Collaborative Filtering for Learning User Preferences for Robotic Tasks

Nichola Abdo*, Luciano Spinello*, Cyrill Stachniss†, Wolfram Burgard*
*Institute of Computer Science, University of Freiburg, Germany
†Institute of Geodesy and Geoinformation, University of Bonn, Germany

***Abstract*—Service robots are envisioned to have an increasing influence on our lives and to support us on a daily basis. From a truly effective and personalized robot, we expect the ability to learn our preferences concerning the requested tasks. Such preferences, however, are often user-dependent so that predefined strategies only match a subset of all users. In this work, we address the problem of tailoring the robot's behavior to the preferences of its user. We present a novel solution to the problem of encoding multiple preferences for individual tasks that leverages the collaborative filtering framework. A key aspect of our method is that it does not require each user to specify preferences for all tasks. From a small number of known preferences, our approach is able to infer the user's taste for other tasks. We present quantitative results based on crowdsourced data from thousands of users. Our results suggest the validity of our approach and demonstrate that we are able to predict user preferences with respect to two service robot scenarios.**

I. Introduction

One of the goals in robotics research is to develop autonomous service robots that assist humans in domestic environments. We envision robots to undertake a variety of tasks like tidying up and attending to the needs of disabled people. As robots get more and more capable of performing such tasks, there is a growing need to take the personal preferences of users into account [11, 12]. Learning user preferences, however, is a non-trivial problem. In a home scenario for example, each user probably has a preferred way of sorting and storing various items. Many of our preferences stem from factors like personal taste, cultural background, activities that currently take place in the environment, or common sense. Such factors are hard to formulate or model a priori. At the same time, it is highly impractical to query the users about their preferences for all tasks assigned to the robot.

In this work, we present a novel approach to the problem of inferring user preferences that borrows ideas from data mining and recommender systems. Our solution is based on collaborative filtering, a successful theoretical framework for learning user preferences in a wide variety of applications such as suggesting movies on Netflix or products on Amazon. We adapt the classical formulation for learning relations between buyers and items to learning relations between users and task preferences. By leveraging this theory, we formulate an active learning system that allows for encoding multiple user preferences for individual tasks. The representation is easy to update and offers the possibility for lifelong learning and improvement. Finally, our method does not require each user to specify their own preferences for all the tasks.

Our approach proceeds in two phases. First, we collect many user preferences using crowdsourced surveys in an offline learning step. This is used to build a model for tastes related to the considered tasks. In the second, active learning phase, the robot queries the user about a few preferences and is able to predict them for all the remaining tasks. We consider two application scenarios that a service robot might be confronted with: arranging grocery items on shelves and scheduling tidying-up tasks according to their relative importance. We collected preferences from over 2,000 surveys and present results that show that our method is able to accurately predict the taste of the users after only a few queries.

II. Related Work

Recent advances in perception and manipulation have allowed service robots to perform a variety of tasks related to tidying up and cleaning [10, 14, 5]. However, for building truly personal robots that operate in domestic environments, we need a better understanding of how end-users prefer their robots to attend to different everyday tasks. Several researchers highlight the differences in people's expectations and preferences regarding tasks performed by service robots [3, 11, 12]. The work of Koppula and Saxena [6] about anticipating human activities can be regarded as a step in this direction.

In this work, we present an approach for predicting user preferences for service robot applications based on collaborative filtering, a successful approach used in recommender systems [4]. Such systems are widely used for making personalized recommendations of products to users, e.g., Netflix and Amazon. A popular approach to collaborative filtering is based on factorization techniques [7, 13]. The main idea is to discover latent patterns or tastes in a ratings matrix through techniques such as singular value decomposition.

Recently, collaborative filtering techniques have been applied in the context of action recognition in videos and selecting good floor coverage strategies to a vacuum-cleaning robot [8, 9]. In contrast to such use cases, we aim at predicting subjective human user preferences. We therefore use crowdsourcing to acquire a large number of individual user preferences and construct a ratings matrix. To the best of our knowledge, our work is the first to apply a combination of crowdsourcing and collaborative filtering to learn and predict task preferences for service robots.

