_l and are pairwise further apart from each other than from the remaining nodes. Thus, we first select the nodes (items) in C_{l-1} , then add the remaining $k - |C_{l-1}|$ from within $C_l - C_{l-1}$. We consider in our implementation and experiments three alternative methods for choosing these additional elements:

- *Max-rating* - the elements having the maximal ratings,
- *Max-diversity* - the elements that are farthest from the previously chosen elements, and
- *Max-coverage* - the elements having the maximal number of descendants.

The next example illustrates the difference between these variants.

Example 3.4: Consider our running example of the priority cover-tree depicted in Figure 2. Assume that we look for three representatives ($k = 3$). The first level that contains at least three items is the level 1. All items in the above level (0) are “automatically” selected, which is only i_6 in the running example. The algorithm then needs to decide which two more items out of the three candidates in $C_1 - C_0$ will be selected

as well. When considering *Max-rating*, items i_3 and i_5 are selected as their rating (4) is higher than the one of item i_2 (3), resulting with the final set $\{i_3, i_5, i_6\}$. Alternative, when considering *Max-diversity*, items i_2 and i_3 are selected as their *dist()* measure with i_6 is higher than the ones of item i_5 , resulting with the final set $\{i_2, i_3, i_6\}$. Finally, when considering *Max-Coverage*, items i_2 and i_5 are selected as the number of their descendants is higher than the ones of i_3 , resulting with the final set $\{i_2, i_5, i_6\}$. We note that although in this example each of the three measures yields a different representative sets, this is not always the case in general and they naturally may agree.

Note that, during the tree construction, we can easily record the number of descendants each node has, and use this information for the implementation of the Max-coverage variant. Thus, all three options are equivalent in terms of computational effort. They may differ however, as we shall see later, in the quality of the generated representatives set.

Algorithm 3 Priority-medoids Selection

Require: $k, treeheight, |C_{treeheight}| \geq k$

```

1:  $l = 0$ 
2: while  $l < treeheight$  do
3:   if  $|C_l| \geq k$  then
4:     break
5:   end if
6:    $l++$ 
7: end while
8:  $PriorityMedoids = C_{l-1}$ 
9:  $C_{rest} = C_l - C_{l-1}$ 
10: sort (descending)  $C_{rest}$  by rating, diversity or coverage
11: while  $|PriorityMedoids| < k$  do
12:    $i = \text{pop first } C_{rest}$ 
13:    $PriorityMedoids = PriorityMedoids \cup \{i\}$ 
14: end while
15: return  $PriorityMedoids$ 

```

IV. RECOMMENDATION REFINEMENT

Alongside each item (representative) i that is presented on the screen, DiRec offers a “more of that” zoom-in button that allows the user to view further related items.

One possible approach for identifying such items in I is to select the ones whose distance from i (as measured by the *dist()* function) is below a certain threshold. However, as it is never obvious which threshold should one choose, we take, instead, an alternative approach based on the following intuition: when a user clicks on a specific item i , she not only signals her interest in item’s i family but also signals that i interests her more than the other $k - 1$ presented representatives. Recalling the clustering formation generated by the presented representatives (priority-medoids), we determine the items in the cluster represented by i as relevant, denoting this set by *relevant*(i). Formally,

$$relevant(i) = \{i' | rep(i') = i \wedge i' \in I\}$$

Here again, the set $relevant(i)$ often contains more items than could fit on the screen and a subset needs to be chosen, to be presented to the user. A straightforward approach to choose k representatives for $relevant(i)$ is to build a priority cover-tree for the set, then use it to select representatives, as described above. We refer this as the *naive* approach. Rather than doing so, we employ instead an optimized, significantly more efficient, algorithm that is based on the following observation: In the priority cover-tree previously constructed for I , most of the elements in $relevant(i)$ already appear in subtree rooted at i . Indeed, the *tight Insertion* used in to the construction algorithm (Section III-A) was employed precisely to increase the number of such elements. We thus use this subtree as a basis for the construction of the priority cover-tree of $relevant(i)$. Note however that the subtree may not include all the members of $relevant(i)$ and also it might include some redundant elements that are not members of $relevant(i)$. We show this on the following example:

