

EFFICIENT COLLABORATIVE FILTERING IN CONTENT-ADDRESSABLE SPACES

SHLOMO BERKOVSKY*, YANIV EYTANI† and LARRY MANEVITZ‡

*Computer Science Department, University of Haifa
31905, Mount Carmel, Haifa, Israel*

**slava@cs.haifa.ac.il*

†eytani@cs.haifa.ac.il

‡manevitz@cs.haifa.ac.il

Collaborative Filtering (CF) is currently one of the most popular and most widely used personalization techniques. It generates personalized predictions based on the assumption that users with similar tastes prefer similar items. One of the major drawbacks of the CF from the computational point of view is its limited scalability since the computational effort required by the CF grows linearly both with the number of available users and items. This work proposes a novel efficient variant of the CF employed over a multidimensional content-addressable space. The proposed approach heuristically decreases the computational effort required by the CF algorithm by limiting the search process only to potentially similar users. Experimental results demonstrate that the proposed heuristic approach is capable of generating predictions with high levels of accuracy, while significantly improving the performance in comparison with the traditional implementations of the CF.

Keywords: Collaborative filtering; recommender systems; K -nearest neighbors search; content-addressable systems.

1. Introduction

In many circumstances, the quantity of available information grows rapidly and exceeds our cognitive processing capabilities. Thus, there is a pressing need for intelligent *personalization* systems providing services tailored to users' real needs and interests. Recommender Systems¹⁶ are one of the commonly used approaches to address this problem. These systems assist a user in selecting a suitable item among a set of potentially selectable items by predicting the user's opinion on the items by applying statistical and knowledge discovery techniques.¹⁹ Currently, Recommender Systems are used in a variety of application domains, e.g. movies,⁷ jokes,⁶ music¹ and others, and they exploit various recommendation techniques, such as Collaborative Filtering,⁹ Content-Based Filtering,¹² Case-Based Reasoning,¹⁷ and numerous hybrid techniques.⁴

Collaborative Filtering (CF) is probably the most familiar and one of the most widely-used techniques to generate predictions in Recommender Systems. It relies

on the assumption that people who agreed in the past will also agree in the future.²¹ The input for the CF algorithm is a matrix of users' ratings on a set of items, where each row represents the ratings provided by a single user and each column represents the ratings provided by different users on a single item. CF aggregates the ratings to recognize similarities between users and generates the prediction for an item by weighting the ratings of similar users on this item.

The CF algorithm is typically partitioned to three generic stages: (1) Similarity Computation: weighting all the users with respect to their similarity with the *active user* (i.e. the user, whose ratings are being predicted), (2) Neighborhood Formation: selecting the most similar users for the prediction generation, and (3) Prediction Generation: computing the prediction by weighting the ratings of the selected users.

One of the major drawbacks of the CF is its limited scalability. The stages of Similarity Computation and Neighborhood Formation require comparing the active users with all the other users over all the available ratings. Hence, the computational effort required by the CF grows linearly both with the number of users and the number of items in the ratings matrix. Thus, for a matrix containing ratings of M users on N items, the required computational effort is $O(MN)$. This poses a problem in real-life systems, where the predictions are generated using millions of ratings on thousands of items, e.g. in Web-based Recommender Systems. Previous studies, (e.g. Refs. 3, 5 and 6 and others) tackle the issue of reducing the computational effort required by the CF either by preprocessing of the ratings matrix or by distributing the heavy computational stages. Nonetheless it remains one of the most important issues in the CF research community.

In this work we develop a fast heuristic variant of the CF algorithm that decreases the computational effort required by the Similarity Computation and the Neighborhood Formation stages. The basic assumption of the proposed heuristic algorithm is that losing general completeness of the exhaustive search (1) has a minor negative effect on the accuracy of the predictions, but (2) significantly decreases the required computational effort. Thus it provides a scalable approach, applicable to real-life scenarios with a high number of users and items, such as in Web-based systems.

The proposed heuristic approach is based on a notion of content-addressable data management¹⁵ that provides an adaptive topology for mapping of users' profiles to a multidimensional space. This mapping implicitly clusters similar users and limits the Similarity Computation and the Neighborhood Formation stages to a heuristic search among the users that are potentially highly similar to the active user.

Experimental evaluation of the proposed approach demonstrates both high efficiency and good accuracy of the proposed algorithm in comparison with the traditional (exhaustive) *K-Nearest Neighbors* (KNN) search of the Neighborhood Formation stage. The evaluation also demonstrates that the algorithm is highly scalable with the number of nearest neighbors to be retrieved.

The rest of the paper is organized as follows. Section 2 describes the CF personalization technique and surveys the studies focusing on the required computational effort reduction. Section 3 describes the CAN, a Peer-to-Peer content-addressable platform for decentralized data management. Section 4 describes the decentralized storage of users' profiles over the CAN platform and elaborates on the proposed heuristic variant of the CF over CAN. Section 5 presents and analyzes the experimental results. Finally, Sec. 6 lists our conclusions and presents some open questions for future research.

2. Collaborative Filtering

Collaborative Filtering (CF) is probably one of the most familiar and widely-used recommendation techniques. An input for the CF is the so-called *ratings matrix*, where each user is represented by a set of explicit ratings given on various items, and each item is represented by a set of ratings given by the users.

CF requires a similarity metric between users to be explicitly defined. The state-of-the-art CF systems exploit three similarity metrics: Cosine Similarity,⁷ Mean Squared Difference (MSD),¹³ and Pearson correlation.¹⁹ This work focuses on the MSD, computing the degree of similarity between users x and y by:

$$\text{sim}_{x,y} = \frac{\sum_{i=1}^{|x \cap y|} (R_{x,i} - R_{y,i})^2}{|x \cap y|} \quad (1)$$

where $|x \cap y|$ denotes the number of items rated by both users (typically, above some minimal threshold), and $R_{x,i}$ denotes the rating of user x on item i . In some sense, $\text{sim}_{x,y}$ can be considered as the *dissimilarity* of the users, as the lower the result of the MSD computation, the greater is the real similarity between the users.

Prediction $P_{a,j}$ for the rating of the user a on item j is computed as a weighted average of the ratings of his/her K most similar users, i.e. K *nearest neighbors*, by:

$$P_{a,j} = R'_a + \frac{\sum_{k=1}^K (R_{k,j} - R'_k) \cdot \text{sim}_{a,k}}{\sum_{k=1}^K |\text{sim}_{a,k}|} \quad (2)$$

where $R_{x,y}$ denotes the rating of user x on item y , R'_z denotes the average rating of user z , and $\text{sim}_{v,u}$ denotes the level of similarity between users v and u .