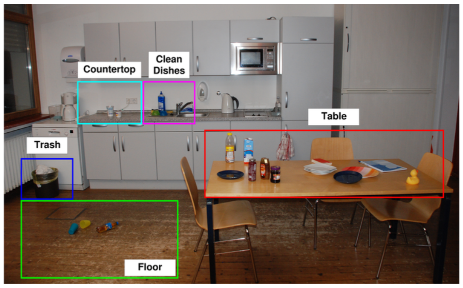

Fig. 1. Example of a kitchen scene that has been presented to the contributors for ranking the urgency of tidying-up tasks given the scene.

III. Collaborative Filtering for Learning Task Preferences from Users

In this section, we formulate the problem of learning and predicting user preferences using the framework of collaborative filtering. This is grounded on the idea that user preferences follow latent patterns such that people with a similar taste have similar expectations about how the robot should solve a task. We assume that such tasks have been previously learned by the robot or pre-programmed by an expert.

A. Learning and Predicting User Preferences

Our collaborative filtering system makes use of an $M \times N$ ratings matrix $\mathbf{R}$ that is composed of ratings for M different tasks $\mathcal{I} = \{i_i, i_2, \ldots, i_M\}$ given by N users $\mathcal{U} = \{u_1, u_2, \ldots, u_N\}$. Each column of $\mathbf{R}$ corresponds to the vector of the ratings made by one user. Since most users only rate a small subset of all tasks, $\mathbf{R}$ is typically incomplete.

We follow the formulation and factorization approach of Bell and Koren [2, 7] and express each element of $\mathbf{R}$ as

$$r_{iu} = b_{iu} + \overline{r}_{iu}. \tag{1}$$

Here, r_{iu} refers to the rating for task i given by user u. The term b_{iu} represents a (baseline) bias rating and $\overline{r}_{iu}$ a residual rating, i.e., the deviation of the user rating from the bias, which we aim to predict. All residual ratings $\overline{r}_{iu}$ are collected in the residual ratings matrix $\overline{\mathbf{R}}$. Each b_{iu} is expressed as the sum of a global, task-specific, and user-specific factor

$$b_{iu} = \mu + \phi_i + \psi_u, \tag{2}$$

where μ is the global mean rating over all entries in $\mathbf{R}$. The term ϕ_i is specific for the task i and ψ_u is the user-specific term for user u. They respectively capture the average rating a task tends to get or a user tends to give, i.e.,

$$\begin{aligned} \phi'_i &= \frac{\sum_u (r_{iu} - \mu)}{N_i}, \quad \psi'_u = \frac{\sum_i (r_{iu} - \mu)}{M_u}, \\ \phi_i &= \frac{\sum_u (r_{iu} - \mu - \psi'_u)}{N_i}, \quad \psi_u = \frac{\sum_i (r_{iu} - \mu - \phi_i)}{M_u}, \end{aligned} \tag{3}$$

where N_i is the number of users who have rated tasks i, and M_u is the number of tasks rated by user u.

We refer to tasks already rated by user v as the *probe tasks*, $\mathcal{I}_v \subset \mathcal{I}$, and their corresponding ratings as the *probe ratings* $\mathbf{r}_v$. The aim of a collaborative filtering system is to predict the rating r_{jv} for a task j by a user v who did not rate this task before. To do so, we factor residual ratings $\overline{r}_{iu}$ in $\overline{\mathbf{R}}$ as the dot product of two K-dimensional factor vectors

$$\overline{r}_{iu} = \mathbf{p}_i^T \cdot \mathbf{q}_u = \sum_{k=1}^{K} (p_{ki} q_{ku}), \tag{4}$$

where $\mathbf{p}_i$ and $\mathbf{q}_u$ are respectively low-dimensional representations of task i and user u. We use $\mathbf{P}^T$ and $\mathbf{Q}$ to denote the $M \times K$ and $K \times N$ matrices whose columns are made up of all $\mathbf{p}_i$ and $\mathbf{q}_u$ vectors, respectively. To compute $\mathbf{p}_i$ and $\mathbf{q}_u$ for all $i \in \mathcal{I}$ and $u \in \mathcal{U}$, we solve an optimization problem aiming at minimizing the sum of the squared errors $e_{iu} = \overline{r}_{iu} - \mathbf{p}_i^T \cdot \mathbf{q}_u$ over all known residual ratings, i.e.,