Example 4.1: We continue with our running example of the priority cover-tree depict in Figure 2, and assume that $\{i_2, i_3, i_6\}$ is the set of representatives. (Note that this scenario was presented at the previous section when $k = 3$ and *Max-diversity* was selected). The set $relevant(i_6)$ consists of the item i_7 as well as i_4 and i_5 , as the rating of item i_6 is higher than their rating and their distance to it is smaller than their distance to the (non selected) items i_2 and i_3 . But the subtree underneath item i_6 consists only of item i_7 (and trivially item i_6), leaving items i_4 and i_5 outside. Note that this scenario may occur when $|C_i| > k$, as some items here do not reside underneath the subtree of any representative.

We also note that some elements that belong to the subtree rooted at the selected representative i may not be members of $relevant(i)$. This scenario can occur if the pairwise distances between the items (the $dist()$ measure) do not form a metric. We thus need to prune redundant elements and add missing ones into the subtree. This is performed by Algorithm 4.

The algorithm starts by pruning redundant elements from the tree (lines 1-5). As such elements may have attached descendants, the remove function removes the entire subtree underneath these elements (line 3). Note that their descendants may or may not be a member of $relevant(i)$. The algorithm then initialize a “to be inserted” priority queue (line 6) and adds all the missing elements to it (lines 7-11). The priority queue is sorted by the item ratings from high to low. The algorithm then iterates over these element and inserts (in order) the missing items into the subtree (lines 13-14). The insertion follows the usual procedure described in the previous section, except that special attention is payed to cases where an element with high rating is to be inserted below one with lower rating. Note that such special treatment is essential if one wants to preserve invariant 4 of the priority cover-tree. In such cases we first prune out the “problematic” subtree (lines 15-16) - that is, the one rooted at the node with the lower rating - and adds its items to the “to be inserted” items priority queue (lines 17-18). The process is guaranteed to terminate since the queue is sorted: when an element is inserted, only items with

lower rating are added to the queue, and the number of such elements is finite. The updated subtree is then fed as input to Algorithm 3 for selecting the item representatives (line 22).

Algorithm 4 Refinement - Optimized approach

Require: $i, subtree$

```

1: for all  $i' \in subtree$  do
2:   if not  $i' \in relevant(i)$  then
3:      $subtree.removeSubtree(i')$ 
4:   end if
5: end for
6:  $init\ priorityQueue$ 
7: for all  $i' \in relevant(i)$  do
8:   if not  $i' \in subtree$  then
9:      $priorityQueue.add(i')$ 
10:  end if
11: end for
12: while not  $priorityQueue.empty$  do
13:    $i' = priorityQueue.pop$ 
14:    $subtree.add(i')$ 
15:   if  $subtree.violated$  then
16:      $subtree' = subtree.removeViolated()$ 
17:     for all  $i' \in subtree'$  do
18:        $priorityQueue.add(i')$ 
19:     end for
20:   end if
21: end while
22: return Algorithm 3 with respect to  $subtree$ 

```

While the worst case complexity of the algorithm equals to that of the naive, full tree construction, in practice many of the items indeed appear in the subtree and only few conflicts are encountered. This yields significant performance improvement, as demonstrated in our experiments in Section V-D.

V. IMPLEMENTATION AND EXPERIMENTS

We have implemented the above algorithms in DiRec , a plug-in that allows CF Recommender systems to diversify the recommendations that they present to users. We first present the system architecture and the settings used in our experiments. Then we describe the experimental results. The experiments study both the representative selection algorithms and the refinement process.