The Similarity Computation stage of the CF requires comparing the active user with every other user in the system. For a ratings matrix storing the ratings of M users on N items, the computational complexity of the Similarity Computation stage is $O(MN)$. This indicates poor scalability of the Similarity Computation stage, as the complexity grows linearly with both the number of users and the number of items in the matrix.

2.1. Reducing the computational effort required by the CF

Many prior works have dealt with decreasing the computational, effort required by the CF. In general, it is achieved either by preprocessing the ratings matrix, or by distributing the computationally intensive stages of the CF among multiple machines.

Various preprocessing techniques for decreasing the computational effort required by the CF (e.g. correlation coefficients, vector-based similarity, and statistical Bayesian methods) are discussed and analyzed in Ref. 3. Another technique, exploiting preclustering of the ratings matrix, is discussed in Ref. 6. There, principal component analysis is used to identify two *discriminative* dimensions of the ratings matrix and all the vectors are projected onto the resulting plane. This inherently partitions the users to clusters or neighborhoods, which are further used to generate the predictions. In Ref. 5, the authors use a tree-like data structure and apply a *divide-and-conquer* approach using an iterative K -means clustering to group the users. This leads to smaller and more homogeneous clustering of users for the following Predictions Generation stage.

An alternative approach is to distribute the computational effort required by the CF among the users, such that every user independently computes its similarity with the active user. This approach was initially proposed in Ref. 22 and elaborated in Ref. 20. The latter also developed a detailed taxonomy of the CF distribution approaches and presented implementation frameworks for different application domains. The PocketLens project¹¹ compared five decentralized distributed architectures for the CF. These comparisons showed that the performance of the decentralized mechanism is similar to the performance of the centralized CF while providing increased robustness and security.

Further improvements to the decentralized CF were discussed in Ref. 8, which proposes the exploitation of Peer-to-Peer platform for a decentralized management of users' profiles. However, this approach approximates the set of the most similar users identified by the Neighborhood Formation stage of the CF, and as a result, the accuracy of the generated predictions is reduced.

This paper is loosely based on the ideas of CAN,¹⁵ a content-addressable Peer-to-Peer platform. We implement a fast heuristic variant of the CF, using a CAN-like multidimensional space for maintaining a connected structure of users. This allows to significantly decrease the computational effort required by the Similarity Computation and Neighborhood Formation stages by limiting the search process to a search among potentially similar users located in close vicinity to the active user.

3. Content-Addressable Data Management

This section presents the general architecture of CAN,¹⁵ a scalable decentralized data management platform. In CAN, the users are represented in a one-to-one manner by the nodes of a virtual N -dimensional coordinate space such that the location

of the user's node is denoted by a vector (v_1, v_2, \dots, v_N) , where v_i represents the numeric coordinate of the node within a dimension number i . In addition to the node, each user continuously manages an N -dimensional subspace, called a *zone*. For example, consider a two-dimensional space partitioned to three zones, managed by users A , B , and C [Fig. 1 (left)]. Note that the figure shows only the zones managed by the users, whereas the nodes themselves are not shown.

In CAN space, two nodes (and also zones) are called *neighbors* if their coordinate spans overlap along $N - 1$ dimensions and adjoin along one dimension. For example, consider the neighbor zones A and C in Fig. 1 (left), whose coordinates partially overlap across the horizontal dimension and adjoin along the vertical. To maintain connectivity in CAN space, each node stores a data structure representing a list of pointers to a set of other nodes, managing the neighbor zones. For example, node A stores the pointers to the nodes managing zones B and C (as, respectively, horizontal and vertical neighbors) in its list of pointers.

Routing of messages in CAN space is based on the Plaxton routing algorithm.¹⁴ This routing iteratively forwards the messages to the nodes that are closer to the target node than the current node using a greedy forwarding. The metric for evaluating the distance between two nodes in the address space is the L_1 metric, i.e. the Manhattan Distance. This metric was chosen due to the fact that CAN space inherently supports it, as every node stores a list of pointers to the nodes, managing the neighbor zones. For example, the distance between the nodes $(1, 2, 3)$ and $(6, 5, 4)$ in three-dimensional CAN space is $(6 - 1) + (5 - 2) + (4 - 3) = 9$. Thus, in N -dimensional CAN space a message is routed between an arbitrary pair of nodes in $O(N)$ routing steps.

In addition, CAN provides a connectivity maintenance algorithm, stable to sporadic joints and departures of new users. When a new user is inserted, it is assigned its own node and the respective zone. This is done by splitting a zone (determined by the content provided by the recently inserted user) of one of the existing neighbors according to the following steps: (1) the new user identifies an existing network node, (2) the new user is routed to the target zone that will be split, and (3) the target zone is split and the neighbors of the new zone are updated to maintain connectivity and facilitate future routings. As a result, only a subset of immediate neighbor zones of the zone that was split is actually affected by the insertion of a new node.

The issue of splitting the target zone (i.e. how to split the existing zone, where the contents of the recently inserted node are mapped?) is one of the important




Fig. 1. Example of a two-dimensional CAN space.

issues affecting the performance of CAN. A number of splitting policies are proposed, analyzed and compared in Ref. 15. The simplest policy for the zones splitting is so-called *ordered* splitting. According to this policy, the number of dimension, across which a zone is split, iteratively increases from 1 to N .

For example, consider user D joining CAN two-dimensional space [Fig. 1 (middle)]. Assuming that the content provided by user D should be located in the right part of the zone managed by node C and this is the zone that will be split, user D is routed to C using the Plaxton routing, and zone C is split across the horizontal dimension (assuming that the previous split of zone C , and also the following split of both zones C and D will be performed across the vertical dimension). Finally, the recently inserted node, managing the zone D notifies its neighbors (i.e. the users “managing zones B and C) about the insertion of a new node, and also their neighbors’ pointer tables are updated. Note that in this case, only the zone managed by user C , which was split and a subset of its neighbor zones (actually, only one zone managed by user B), are affected by the insertion of a new user D , whereas other zones are not affected.

Disconnections of the users are handled in a similar manner. The disconnecting user identifies one of the neighbor nodes that will takeover managing its zone, and updates other neighbor zones about the departure and the management takeover. For example, consider the user managing zone B disconnecting from CAN space [Fig. 1 (right)]. As a result of the disconnection, the user managing zone D takes over the management of the zone previously managed by user B .