$$\operatorname*{argmin}_{\mathbf{P},\mathbf{Q}} \sum_{i \in \mathcal{I}, u \in \mathcal{U}} (\overline{r}_{iu} - \mathbf{p}_i^T \cdot \mathbf{q}_u)^2 + \frac{\lambda}{2}(\|\mathbf{P}\|^2 + \|\mathbf{Q}\|^2), \tag{5}$$

where λ is a regularizer. After computing $\mathbf{P}$ and $\mathbf{Q}$, we can predict the residual rating $\hat{\overline{r}}_{jv}$, i.e., the predicted deviation of the rating of user v about task j from the bias as

$$\hat{\overline{r}}_{jv} = \mathbf{p}_j^T \cdot \mathbf{q}_v. \tag{6}$$

Finally, we obtain the predicted rating $\hat{r}_{jv}$ by adding the corresponding biases, i.e., $\hat{r}_{jv} = \mu + \phi_j + \psi_v + \hat{\overline{r}}_{jv}$.

B. Active Collaborative Filtering for Robotics Tasks

After building a ratings matrix $\mathbf{R}$ from the ratings of different users, the aim of the robot is to attend to tasks with respect to a new user or environment. Initially, the robot has no knowledge about the new user (the so-called "cold-start" problem in the collaborative filtering literature). Therefore, in the second online stage, it actively *asks/probes* the user about a few random tasks. Using these probe ratings, the robot uses $\mathbf{R}$ as described in Sec. III-A to predict all remaining task ratings for the new user.

Our collaborative filtering formulation allows us to easily incorporate new information about users. For adding a new user, the ratings matrix $\mathbf{R}$ is augmented with an additional column. For an existing user, only a single element in $\mathbf{R}$ must be updated. This increases the knowledge of the robot about the user's tastes and allows for active online learning. Without complicated modeling, new patterns of preferences can be created or less granular tastes can be encoded. In the long run, the more the users contribute to the database, the finer the encoded spectrum of preferences will be.

IV. Experimental Evaluation

We evaluated our approach on two service robot scenarios. For both scenarios, we gathered user preferences through surveys conducted using CrowdFlower [1]. In scenario 1, we considered a butler robot that has to arrange grocery items on different shelves. This is an example where spatial arrangements of objects are motivated by complex factors like user taste and object types so that manually designing a set of hand-crafted rules is impractical and unlikely to match with the tastes of all users. We defined a task as placing an item of

TABLE I
GROCERY SCENARIO DETAILED EVALUATION WITH $P = 12$ PROBES

		no	maybe	yes
Baseline 1	Precision	59.5%	34.0%	69.2%
	Recall	9.3%	95.0%	2.7%
	F-score	0.16	0.50	0.05
Baseline 2	Precision	41.7%	32.5%	19.8%
	Recall	22.5%	46.4%	22.2%
	F-score	0.29	0.38	0.21
CF	Precision	67.1%	45.3%	58.1%
	Recall	62.9%	55.2%	38.2%
	F-score	0.65	0.50	0.46

TABLE II
KITCHEN SCENARIO, DETAILED EVALUATION WITH $P = 6$ PROBES

		no	maybe	yes	yes urgently
Baseline 1	Precision	0%	24.45%	34.05%	64.06%
	Recall	0%	35.38%	71.10%	29.79%
	F-score	0	0.29	0.46	0.41
Baseline 2	Precision	15.3%	24.9%	27.2%	32.7%
	Recall	16.7%	33.6%	32.5%	15.6%
	F-score	0.16	0.29	0.30	0.21
CF	Precision	28.5%	75.2%	83.2%	85.3%
	Recall	85.5%	43.8%	40.3%	41.5%
	F-score	0.43	0.55	0.54	0.56

one type next to another item on the same shelf. Overall, we gathered ratings from 1,064 users on 100 different tasks by posing questions like '*Would you place item A next to item B on the same shelf?*', where we randomly sampled the two types from a set of 20 common grocery items (bread, cereals, spices, etc). Each question could be answered with no, maybe, or yes, corresponding to ratings of 0, 1, and 2. The resulting ratings matrix has only around 18% of its entries filled with ratings, and each task was rated by 190 users on average.