System Architecture: DiRec is implemented in Java and PHP, and designed to be deployed alongside any existing CF-based recommender system. Figure 3 illustrates the system architecture, divided into operating modules. When a user logs in, her id is passed to the *CF Recommender System* module which generates a customized list of recommended items, along with their ratings, and computes pair-wise item similarities. This information is then passed to the *Priority Cover-Tree* module, which constructs the corresponding tree. The *Priority-Medoid Approximation* module receives this tree and selects the representative items. These are presented to the user via an intuitive *User Interface* (UI). When a user

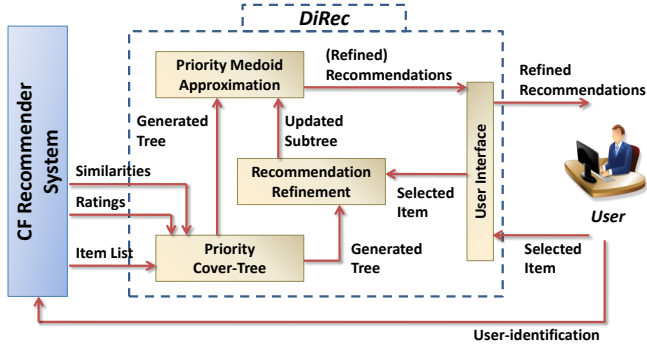


Fig. 3. DiRec architecture

clicks on a “more of that” button, the corresponding item id is passed to the *Recommendation Refinement* module. This module uses the previously generated priority cover-tree, to efficiently compute a refined tree. The tree is again passed to the Priority Medoid Approximation module for choosing item representatives, and the result is presented to the user.

Experiments setting: In our experiments, DiRec was employed on top of a CF-based recommender system (C2F [20]). For estimating item similarity, C2F (and thus DiRec) uses Pearson’s correlation coefficient [19]. We used a natural linear inverse mapping to compute items distance out of their Pearson correlation value: $dist(i, j) = (1 - Pearson(i, j))/2$. The experiments were performed on an Intel quad-core machine (Q9400) with 2.66 GHz CPUs, (using only a single core of the CPU), 4GB memory and windows XP x64 edition. As data, we used a real-life data set from the cinema domain, provided by Netflix [10]. This data set contains over 100 million distinct raw movie ratings (such as 1 to 5 “stars”), by approximately 500,000 users, and no semantic information on the movies.

Algorithms: Our algorithm for representatives selection first constructs a priority cover tree, then chooses representatives out of this tree. We consider the three variants of this choice: Max-rating, Max-diversity and Max-coverage, denoted below *PCT-R*, *PCT-D* and *PCT-C*, respectively. (PCT stands for Priority Cover-Tree). As we have mentioned, our use of *priority* cover-tree, instead of the classical cover-tree of [18], allows to take items rating into consideration. To illustrate the resulting improvement in the quality of the generated recommendations we had also implemented the cover-tree algorithm of [18], denoted below *CT*.

Recall that previous works typically use semantic information, to diversify items, and predefined weights/thresholds, to balance between rating and diversity. Since no semantic information is available in our context, one can employ here only algorithms that operate without it. We have implemented the state of the art such algorithms from [6] (*Greedy* and *Swap*) and compared them to ours. Note that Algorithm *Swap*, like all other previous works (except for Algorithm *Greedy*), use predefined thresholds to balance rating and diversity. We had thus first optimized the thresholds (to be explained in Section V-E) and had our algorithms compete against these optimized

versions. The results for *Greedy* and *Swap* were similar and we thus show below only those of *Swap*.

Finally, as a base line, we had also implemented the two “extreme case” algorithms, *optR* (for Optimal Rating), which operated like a standard CF algorithm, selecting the k items having highest rating and ignoring diversity, and *optD* (for Optimal Diversity), which ignores the rating of items and selects k items whose pairwise distance values is the highest.¹

A. Users feedback

To assess the quality of the item sets that DiRec generates (and thereby the quality of our algorithms), we first ran a set of experiments with real users who were asked to evaluate the quality of the presented movie recommendations. A preliminary set of experiments was used to select the optimal threshold value for *Swap* (and thus, represents also Algorithm *Greedy*). Then in the remaining experiments we compared our *PCT* algorithms to *Swap* (with this value), *CT*, *optR* and *optD*.