Nonetheless, it should be mentioned that the dimensionality of the above CAN space can be barely extended. Such extension requires remapping of the existing N -dimensional nodes to a new $(N + 1)$ -dimensional space. This is an expensive procedure in a decentralized P2P environment, which can also involve multiple interactions with the users. Thus, in this work we assume that the dimension of CAN space is fixed.

Thus, CAN provides a decentralized platform, supporting (1) dynamic space partitioning and zones allocation, (2) efficient routing algorithm, and (3) connectivity maintenance algorithm over virtual N -dimensional coordinate space. Note that the distributed structure of CAN is not robust against sudden departures of users, as fault-tolerance is not one of the main goals of the platform. However, CAN facilitates a decentralized self-manageable platform for content-addressable data management in a distributed environment.

4. CF over Content-Addressable Space

This work proposes an efficient heuristic variant of the CF algorithm. It uses a content-addressable architecture for the purposes of optimizing traditional exhaustive K -Nearest Neighbors (KNN) search to a search among potentially similar users only. Although our algorithm is a heuristic one by nature, experimental results demonstrate that it facilitates efficient search process without hampering the accuracy of the generated predictions.

4.1. Mapping user profiles to content-addressable space

The input for the CF algorithm is a matrix of users' ratings on items, where each row (ratings vector) represents the ratings of a single user and each column represents the ratings on a single item. The total number of items (N) defines an N -dimensional space, where the coordinates range in each dimension corresponds to the range of ratings on the respective item.

To handle the ratings matrix in a content-addressable manner, we map it to a CAN-like multidimensional space. Each rating is projected using a uniform *injective (one-to-one) mapping* onto the appropriate dimension, such that the whole vector of length N is mapped to a single point in an N -dimensional space. For example, consider a system storing the ratings of users on three different items. In such a system, the evolving CAN-like space will be a three-dimensional cube, where the range of coordinates within every dimension corresponds to the range of possible ratings on the respective item.

As already mentioned, each user is represented in a CAN-like space by a single node whose location corresponds to the set of user's ratings and by the respective zone (storing a list of immediate neighbor zones). For example, consider a user U that rated all three items in the above three-dimensional cube: item i_1 was rated as r_1 , item i_2 as r_2 , and i_3 as r_3 . The user will be mapped to a location (r_1, r_2, r_3) of the space and will have exactly two neighbors in each dimension. For example, in the dimension corresponding to item i_1 , the user U will have two neighbors, $N_1 = (r_1 - x, r_2, r_3)$ and $N_2 = (r_1 + y, r_2, r_3)$, such that both N_1 and N_2 rated i_2 as r_2 and i_3 as r_3 , N_1 rated i_1 below r_1 , and N_2 rated it above r_1 , and there is no other user that rated i_1 as r' , where $r_1 - x < r' < r_1$ or $r_1 < r' < r_1 + y$. Similarly, user U will have two neighbors in the dimension corresponding to item i_2 and to item i_3 . If there is no user that provided the required combination of ratings on the available items, CAN space will maintain connectivity by connecting user U to a further node, which will serve as its immediate neighbor (i.e. both zones will keep mutual pointers to the relevant neighbor zone).

Note that in the evolving CAN space, the users (through their ratings vectors) can be dynamically inserted and removed not only during the initialization, but also during the life cycle of the system. This is explained by the observation that the above connectivity maintenance algorithm guarantees that the structure remains connected regardless of the sudden joints and disconnections of the nodes. Nevertheless, CAN spaces can barely manage insertions of new items, as the dimension of the space should remain fixed. Thus, the proposed heuristic search (that will be discussed in the following subsection) is applicable only over a stable matrix of users' ratings, where no new items are inserted.

Deciding on the zones split policy affects the evolving structure of the ratings vectors. In our implementation, we used the above mentioned ordered splitting policy. This policy may be suboptimal in terms of the number of neighbor zones, resulting in a less efficient algorithm, i.e. more comparisons or retrieving less similar neighbors. However, our experiments demonstrate that even this simple

policy considerably increases the efficiency of the proposed K -Nearest Neighbors (KNN) search, in comparison with the traditional exhaustive search. Evaluating other splitting policies is beyond the scope of this work.

In addition to the guaranteed connectivity, content-addressable space *inherently clusters* similar users, such that the distance between two similar users (in our case, according to the MSD similarity metric) is lower than the distance between two arbitrary users. This is achieved due to the use of an injective mapping of the ratings vector to the multidimensional CAN-like space, which preserves the users' similarity while mapping the ratings vectors to the numeric coordinates in the space. The following subsection shows a use of the above inherent clustering property for the purposes of developing fast heuristic variant of the KNN search.

4.2. *Heuristic nearest-neighbors search*

The Neighborhood Formation stage of the CF over the evolving N -dimensional space can be schematically described as a heuristically expanding breadth-first search. The algorithm for retrieving K -Nearest Neighbors of a user x is briefly explained by the following pseudo-code. The code uses two lists of size K : (1) CANDIDATES — list of candidates for being one of the K -nearest neighbors, and (2) NEIGHBORS — list of real K -Nearest Neighbors. In principle, the algorithm needs the CANDIDATES list only, as the NEIGHBORS list only increases during the execution of the algorithm until it reaches its maximal length and contains the real K -Nearest Neighbors. For the sake of clarity, we show an algorithm that uses two lists instead of only one.

K _Nearest_Neighbors (*user* x)

- (1) let *NEIGHBORS* and *CANDIDATES* be empty lists, each of size K
- (2) let Z be the zone, to where x would be mapped in the CAN space
- (3) foreach $u \in (Z \cup \text{neighbors}(Z))$
- (4) compute $\text{distance}(x, u)$
- (5) insert u into *CANDIDATES*, s.t. *CANDIDATES* is sorted according to the values of $\text{distances}(x, u)$
- (6) for $i=1$ to K
- (7) choose v from *CANDIDATES*, s.t. $\text{distance}(x, v)$ is smallest
- (8) for each $w \in \text{neighbors}(v)$ s.t. $\text{distance}(x, w)$ is unknown
- (9) compute $\text{distance}(x, w)$
- (10) insert w into *CANDIDATES*, s.t. it remains sorted according to the values of $\text{distances}(x, v)$
- (11) move v from *CANDIDATES* to *NEIGHBORS*
- (12) return *NEIGHBORS*

Initially, the algorithm pretends to map the active user x to its location in the N -dimensional space (step 2). Next, the algorithm identifies the zone x is mapped

to, and its neighbors, i.e. users managing the neighbor zones (step 3). For each of these zones, the degree of similarity, i.e. the distance between x and the relevant user, is computed (step 4). Then, the neighbor users are inserted into the CANDIDATES list such that the whole list of candidates users is sorted according to the distances of the users from the active user x (steps 4 and 5).