In scenario 2, we considered a robot tidying up a kitchen environment with four main locations: table, countertop, floor, and sink. We considered three object categories: kitchenware, food, and miscellaneous items (e.g. books, toys, etc). We specified 15 tasks the robot can perform: remove dirty kitchenware / food / miscellaneous items from the table / floor / countertop, wipe the table / countertop, vacuum-clean the floor, set the table for dinner, clear away the clean dishes from the sink, and take out the trash. We created 200 kitchen scene images with random numbers of objects at the different locations (see Fig. 1 for an example). Each contributor was shown an image and asked to rate the tasks in terms of urgency as follows: no (0), if possible (1), yes (2), and yes urgently (3). We collected the ratings in a matrix with 15 task rows and 1,100 user columns, with around 60% of the entries filled with ratings.

A. Predicting User Preferences

We evaluated the ability of our approach to predict preferences in both scenarios by conducting a leave-one-out cross-validation. We randomly queried a user about their ratings for P tasks and predicted all the other tasks rated by that user. We rounded the predicted ratings to the closest integer to get a value from the rating scale of each scenario. We compared our predictions to the known user ratings and evaluated the performance by computing precision and recall values as well as the F-score. We also compared our results to two baseline approaches. *Baseline 1* predicts a preference of a task using the mean rating for that task over all users who have rated it. *Baseline 2* predicts a preference by randomly selecting a value from the rating scale.

The results for scenario 1 are shown in Tab. I for $P = 12$. Here, our collaborative filtering approach (CF) largely outperforms the baselines with an F-score of 0.53 on average versus 0.23 and 0.29 for the baseline approaches 1 and 2, respectively. Note that baseline 1 is only able to obtain an high recall in the maybe case and produces false negatives in the other two ratings. This is due to the tendency of this baseline to predict the preferences as maybe due to the highly tri-modal preference distribution in most tasks. On the other hand, despite partial, multimodal ratings, our method is able to obtain high recall and precision on all cases. In particular, we are able to correctly predict significantly more cases where users explicitly rated tasks as no or yes.

The results for scenario 2 are shown in Tab. II for $P = 6$. With six known ratings about the users, our approach is able to outperform both baselines with an average F-score of 0.52 over all rating classes on the scale. Notice the low recall value for baseline 1 in predicting tasks as no. This is due to the fact that most tasks received high ratings in terms of urgency.

B. Influence of the Number of Probe Ratings

Furthermore, we evaluated the performance of our approach with respect to the number of probe ratings for $P = 2 \dots 10$ and $P = 4 \dots 20$ for the first and second scenarios, respectively. The results are shown in Fig. 2 and Fig. 3. As expected, and as opposed to the baselines, the performance of our approach improves given more probe ratings, which is in line with the ability of collaborative filtering to use knowledge about users to predict their preferences. This highlights the fact that it is suboptimal to rely on heuristics like the baselines for predicting preferences that match all users.

V. CONCLUSION

We presented an approach that utilizes collaborative filtering for encoding robotic task preferences. Our technique leverages active learning, is easy to update, and is able to encode multiple tastes for each task. To the best of our knowledge, this is the first time that all these objectives are simultaneously achieved in a robotics context. We performed extensive experiments in the context of two typical service robotic applications. For training, we collected thousands of user preferences using a crowdsourcing platform. Experimental results demonstrate that our method is able to reliably estimate task preferences after querying the user only a few times.

VI. ACKNOWLEDGEMENTS

This work has partly been supported by the German Research Foundation under research unit FOR 1513 (HYBRIS).

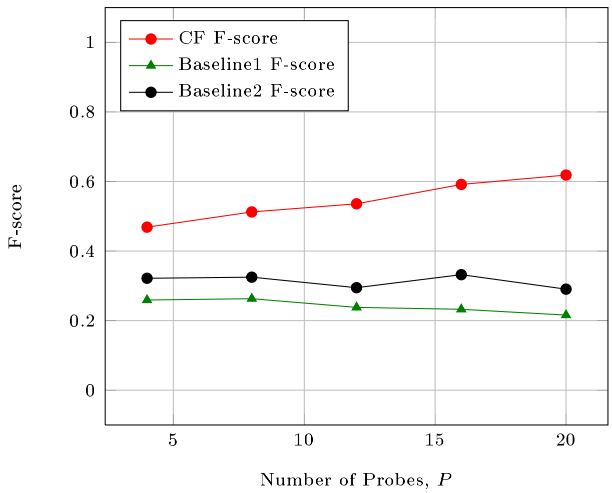

Fig. 2. F-score evaluation with respect to number of probes in the scenario of organizing grocery items. The scores are averaged over all rating categories: no, maybe, and yes. The performance of our method (CF) improves with more knowledge about how users have previously rated tasks. This is in contrast to the baselines that use the average ratings over users.