In each experiment, we first presented to the user a list of 50 highly ranked movies, with their prediction values, selected by the underlying CF system. We then asked the user to consider various subsets, each consisting of 5 movies (recommendations), generated by the different algorithms. 50 users participated in the trial. Each time the user was presented seven recommendation sets, generated by the different algorithms, and was asked to rate each one on a scale of 1 to 10 (10 being the highest grade). The average grade was computed for each algorithm and rounded to an accuracy of 0.1. Algorithm *PCT-R* got the highest results (8.9), with *PCT-C* (8.4), *Swap* (8.3) and *PCT-D* (8.1) not so far behind. Not surprisingly, these four algorithms got much higher results than *optR* (5.2), *optD* (4.7) and *CT* (3.8), as they consider both the relevance and the diversity of the items (while the others focus solely on one property). While the grade difference between *PCT-R* and *Greedy* is smaller, recall that the former has further the advantage of supporting a natural zoom-in mechanism, as demonstrated below in Section V-D.

B. Rating vs. diversity

To better understand the results and how the *PCT* algorithms (and the preferred *PCT-R* in particular) balance ranking and diversity, we run a second set of experiments over the Netflix data where we analyzed the quality the item sets generated by all the algorithms, in terms of their rating and diversity. In each set of experiments reported below we randomly chose a set of 1000 users from the Netflix data set and run the experiment for each of them. The results presented here are the average of all 1000 users. (The variance was less than 5%). We ran all experiments for varying size I of item candidates returned by the CF-system (I ranging from 10 to 1000) and varying number k of representative items selects out of them (k

¹Note that *optD* require solving an NP-hard problem [5]. Its EXP-time algorithm iterates over all subsets of size k to find the best one. In our experiments this became infeasible for items sets of size ≥ 100 , and we thus show results for them only up to that size.

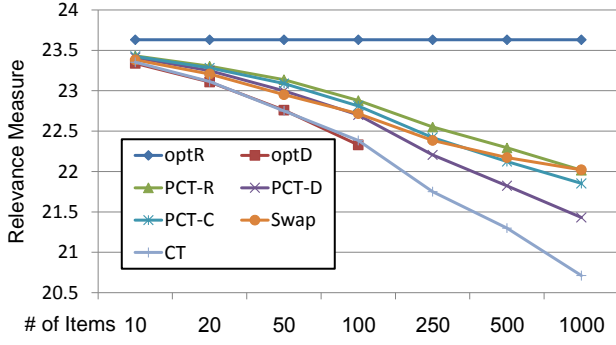


Fig. 4. Rating measure




Fig. 5. Distance-based diversity

ranging from 3 to 20). The results were generally independent of the number k of selected representative. Thus, except when stated otherwise, we present below a representative sample for $k = 5$, the common number of recommendations given by typical recommender systems. For a generated set I_k of representatives, we evaluated its quality w.r.t to the following measures.

Rating: We measure the relevance of the suggested set of item I_k to the current user by the sum of the element ratings, i.e. $rating(I_k) = \sum_{i \in I_k} rate(i)$. (Higher value means better results).




Fig. 6. Semantic (sequel) diversity




Fig. 7. Sequel diversification vs. distance diversification

Diversity: In the absence of semantic information, the diversity of the items in I_k may be measured by the pairwise-distances of the items, namely $diversity(I_k) = \sum_{i,j \in I_k} dist(i,j)$. (Here too, higher value means better results).

Although our algorithms do not require semantic information, it is interesting to examine the actual semantic diversity of the selected items. To get a sense of this, we extracted for each movie information from the Internet Movie Databases IMDb [21] and used it to determine, for each movie, to which of series it belongs (if any). We also extracted the movie's genre (action, drama, etc.). For diversification, it is intuitively preferable not to have multiple movie sequels in the set presented to the user, and to have movies from different genres. We quantify this as follows.