Afterwards, the algorithm iteratively performs the following operations:

- Selects v , the nearest neighbor stored in the CANDIDATES list (step 7).
- Identifies the neighbors of v that are not in the CANDIDATES list yet, computes their distances from x , and inserts them into the CANDIDATES, while keeping the list sorted (steps 8–10).
- Removes v from the CANDIDATES list and inserts it into the NEIGHBORS list.

Finally, the algorithm returns the resulting NEIGHBORS list (step 12).

Consider an example execution of the KNN search as illustrated in Fig. 2. The initial structure of two-dimensional space is depicted in Fig. 2(a). Nine users, named from a to i , are inserted into the space and manage the respective zones. Note that also this figure shows only the zones managed by the users, whereas the nodes representing the users are not shown. Assume that the active user is mapped to the zone managed by user e .

Thus, e and its neighbors, i.e. users managing zones c , d , f and i , are the first candidates for being the nearest neighbors and they are inserted into the CANDIDATES list. Assume that the user managing zone e is the closest one. It is moved from the CANDIDATES list to the NEIGHBORS list [Fig. 2(a)]. Since all the neighbors of e are already known, the next closest neighbor is chosen among its neighbors. Assume that the next closest neighbor is the user managing zone f . It is moved from the CANDIDATES list to the NEIGHBORS list, and its only new neighbor, the user managing zone g , is inserted into the CANDIDATES list

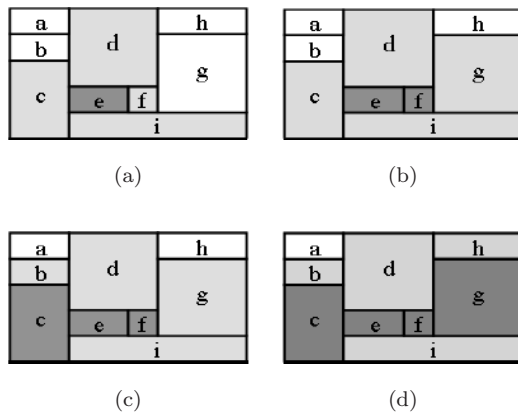


Fig. 2. Stages of the KNN search over two-dimensional CAN space (zones managed by users from the CANDIDATES are indicated with a light tone and from the NEIGHBORS — in a dark gray tone).

[Fig. 2(b)]. The next closest neighbor is the user managing zone c , inserting the user managing zone b into the CANDIDATES list [Fig. 2(c)]. Assume that the next closest neighbor is the user managing zone g (not an immediate neighbor of e). As a result, the user managing zone h is inserted into the CANDIDATES list [Fig. 2(d)]. This process is repeated until the NEIGHBORS list contains K -Nearest Neighbors.

The proposed algorithm reduces the computational effort required by the Similarity Computation and the Neighborhood Formation stages, in comparison with the traditional CF algorithm, where an active user is compared with all the available users. Conversely, the proposed heuristic algorithm compares the active users with potentially similar users only, located in close vicinity to the active user.

Since every user in the N -dimensional space continuously maintains an updated list of its immediate neighbors, any neighbor of a given user is accessed through a single network hop. This is true regardless of the physical (geographical) and logical (similarity) distances between the neighbors. Thus, the algorithm will also work in sparse spaces, where the distance between neighbors in the underlying network might be very high.

4.3. *Heuristic completions of user profiles*

In the former sections, we assumed that the user's ratings were represented as a *complete vector*, i.e. explicit ratings on all the items are available. Thus the mapping of the user's ratings vectors to the underlying content-addressable space is straightforward. However, this assumption is unachievable in most real-life applications and scenarios, where an average user rates only a portion of the available items. This raises a need for developing a mapping mechanism capable of mapping *incomplete vectors*, where a subset of the ratings is missing, to the content-addressable space.

In this subsection we propose three mappings to handle this task. However, instead of developing a new mapping of incomplete vectors to the content-addressable space, we propose to convert the incomplete vectors to complete ones by heuristically filling-in the missing ratings in the incomplete vectors.² Thus, the proposed completion heuristics are designed to reuse the above injective mapping of complete vectors, while employing it on the modified vectors with heuristically filled-in ratings.

As the completion heuristics are not the main focus of the current work, we suffice with three relatively simple heuristics that demonstrate the applicability of the proposed vectors' completion. The heuristics are as follows:

- *User-average* — The missing rating on an item in the user's vector is substituted with the average of the real ratings, explicitly provided by this user.
- *Item-average* — The missing rating on an item in the user's vector is substituted with the average of the real ratings, explicitly provided by the other users on this item.

- *Conditional* — Integrates both the user-average and the item-average heuristics and decides on a run-time regarding the specific completion heuristic to be used according to a certain predefined condition.

Clearly, the *user-average* heuristic can be considered as an accurate personalized completion heuristic, as the missing ratings are substituted with a value, produced by the real ratings of the given user. Thus, it reflects the real preferences and tendencies of the user, such as over- or under-rating of items, natural intensity of expressions and so forth. Conversely, the *item-average* heuristic can be considered as the most accurate non-personalized completion heuristic, as the missing ratings are substituted with a value, produced by numerous real ratings on the given item. As such, it reflects a general (and relatively reliable) opinion of many other users on the item.

We conjecture that the *user-average* heuristic is preferable when the knowledge about the user's preferences is reliable, i.e. the number of ratings explicitly provided by the user is relatively high. On the other hand, when the number of user's explicit ratings is low, the *item-average* heuristic will exploit other users' ratings for filling-in the missing rating and it should be preferred. Based on these considerations, we defined another *conditional* heuristic, which will autonomously decide which of the above completion heuristics should be exploited for filling-in the missing ratings of every user.

In summary, each of these heuristics allows the filling-in of the missing ratings, converting the incomplete vectors to the complete ones, and then mapping them to the content-addressable space using the abovementioned injective mapping mechanism.

5. Experimental Results

In the experimental part of our work we used the Jester dataset of jokes' ratings.⁶ Jester is a Web-based jokes Recommender System, containing 4.1 millions of ratings (on a continuous scale from -10.00 to $+10.00$) of 73,421 users on 100 jokes. A significant portion of the users rated all the jokes, so the Jester dataset is relatively dense. Overall, approximately 56% of all the possible ratings in the matrix are present.