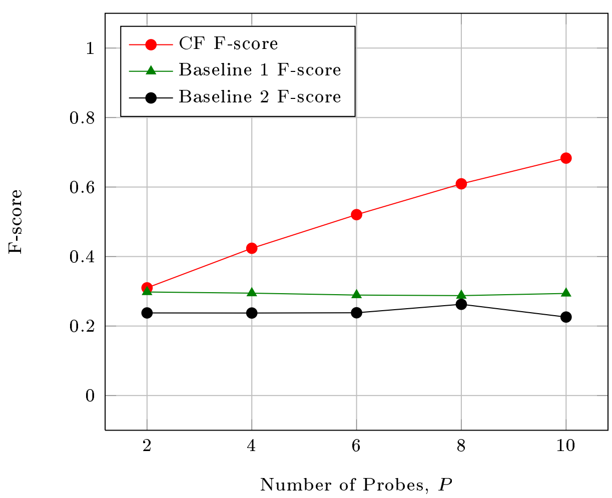

Fig. 3. F-score evaluation with respect to number of probes in the kitchen scenario. The scores are averaged over all rating categories: no, maybe, yes, and yes urgently. As in the groceries scenario, the performance of our method (CF) improves with more knowledge about the users. This is in contrast to the behaviour of both baselines.

References

[1] Crowdflower crowdsourcing service. URL http://crowdflower.com/.

[2] Robert M. Bell and Yehuda Koren. Improved neighborhood-based collaborative filtering. *KDD-Cup and Workshop*, 2007.

[3] Maya Cakmak and Leila Takayama. Towards a comprehensive chore list for robots. In *Int. Conf. on Human-Robot Interaction (HRI)*, 2013.

[4] David Goldberg, David Nichols, Brian M. Oki, and Douglas Terry. Using collaborative filtering to weave an information tapestry. *Comm. of the ACM*, 1992.

[5] Jürgen Hess, Gian D. Tipaldi, and Wolfram Burgard. Null space optimization for effective coverage of 3D surfaces using redundant manipulators. In *Int. Conf. on Intelligent Robots and Systems*, 2012.

[6] Hema Koppula and Ashutosh Saxena. Anticipating human activities using object affordances for reactive robotic response. In *Proc. of Robotics: Science and Systems*, 2013.

[7] Yehuda Koren. Factor in the neighbors: Scalable and accurate collaborative filtering. *ACM Trans. on Knowledge Disc. from Data (TKDD)*, 2010.

[8] Pyry Matikainen, Rahul Sukthankar, and Martial Hebert. Model recommendation for action recognition. In *Int. Conf. on Computer Vision and Patt. Recognition (CVPR)*, 2012.

[9] Pyry Matikainen, Padraig Michael Furlong, Rahul Sukthankar, and Martial Hebert. Multi-armed recommendation bandits for selecting state machine policies for robotic systems. In *Int. Conf. on Robotics & Automation*, 2013.

[10] Stephen Miller, Jur Van Den Berg, Mario Fritz, Trevor Darrell, Ken Goldberg, and Pieter Abbeel. A geometric approach to robotic laundry folding. *Int. J. of Robotics Research*, 2012.

[11] Caroline Pantofaru, Leila Takayama, Tully Foote, and Bianca Soto. Exploring the role of robots in home organization. In *Int. Conf. on Human-Robot Interaction (HRI)*, 2012.

[12] Celine Ray, Franceso Mondada, and Roland Siegwart. What do people expect from robots? In *Int. Conf. on Intelligent Robots and Systems*, 2008.

[13] Jasson D. M. Rennie and Nathan Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Int. Conf. on Machine Learning (ICML)*, 2005.

[14] Ashutosh Saxena, Justin Driemeyer, and Andrew Y. Ng. Robotic grasping of novel objects using vision. *Int. J. of Robotics Research*, 2008.