- We denote below by $sequel(i)$ the series to which an item (movie) i belongs. We count the number of distinct series to which the movies in I_k belong and divide this by k , the number of movies in I_k . Namely, $sequelDivers(I_k) = \frac{|\{sequel(i) | i \in I_k\}|}{k}$. Observe that $sequelDivers$ is a number in the range of $(0 : 1]$. Higher values reflect higher diversity, hence better result.
- We assume some total order (e.g. lexicographical) on the possible move genres and denote by $genres(i)$ a boolean vector that records the (possibly multiple) genres of the movie i . (A value 1 in entry j means that the movie i belongs to genre j). To measure diversity we consider the pairwise distances (by the Cosine measure) between the vectors. Then, $genreDivers(I_k) = 1 - average\{cosine(genres(i), genres(j)) \mid i, j \in I_k\}$. Here again $genreDivers$ is a number in the range of $(0 : 1]$ with higher values reflecting higher diversity.

Figures 4, 5 and 6 shows the values for the relevance, distance-based diversity, and sequel-based diversity, correspondingly, for the various algorithms, for varying sizes of item sets I . The results for genre-based diversity are similar to those of sequel-based one and are thus omitted for space constraints. The figure depicts these values, for the sets computed by the seven algorithms mentioned before: *PCT-R*, *PCT-D*, *PCT-C*, *Swap* (with its tuned threshold), *CT*, *optR* and *optD*. k here equals 5 and the size of the set I ranges from 10 to 1000.²

Recall that higher values here mean better results. Diversity

²Recall that *optD* became infeasible for items sets of size greater than 100, and we thus show results for them only up to that size.

	Rating	Distance-based diversity	Sequel diversity
PCT-D	-	+	=
PCT-C	-	=	=
Swap	-	+	-
CT	- -	=	=
optR	++	- -	- -
optD	- -	++	-

TABLE I
SUMMARY OF THE RESULTS (RELATIVE PCT-C)

generally increases when I grows, as there are more items to choose from; Rating decreases, as more diverse but lower ranked items are chosen. Table I provides a summary of the graphs in Figures 4, 5 and 6 that highlights how the various algorithms perform (for the various measures) relative to *PCT-R*, and helps to explain why users may have found its results superior. For an algorithm A and measure M , an “+” (resp. “++”) entry indicates that Algorithm A generally performed slightly (significantly) better than *PCT-R* in measure M . Similarly, “-” (resp. “--”) indicates that A performed slightly (significantly) worse than *PCT-R* in M . An “=” entry indicates similar performance.

As expected *optR* have the highest rating value and lowest diversity. Analogously, *optD* has lowest rating value and highest distance-based diversity. *CT* too has lowest rating value as it ignores item ranks. An interesting point to observe is that although some algorithms (and in particular *optD*) achieve better distance-based diversity than *PCT-R*, no algorithm achieves better sequel diversity. To understand this discrepancy between distance-based and semantic (sequel) diversity, consider the following example.

Example 5.1: Consider a simple scenario where the set I consists of six movies, out of which 5 need to be selected. Assume that two of the six movies are sequels. Let us denote them by i_1 and i'_1 and the remaining movies by i_2-i_5 . Figure 7 depict this scenario with two different selected subsets. Intuitively, the most diversified, sequel-wise, set of representatives, is the one that includes i_2-i_5 and either i_1 or i'_1 (as no sequels are present in this selection). This selection is illustrated on the left-hand side of the figure. The filled circles denote the selected movies. Assume that the euclidian distance between the items reflects their actual $dist(i, j)$ measures. Since i_1 and i'_1 are sequels, they are fairly close. When applying an algorithm which maximize the abstract diversity measure, however, we obtain here a different set of representatives, that maximizes the distance between the elements, as shown in the right part of the figure. Here, although i_1 and i'_1 are close to each other, they are nevertheless both selected as they are relatively far from all other items, and the algorithm aims to the sum of $dist(i, j)$ values for all pairs of representatives.