For the complete vectors experiments, we selected a subset of 14,192 users that rated all 100 jokes, producing a matrix, where every value corresponds to a real rating, explicitly provided by a user. The average rating of a single joke in the data is 0.807, and the overall standard deviation of the ratings in the matrix is 4.267. We implemented a centralized simulation of a 100-dimensional CAN space (note that the space dimension equals to the number of rated jokes in the ratings vectors) and inserted the above 14,192 users into the space. Insertions of the users into the space were done using the ordered splitting policy.

5.1. Scalability of the search

These experiments were designed to evaluate the scalability of the proposed heuristic variant of the KNN search. The efficiency of CAN-based KNN is measured by the number of comparisons performed during the Neighborhood Formation stage of the CF.

In this experiment we measured number of comparisons during the Neighborhood Formation stage. For this, we gradually increased the number of users inserted into the system from $M = 1000$ to $M = 14,000$. For each M , we computed the number of comparisons performed in the traditional exhaustive KNN search and in CAN-based heuristic variant of KNN. Both searches were aimed at retrieving $K = 5$ nearest neighbors. For each value of M , the experiments were repeated 1000 times for different active users. The experimental results are shown in Fig. 3. The horizontal axis stands for M , the number of users inserted into the system, and the vertical axis reflects the average number of comparisons during a single KNN search, for both exhaustive and heuristic searches.

As expected, the number of comparisons in CAN-based KNN is significantly lower than in traditional KNN and it grows at a logarithmic-like manner with the number of users. This is explained by the fact that in CAN-based KNN the active user is compared only with a subset of highly similar users (located in close vicinity in a content-addressable space), whereas in traditional KNN it is exhaustively compared with all the available users.

To achieve a better understanding of comparison-based scalability of the proposed approach, we computed the ratio between the number of comparisons in CAN-based KNN and the number of comparisons in the exhaustive KNN. This ratio was computed for different values of M and the results are shown in Fig. 4. It can be clearly seen that the ratio steadily decreases with M . This allows us to




Fig. 3. Average number of comparisons versus the number of users inserted.




Fig. 4. Ratio between the number of comparisons versus the number of users inserted.




Fig. 5. Average number of comparisons versus the number of retrieved neighbors.

conclude that the proposed algorithm is applicable in large-scale systems with high number of users and items, e.g. on the Web.

The second experiment was designed to evaluate the scalability of CAN-based KNN with the number of nearest neighbors (K) to be retrieved. We gradually increased the value of K from $K = 1$ to $K = 50$. For each value of K , we measured the number of comparisons needed to retrieve K nearest neighbors for $M = 1000, 2000, 4000, 8000,$ and $14,000$ users. For each value of M and K , the experiments were repeated 1,000 times for different active users. The number of comparisons as a function of K for the above values of M is shown in Fig. 5. The horizontal axis stands for K , the number of nearest neighbors to be retrieved, whereas the vertical reflects the average number of comparisons during the KNN search.

As can be clearly seen, the number of comparisons in CAN-based KNN remains roughly unchanged when K increases. This is explained by the observation that most of the KNN users are located in close vicinity to the active user (this characterizes a real-life naturally clustered data). Thus, the similar users are discovered in the early stages of the KNN search, while further expansions contribute very few new similar users.

Both experiments show good scalability of CAN-based KNN with K . This means, that practical Recommender Systems can use higher values of K , to form moderately larger and more reliable neighborhoods, and generate more accurate predictions with only a very minor computational overhead.

5.2. Accuracy of the search

The following experiments were designed to evaluate the accuracy of the results obtained by the proposed heuristic variant of KNN search. In the first experiment we compared the sets of users, i.e. the neighborhoods, retrieved by the traditional (exhaustive) KNN and by the CAN-based variant of KNN.

Let us denote by KNN_e the set of users retrieved by the traditional exhaustive KNN search and by KNN_h the set of users retrieved by the CAN-based heuristic variant of KNN. Since the CAN-based KNN is a heuristic approach, a sub-optimal structure of zones may lead to a situation, where $\text{KNN}_e \neq \text{KNN}_h$, i.e. the heuristic search retrieves only a subset of the real K nearest neighbors. As the collaborative predictions are generated by aggregating the ratings of similar users, identifying the set of most similar users is essential for generating accurate predictions.

To evaluate the accuracy of the proposed heuristic KNN search, we adapt the traditional Information Retrieval metric of precision.¹⁸ In fact, the computed accuracy metric is not a classical precision, but rather precision@ K , since the overall search procedure is limited to K most similar users only. However, this metric also provides some indication about the recall of the search, as it can be considered as the recall of the search for a limited number of the most similar users to be retrieved. For the sake of clarity, this metric is referred to in the paper as precision. The precision is computed by:

$$\text{precision} = \frac{|\text{KNN}_e \cap \text{KNN}_h|}{|\text{KNN}_e|} = \frac{|\text{KNN}_e \cap \text{KNN}_h|}{K}. \quad (3)$$

The cardinality of the KNN_e set was $K = 10$, while the cardinality of the KNN_h set was gradually increased from $K' = 1$ to $K' = 100$. The precision was computed for $M = 1000, 2000, 4000, 8000$ and $14,000$ users inserted into the system. For each value of M and K' , the experiments were repeated 1000 times for different active users. Figure 6 shows the precision as a function of K' for the above values of M . The horizontal axis stands for M , the number of users inserted into the system, whereas the vertical reflects the average precision of the heuristic KNN search.




Fig. 6. Precision of CAN-based KNN.

As can be clearly seen, the curves behave similarly and the accuracy increases with K' , such that for $K' > 50$, it is over 0.9 for all the given values of M . Previous experiments presented in previous subsection show that the algorithm is highly scalable with K . Thus, retrieving a larger set of users (i.e. higher values of K') leads to a minor increase in the computational overhead. Hence, it is feasible to moderately increase the number of neighbors retrieved by CAN-based search in order to achieve a higher accuracy and generate better predictions.

Since the precision of the heuristic CAN-based KNN search may seem low for small values of K' , we conducted another two experiments, aimed at evaluating the quality of the neighborhood retrieved by the heuristic search. In the first, this was done by computing the average similarity between the nearest neighbors retrieved by the heuristic search and the active user. The computed average similarity was compared to the average similarity of neighborhood retrieved by the traditional search.