The same phenomenon, illustrated by that particular example, occurs more generally in many real life cases and explains why *PCT-R* is not handicapped by its non optimal distance

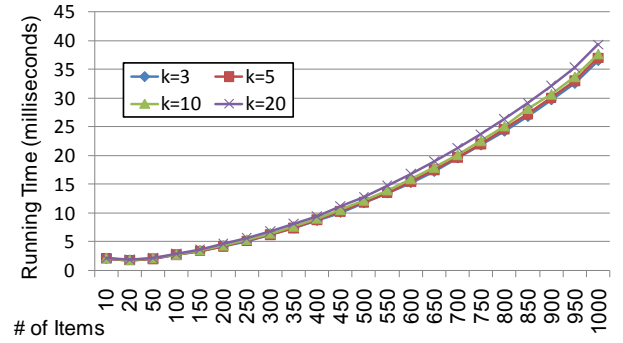


Fig. 8. Running time

diversity. Indeed, its sequel diversity is close to 1 (meaning that, as desired, the result contains no sequels) and we can see that no other algorithm was able to achieve better semantic (sequel) diversity. Furthermore, among the algorithms that achieve sequel diversity equal to that of *PCT-R*, it is the one with highest rating score, which can explain why its recommendations were more favorable than the rest.

It is important to note that, for *Swap*, the value for the above measures depends on the chosen threshold: The threshold determines what is the lowest rating that is allowed to be considered for diversification, so lower threshold values yield lower rating score and, correspondingly, higher diversity score. We present in Section V-E experimental results that illustrate this drop (resp. increase) in relevance (diversity) score, as the threshold decreases. Interestingly, as seen in Figures 4 and 5, the rating and diversity values yielded by the tuned *Swap* are rather close to those of our *PCT* algorithms (obtained naturally, with no need for tuning).

C. PCT Performance

To conclude the discussion of our *PCT* algorithms we show that they operate fast enough to guarantee a pleasant user experience. Figure 8 depict the average time (in milliseconds) it takes to compute I_k , for a set I of size ranging from 10 to 1000 and varying k values, for Algorithm *PCT-R*. The running time for Algorithms *PCT-D* and *PCT-C* is similar (and thus omitted). This is because most of the running time of the algorithms is spent on the construction of the priority cover-tree, which is similar for all *PCT* algorithms. This is also the reason why the results are only marginally effected by the number k of selected representatives. We can see that the running time naturally increases (polynomially) with the number of items, but even for 1000 item candidates it takes less 0.04 of a second to select a representative subset.

D. Recommendations Refinement

When a user zooms-in on an item i , DiRec invokes the refinement algorithm to computes a set of representatives for $relevant(i)$. We next analyze the efficiency of the algorithm and the quality of its output.

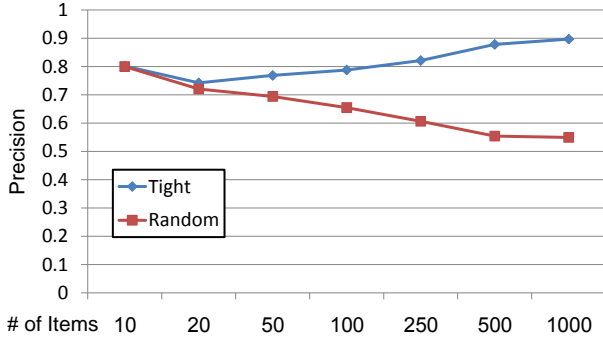


Fig. 9. Subtree precision




Fig. 10. Subtree recall

Efficiency: Our optimized refinement algorithm is based on the observation that much of the items in $relevant(i)$ already belong to the subtree rooted at i , in the existing priority cover-tree. This phenomenon is due to the *tight insertion* employed by our algorithm (as opposed to the random insertion used in the classical cover-tree construction). To demonstrate the effectiveness of tight insertion, relative to a random one, we implemented a variant of our algorithm that uses random insertion. In the following we refer to our algorithm as *Tight* and to its variant with random insertion as *Random*. We compared the precision and recall of the subtrees generated with the two variants. Namely, the average fraction of the relevant items, out of all the items in the subtree (resp. out of all the desired ones). Formally,

$$precision(I_k) = average\left\{\frac{|subtree(i) \cap relevant(i)|}{|subtree(i)|} \mid i \in I_k\right\}$$

$$recall(I_k) = average\left\{\frac{|subtree(i) \cap relevant(i)|}{|relevant(i)|} \mid i \in I_k\right\}$$