In the experiment, we gradually increased the number of users inserted into the system from $M = 1000$ to $M = 14,000$. For each value of M , we compared the average similarity of heuristically retrieved neighbors with the average similarity of exhaustively retrieved neighbors for $K = K' = 10$. For each value of M , the above experiments were repeated 1000 times for different active users. The results of the experiment are shown in Fig. 7 (they are discussed after Fig. 8). The horizontal axis stands for the number of users inserted into the system, whereas the vertical reflects the average similarity value between the users in KNN set and the active user for both exhaustive and heuristic searches.

The second experiment was designed to evaluate the quality of the heuristically retrieved neighborhood by comparing the accuracy of the generated predictions. The final goal of the KNN search is to retrieve a set of the most similar users, whose ratings will be aggregated when generating the predictions. Thus, we generated the




Fig. 7. Average similarity versus the number of users inserted.

predictions using both exhaustively and heuristically retrieved sets of K -Nearest Neighbors and evaluated the accuracy of the predictions using well-known Mean Average Error (MAE) metric⁹:

$$\text{MAE} = \frac{\sum_{i=1}^N |p_i - r_i|}{N} \quad (4)$$

where N denotes the number of predicted items, and p_i is the predicted, and r_i is the real rating on item i .

Also in this experiment the number of users inserted into the system was gradually increased from $M = 1000$ to $M = 14,000$. For each value of M , the experiment was repeated 1000 times for various, randomly chosen active users. For each active user chosen, the following operations were conducted: (1) a single randomly selected rating in the user's profile was hidden and served as a rating to be predicted, while the remaining *all-but-one* ratings served as the user's profile, (2) based on the all-but-one user's profile, the set of $K = K' = 10$ nearest neighbors was retrieved using both traditional exhaustive and heuristic retrievals, (3) predictions were generated using both heuristically and exhaustively retrieved neighborhoods, and (4) the MAE error of the generated predictions relatively to the original hidden rating was computed. The average values of the MAE computed for certain values of M are shown in Fig. 8. The horizontal axis stands for the number of users inserted into the system, whereas the vertical reflects the MAE values for both exhaustive and heuristic searches.

The results show that the average similarity (which is actually the dissimilarity) and the MAE of the predictions decrease with M . This is explained by the observation that the probability of discovering a similar user increases with the number of users inserted into the system. Thus, the average dissimilarity of the retrieved




Fig. 8. Mean average error of the predictions versus the number of users inserted.

K -Nearest Neighbors decreases with M , while the accuracy of the generated predictions increases, and the MAE decreases as well.

Although both the similarity and the MAE of CAN-based heuristic search are higher (i.e. the retrieved neighbors are more dissimilar and the accuracy is actually lower), the curves are very close and the results are quite similar. Average deviation of the similarities is 2.93% and of the MAEs is only 0.38%. Note that the average deviation of the MAE is significantly lower than the average deviation of the similarities, as the generated predictions are barely affected by the changes in the retrieved neighborhoods. These experiments allow us to conclude that the proposed heuristic algorithm succeeds in both retrieving similar neighborhoods and generating accurate predictions.

5.3. Inherent clustering

One of the basic assumptions, that allows us to limit the heuristic search to users, located in close vicinity to the active user, is the *inherent clustering*. That means that the distance between two similar users is lower than the distance between two arbitrary users. Thus, the following experiment was designed to verify the property of inherent clustering in the underlying content-addressable space.

For this, we computed the average and the standard deviation of the similarity of the users located $R = 1, 2$, and 3 routing hops from the active user. The experiments were conducted for $M = 1000, 2000, 4000, 8000$ and 14,000 users inserted into the system. For each value of M , the experiments were repeated 1000 times for different random orders of inserting the users into the system and for different active users. Figure 9 shows the average similarity and the standard deviation as a function of R for the above values of M . The horizontal axis stands for M , the number of users inserted into the system, whereas the vertical reflects the average and the standard




Fig. 9. Average similarity versus number of hops from the active user.

deviation of the similarity of the retrieved users, located within a given number of hops from the active user.

It can be seen that for any given value of M the similarity increases with R . This means that the similarity of users, located close to the active user is higher than the similarity of those located far. Thus, this experiment verifies our assumption on the clustering in content-addressable space. For any R , the average similarity and the standard deviation steadily decrease with M . This observation is explained by the fact that higher number of users leads to a better organization of zones, where zones managed by more similar users *block* the zones managed by dissimilar users. Thus, the average similarity (and the standard deviation) of users located within a given number of hops decreases with R .

Moreover, this experiment demonstrates the stability of the proposed CAN-based structure of users. This experiment was repeated 1000 times, for different random orders of inserting the users into the system. Low values of the standard deviation, and the steady decrease of it with the number of users in the system, show that the inherent clustering holds regardless of the different types of organization of the CAN zones, imposed by the different orders of inserting the users. Thus, we can conclude that the proposed heuristic KNN search will also succeed in retrieving accurate neighborhoods of users for different system usage scenarios.

5.4. Completion heuristics

The following experiments were designed to evaluate the proposed completion heuristics for filling-in the missing values in the incomplete ratings vectors. To run the experiment with the incomplete vectors, we used the full Jester dataset.⁶ In previous experiments we used a partial dataset of complete vectors, built by 14,192 users that rated all 100 jokes. In addition, the full dataset also contains the ratings of 59,229 users that rated on average 45.26 jokes. The full Jester dataset

(i.e. the dataset containing both complete and incomplete vectors) was used in the completion heuristics experiments.

We implemented the *user-average* and the *item-average* heuristics that were discussed in Sec. 4.3. As for the *conditional* heuristic, the decision regarding the chosen completion heuristic was based on the number of explicitly rated items in user's ratings vector. Since in the full Jester dataset the average number of items rated by a user was 45.26, in our implementation of the conditional heuristic the threshold for choosing an appropriate heuristic was set to 20 items. That means that if a user rated less than 20 items, his/her ratings vector is not considered as a reliable one, and vector completion exploits the *item-average* heuristic, which substitutes each missing rating with the average rating of the other users on the given item. However, if a user rated 20 items or more, the *user-average* heuristic is exploited, the missing ratings are substituted with the average rating of the given user on the other items.

To evaluate the accuracy of the proposed three completion heuristics, we conducted two types of experiments. In the first, we compared the average similarity value between the active user and the K -Nearest Neighbors retrieved by the heuristic search and by the traditional exhaustive search. The experiment was repeated three times, for the different completion heuristics exploited before inserting the completed vectors to the underlying content-addressable space.

In the experiment, we gradually increased the number of users inserted into the system from $M = 5000$ to $M = 50,000$. For each value of M , we compared the average similarity of the retrieved neighbors (using both exhaustive and heuristic retrieval techniques) for $K = K' = 10$. For each value of M , the above experiments were repeated 1000 times for different active users. The results of the experiment are shown in Fig. 10. The horizontal axis stands for M , the number of users inserted into the system, whereas the vertical reflects the average similarity value between the users in KNN set and the active user for both exhaustive and heuristic searches. Note that the heuristic retrieval was conducted three times, according to the slightly different datasets inserted into the content-addressable space, as imposed by the completion heuristics being exploited.

The curves show, that similarly to the accuracy results in previous subsections, the average similarity (i.e. dissimilarity) of the retrieved KNN users decreases with M , the number of users inserted into the system. Comparison of the proposed completion heuristics yields that the personalized *user-average* heuristic outperforms the nonpersonalized *item-average* heuristic. Average similarity deviation of the KNN set exploiting the *user-average* heuristic from the exhaustively retrieved KNN is 4.43%, while the similarity deviation of the *item-average* KNN set is 6.21%. Since the *conditional* heuristic is a smarter combination of the above heuristics, it slightly outperforms the *user-average* heuristic as well, and for it the average similarity deviation from the exhaustively retrieved KNN set is 4.11%.

Since the ultimate goal of the Collaborative Filtering is to generate predictions, the second experiment was designed to evaluate the quality of the completion




Fig. 10. Average similarity versus the number of users inserted.

heuristics by comparing the accuracy of the generated predictions. To do this, we generated the predictions using both exhaustively and heuristically retrieved sets of K -Nearest Neighbors and evaluated the accuracy of the predictions using the MAE metric. In the experiment the number of users inserted into the system was gradually increased from $M = 5000$ to $M = 50,000$. For each value of M , the experiment was repeated 1000 times for various, randomly chosen active users. The experimental setting was similar to one described in previous subsections: the user's profile was partitioned to the predicted rating and all-but-one profile, the sets of $K = K' = 10$ nearest neighbors were retrieved using both exhaustive and heuristic retrievals, the predictions were generated using both neighborhoods, and the MAE of the generated predictions relatively to the original rating was computed. The average values of the MAE are shown in Fig. 11. The horizontal axis stands for the number of users inserted into the system, while the vertical reflects the MAE values for both exhaustive and heuristic searches. Note that the heuristic retrieval was conducted three times, according to the completion heuristics being exploited.

Similarly to the results in Sec. 5.2, this experiment shows that the MAE of the prediction decreases with M , the number of users inserted into the system. Comparison of the proposed completion heuristics yields that the accuracy of the predictions exploiting personalized *user-average* heuristic is better than that of the nonpersonalized *item-average* heuristic. However, for both heuristics the average increase of the MAE values is minor: for the *user-average* heuristic it is 0.69%, whereas for the *item-average* heuristic it is 1.37%. As can be seen from the chart, also in this experiment the *conditional* heuristic slightly outperforms both of them, as for the *conditional* heuristic the increase of the MAE is only 0.46%. Hence, out of the proposed three completion heuristics, the *conditional* heuristic retrieves the most similar KNN set and generates the most accurate prediction. This allows us to conclude that this heuristic should be used for converting the incomplete vectors to




Fig. 11. Mean average error of the predictions versus the number of users inserted.

the complete ones, and naturally leads to future research, dealing with developing more accurate completion heuristics.

6. Conclusions and Future Research

One of the major drawbacks of the state-of-the-art CF implementations is their high computational complexity, which grows linearly both with the number of users and items in the system. In this work we proposed to heuristically decrease the required computational effort by implementing the CF over content-addressable CAN-like N -dimensional space.

6.1. Conclusions

Experiments conducted over the Jester dataset of jokes ratings show that in general the proposed heuristic algorithm outperforms the traditional exhaustive KNN search as the computational overheads are significantly decreased, while the accuracy remains roughly unchanged. Our algorithm decreases the number of required comparisons, while the ratio between the numbers of comparisons steadily decreases with the number of users. For example, for 14,000 users the number of comparisons was decreased by almost an order of magnitude (precisely, by 87%). Other experiments show that the number of comparisons roughly remains unchanged when K increases. This allows us to increase the number of nearest neighbors to be retrieved (and to potentially improve the accuracy of the generated predictions) with a very minor computational overhead.

In the accuracy experiments we qualitatively compared the neighborhoods retrieved and the predictions generated by the CAN-based heuristic and by the traditional exhaustive KNN searches. The retrieved neighborhoods were similar and the predictions were very close, which indicates good accuracy of the proposed

algorithm. In summary, comparing the proposed heuristic KNN search with traditional exhaustive search shows that our algorithm achieves high accuracy (very similar to the accuracy of the traditional exhaustive KNN search), while significantly decreasing the required computational effort.

Another set of experiments were aimed at validating the inherent clustering property of content-addressable spaces. The results showed that this property holds in the CAN-like space, as the dissimilarity of users, located in a certain number of network hops from the active user increased with the number of network hops. The experiments also showed that the inherent clustering property holds regardless of the number of users inserted into the system and the order of their insertion.

The last set of experiments were aimed at comparing three heuristic for converting the incomplete vectors to complete ones by filling-in the missing ratings. Three simple heuristics were compared: two heuristics that substitute the missing ratings either with the average rating of the given user, or with the average rating on the given item, whereas the third heuristic integrates the first two. The experiments showed that the heuristic, which conditionally integrates two other heuristics, outperforms them both in terms of the retrieved neighborhoods' similarity and of the generated predictions' accuracy.

Comparing the MAE of the predictions generated by the complete and heuristically completed vectors yields that the accuracy of the predictions generated by the complete vectors is slightly better. This conclusion is reasonable, since the proposed completion heuristics insert some extent of noise into the original ratings. However, the increase in the MAE is minor, allowing us to conclude that the achieved computational optimization is preferential than the minor noises in the generated predictions caused by the artificial ratings inserted by the completion heuristics.

6.2. Future research

In this work, we inherently assumed that the system assigns equal relative weights to the ratings on each item. However, this assumption is not true in many real-life personalization applications. For example, this assumption might be false in a situation, where different criteria affect differently on the similarity values, e.g. when the similarity values between the items are known. Developing a weighted prediction algorithm will result in a more accurate Recommender System.

Also, we assumed that either the user's ratings on the items are available or they can be easily filled-in using one of the proposed simple completion heuristics. However, in some real-life scenarios, this completion is hard to achieve, since the matrix is very sparse (e.g. density of 2–3% in typical Collaborative Filtering datasets such as in Refs. 9 and 10) and the substitution of the missing values may require exploiting more intelligent techniques. In the future, we plan to study the use of various completion heuristics, exploiting statistical and Machine Learning techniques.