Here, $subtree(i)$ is the set of all the items in the subtree rooted at i . The results are depicted in Figures 9 and 10 for $k = 5$ and I ranging from 10 to 1000 (the results for other k values were similar). We can see that tight insertions yields significantly better precision and recall than the random one and the gap increases as I grows. Note that while the precision

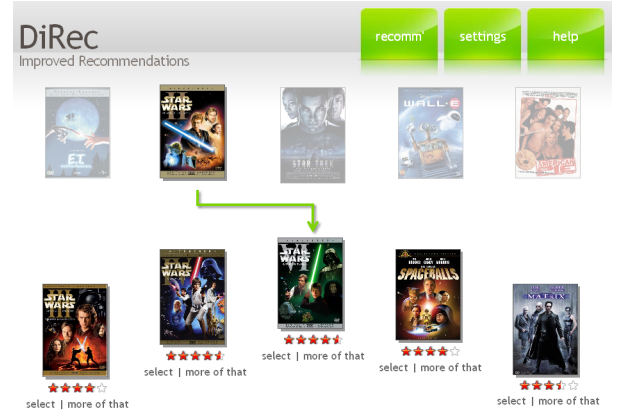


Fig. 11. DiRec “more of that” (zoom-in) mechanism

of *Tight* increased with the growth of I , the recall decreases. This is expected as the increased number of items generates wider levels in the constructed tree; This causes more items to not reside under any representative subtrees, thus decreasing the recall values.

Yet, the relatively high precision and recall of the tight insertion makes our optimized, refinement algorithm very efficient. To demonstrate this we compared the running time of our algorithm to that of the naive one which constructs a new priority cover-tree from scratch. We measured the time it takes each of the two algorithms to refine all k representatives (simulating a user click on each of the presented items) and computed, for each algorithm, the average time for a single refinement. Our optimized algorithm was in all cases at least twice faster than the naive one. The improvement further increased to a factor of three when the size of I was between 50 and 500. This is because the ratio between the precision and the recall of the subtree is better in this interval, as shown in Figures 9 and 10. These good results persisted for all k values.

Quality: Finally, to evaluate the quality of our refinement process, we checked how many of the sequels of a given movie are effectively retrieved by our zoom-in mechanism. Recall from section V-B that $sequel(i)$ denotes the series to which a movie i belongs. Let $sequels(i) = \{j \mid sequel(j) = sequel(i) \wedge j \in I\}$. We checked in this experiment what portion of $sequels(i)$ are included in $relevant(i)$. Interestingly, we found that in over 90% of the cases, $sequels(i) \subseteq relevant(i)$, namely, all sequels are indeed identified. Moreover, all the other cases were movies with many sequels, with just one of them not appearing in $relevant(i)$. These positive results are particularly interesting given the fact that no semantic information is used by our algorithms.

For the readers convenience, Figure 11 presents an example for a refined recommendations set generated by DiRec. In this specific case, the user clicked on the “more of that” button of the “Star Wars II: Attack of the Clones” movie. DiRec successfully identifies the Star Wars sequels and presents in response three additional sequels. It also presents two




Fig. 12. Algorithm Swap - Rating Measure




Fig. 13. Algorithm Swap - Diversity Measure

additional movies that do not belong to this series, yet are related, and were chosen by DiRec to provide a more diverse set of recommendations.

E. Tuning Algorithm Swap

We first give a brief description of Algorithm *Swap* (for further details see [6]). Then we illustrate the effect of different threshold values on its operations.

The basic idea behind Algorithm *Swap* is to start with the top- k most relevant items, and then swap the item which contributes the least to the diversity of the entire list with the highest ranked further item (not yet in the list). As we initialized the list with the highest ranked items, the swapped candidates naturally would reduce the overall relevance measure of the list. To prevent large drop in this value, the algorithm employs a hard threshold representing the minimal relevance value required from the swapped items. The difficulty is to determine the best threshold to use (which naturally differ from different data sets).