In addition to decreasing the computational effort, the proposed algorithm can naturally be extended to distribute it among multiple users. In traditional centralized implementations of the CF, the Similarity Computation and the Neighborhood Formation stages are performed in a single central location. However, as the underlying CAN platform is originally distributed Peer-to-Peer platform, it inherently allows distributed and fully decentralized storage of the ratings matrix. In the future, we plan to implement a distributed variant of the algorithm and to investigate the distribution issues.

The current work is limited to the Mean Squared Difference (MSD) similarity metric, since the injective mapping to a multidimensional CAN-like space inherently supports it. However, for other metrics, such as Cosine Similarity or Pearson Correlation, CAN space might be inappropriate and new types of topologies and respective mappings should be developed. We plan to study other metrics and to produce a general framework for efficient heuristic Collaborative Filtering.

Acknowledgments

The authors gratefully acknowledge the support of the Caesarea Edmond Benjamin de Rothschild Foundation Institute for Interdisciplinary Applications of Computer Science (CRI) and the Haifa Interdisciplinary Research Center for Advanced Computer Science (HIACS), both at the University of Haifa.

References

1. S. Aguzzoli, P. Avesani and P. Massa, Collaborative case-based recommender system, *Proc. ECCBR Conf.* (1997).
2. S. Bogaerts and D. Leake, Facilitating CBR for incompletely-described cases: distance metrics for partial problem descriptions, *Proc. ECCBR Conf.* (2004).
3. J. Breese, D. Heckerman and C. Kadie, Empirical analysis of predictive algorithms for collaborative filtering, *Proc. UAI Conf.* (1998).
4. R. Burke, Hybrid recommender systems: survey and experiments, *User Modeling and User-Adapted Interaction* **12**(4) (2002) 331–370.
5. S. H. S. Chee, J. Han and K. Wang, RecTree: an efficient collaborative filtering method, *Proc. DaWaK Conf.* (2001).
6. K. Goldberg, T. Roeder, D. Gupta and C. Perkins, Eigentaste: a constant time collaborative filtering algorithm, *Inform. Retr. J.* **4**(2) (2001) 133–151.
7. N. Good, J. B. Schafer, J. A. Konstan, A. Borchers, B. Sarwar, J. Herlocker and J. Riedl, Combining collaborative filtering with personal agents for better recommendations, *Proc. AAAI Conf.* (1999).
8. P. Han, B. Xie, F. Yang and R. Shen, A scalable P2P recommender system based on distributed collaborative filtering, *Exp. Syst. Appl. J.* **27**(2) (2004) 203–210.
9. J. L. Herlocker, J. A. Konstan, A. Borchers and J. Riedl, An algorithmic framework for performing collaborative filtering, *Proc. SIGIR Conf.* (1999).
10. P. McJones, Eachmovie collaborative filtering data set (1997), available online at <http://research.compaq.com/SRC/eachmovie/>.
11. B. N. Miller, J. A. Konstan and J. Riedl, PocketLens: toward a personal recommender system, *ACM Trans. Inform. Syst.* **22**(3) (2004) 437–476.

12. M. Morita and Y. Shinoda, Information filtering based on user behavior analysis and best match retrieval, *Proc. SIGIR Conf.* (1994).
13. D. M. Pennock, E. Horvitz and C. L. Giles, Social choice theory and recommender systems: analysis of the axiomatic foundations of collaborative filtering, *Proc. AAAI Conf.* (2000).
14. C. Plaxton, R. Rajaraman and A. Richa, Accessing nearby copies of replicated objects in a distributed environment, *Proc. ACM SPAA Conf.* (1997).
15. S. Ratnasamy, P. Francis, M. Handley, R. Karp and S. Shenker, A scalable content-addressable network, *Proc. Conf.* (2001).
16. P. Resnick and H. R. Varian, Recommender systems, *Commun. ACM* **40**(3) (1997).
17. F. Ricci, A. Venturini, D. Cavada, N. Mirzadeh, D. Blaas and M. Nones, Product recommendation with interactive query management and twofold similarity, *Proc. ICCBR Conf.* (2003).
18. G. Salton and M. McGill, *Introduction to Modern Information Retrieval* (McGraw-Hill Publishers, 1983).
19. B. Sarwar, G. Karypis, J. Konstan and J. Riedl, Analysis of recommendation algorithms for e-commerce, *Proc. EC Conf.* (2000).
20. B. M. Sarwar, J. A. Konstan and J. Riedl, Distributed recommender systems: new opportunities for internet commerce, in *Internet Commerce and Software Agents: Cases, Technologies and Opportunities* (Idea Group Publishers, 2001).
21. U. Shardanand and P. Maes, Social information filtering: algorithms for automating “word of mouth”, *Proc. CHI Conf.* (1995).
22. A. Tveit, Peer-to-Peer based recommendations for mobile commerce, *Proc. WMC Workshop* (2001).



Shlomo Berkovsky is a Ph.D. student in the Computer Science and Management Information Systems Departments at the University of Haifa. He received his M.Sc. degree from the Computer Science Department of the Uni-

versity of Haifa in 2004 in the area of semantic data management in peer-to-peer networks. He has published over 20 papers in scientific journals and conferences.

His research focuses on user models mediation for a better personalization in recommender systems.



Yaniv Eytani is a Ph.D. candidate in the Computer Science Department at the University of Illinois at Urbana-Champaign and is involved in research under the direction of Prof. Grigore Rosu. He is a member of the Formal Systems Laboratory. He received an M.Sc. degree from the Computer Science Department of the University of Haifa in 2005 in the field finding concurrent bugs in Java. He has published about 20 papers in scientific journals and conferences.



Larry M. Manevitz

received his B.Sc. degree from Brooklyn College, and his M.Phil. and Ph.D. degrees from Yale University, Mathematics Department in Mathematical Logic. He has held positions at Hebrew University, Bar-

Ilan University, Oxford University, NASA Ames, University of Texas at Austin, Baruch College, CUNY, University of Maryland and University of Wisconsin. Currently he is in the Department of Computer Science, Univ. Haifa, and is the Director of the HIACS Research Laboratory and the Neurocomputation Laboratory there. He has about 60 scientific publications.

He currently specializes in artificial neural networks, machine learning, and brain modeling and has done substantial work in both theoretical and applied mathematical logic (especially non-standard analysis) and in the theory of combining uncertain information.