We tested the affect of different threshold values on the relevance and the diversity of the generated lists. Figures 12 and 13 depict these values, respectively. Lower threshold values allow the algorithm to use items with lower rating values, which naturally decrease the overall rating measure. This is illustrated in Figure 12. Another point to note is the decrease of the rating measure as the number of item increases. This is because the items are sorted by their relevance, and

thus, when more are added, the overall rating measure can only decrease. On the other hand, the diversity of the items, as shown in Figure 13, analogously increases as the threshold decrease and as the number of items increases. In our users study the recommendations obtained with value of 4.2 were most favorable and this is what we used for the remaining experiments.

VI. CONCLUSION

The DiRec plug-in presented in this paper allows CF recommender systems to diversify the recommendations that they present to users. It is based on a novel notion of priority-medoids that declaratively balances the rating and the diversity of the recommended items. We introduced priority cover-trees as a tool for efficient selection of item representatives. Our solution further allows for an effective realization of a natural “zoom-in” mechanism, presenting to users similar yet diversified items. Our experiments demonstrated the effectiveness of our recommendation diversification technique and its superiority over previous proposals.

DiRec provides an effective solution in common scenarios where semantic information is unavailable. Combining our ratings-based (quantitative) approach with a semantic (qualitative) one, when such semantic data is available, is an intriguing future research direction.

REFERENCES

- [1] G. Adomavicius and A. Tuzhilin, “Towards the next generation of recommender systems,” *IEEE TKDE*, 2005.
- [2] X. Su and T. Khoshgoftaar, “A survey of collaborative filtering techniques,” *Advances in Artificial Intelligence*, 2009.
- [3] C.-N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen, “Improving recommendation lists through topic diversification,” *WWW*, 2005.
- [4] C. Yu, L. Lakshmanan, and S. Amer-Yahia, “Recommendation diversification using explanations,” *ICDE*, 2009.
- [5] M. Drosou and E. Pitoura, “Search result diversification,” *SIGMOD Record*, 2010.
- [6] C. Yu, L. Lakshmanan, and S. Amer-Yahia, “It takes variety to make a world: Diversification in recommender systems,” *EDBT*, 2009.
- [7] L. Kaufman and P. Rousseeuw, “Finding groups in data: An introduction to cluster analysis,” *Wiley's Series in Probability and Statistics*, 1990.
- [8] A. Beygelzimer, S. Kakade, and J. Langford, “Cover trees for nearest neighbor,” *ICML*, 2006.
- [9] R. Boim, T. Milo, and S. Novgorodov, “Direc: Diversified recommendations for semantic-less collaborative filtering,” *ICDE*, 2011.
- [10] J. Benet and S. Lanning, “The netflix prize,” *KDD Cup*, 2007.
- [11] M. Zhang and N. Hurley, “Evaluating the diversity of top- n recommendations,” *ICTAI*, 2009.
- [12] S. Gollapudi and A. Sharma, “An axiomatic approach for result diversification,” *WWW*, 2009.
- [13] M. Zhang and N. Hurley, “Avoiding monotony: Improving the diversity of recommendation lists,” *RecSys*, 2008.
- [14] Z. Chen and T. Li, “Addressing diverse user preferences in sql-query-result navigation,” *SIGMOD*, 2007.
- [15] G. Koutrika, A. Simitsis, and Y. Ioannidis, “Precis: The essence of a query answer,” *ICDE*, 2006.
- [16] K. Stefanidis, M. Drosou, and E. Pitoura, “Perk: Personalized keyword search in relational databases through preferences,” *EDBT*, 2010.
- [17] J. Stoyanovich, S. Amer-Yahia, and T. Milo, “Making interval-based clustering rank-aware,” *to appear in EDBT*, 2011.
- [18] B. Liu and H. Jagadish, “Using trees to depict a forest,” *VLDB*, 2009.
- [19] J. L. Rodgers and W. A. Nicewander, “Thirteen ways to look at the correlation coefficient,” *The American Statistician*, 1988.
- [20] R. Boim, H. Kaplan, T. Milo, and R. Rubinfeld, “Improved recommendations via (more) collaboration,” *WebDB*, 2010.
- [21] “Imdb interface,” <http://www.imdb.com/interfaces/>.