Expert Systems With Applications 61 (2016) 246–261

Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

Reducing preference elicitation in group decision making

Lihi Naamani-Dery[a,*], Meir Kalech[b], Lior Rokach[b], Bracha Shapira[b]

[a] Industrial Engineering and Management, Ariel University, Ariel 40700, Israel
[b] Information Systems Engineering, Ben-Gurion University of the Negev, P.O.Box 653 Beer-Sheva 8410501, Israel

ARTICLE INFO

Article history:
Received 14 December 2015
Revised 9 May 2016
Accepted 27 May 2016
Available online 30 May 2016

Keywords:
Preference elicitation
Group decision making
Computational social choice

ABSTRACT

Groups may need assistance in reaching a joint decision. Elections can reveal the winning item, but this means the group members need to vote on, or at least consider all available items. Our challenge is to minimize the amount of preferences that need to be elicited and thus reduce the effort required from the group members. We present a model that offers a few innovations. First, rather than offering a single winner, we propose to offer the group the best top-*k* alternatives. This can be beneficial if a certain item suddenly becomes unavailable, or if the group wishes to choose manually from a few selected items. Secondly, rather than offering a definite winning item, we suggest to approximate the item or the top-*k* items that best suit the group, according to a predefined confidence level. We study the tradeoff between the accuracy of the proposed winner item and the amount of preference elicitation required. Lastly, we offer to consider different preference aggregation strategies. These strategies differ in their emphasis: towards the individual users *(Least Misery Strategy)* or towards the majority of the group *(Majority Based Strategy)*. We evaluate our findings on data collected in a user study as well as on real world and simulated datasets and show that selecting the suitable aggregation strategy and relaxing the termination condition can reduce communication cost up to 90%. Furthermore, the commonly used Majority strategy does not always outperform the Least Misery strategy. Addressing these three challenges contributes to the minimization of preference elicitation in expert systems.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

A group of people wishing to reach a joint decision faces the task of selecting the alternative that best suits the group out of all available candidate items. When all users' preferences are known, some voting aggregation strategy is used to compute and output the winning item to the group (Rossi, Venable, & Walsh, 2011). When the preferences are not available, a preference elicitation process is required.

Preference elicitation requires time and effort, so our goal is to stop the elicitation as soon as possible. In the worst case, for most voting protocols all the preferences are needed in order to determine a winning item, i.e., an item that most certainly suits the group's joint preferences (Conitzer & Sandholm, 2005). Nevertheless, in practice it has been shown that the required information can be cut in more than 50% (Kalech, Kraus, Kaminka, & Goldman, 2011; Lu & Boutilier, 2011). Given partial preferences, it is possible to define the set of the necessary winners, i.e., items which must necessarily win, as well as the set of possible winners, i.e., items which can still possibly win (Konczak & Lang, 2005). Using these definitions the elicitor can determine whether there is need for more information concerning the voters' preferences. Previous studies provide algorithms for preference elicitation of a single winner under the Range and the Borda protocols (Lu & Boutilier, 2011; Naamani-Dery, Golan, Kalech, & Rokach, 2015; Naamani-Dery, Kalceh, Rokach, & Shapira, 2014). In this paper we define two tradeoffs that enable less elicitation: *Selection* and *Approximation*. Furthermore, we propose to examine different preference *Aggregation* techniques.

Selection: a tradeoff exists between the amount of items outputted to the group and the cost of preferences elicitation required. Less elicitation effort is required for outputting k items where one of them is the winner with a high probability (top-k items) than for outputting one necessary winner (i.e., $k = 1$). Although outputting a definite winner is the most accurate result, there are advantages to outputting the top-k items. Not only is the communication cost reduced, it may actually be preferred to present a *few* alternatives to the user since if one of the alternatives is unavailable the group members can quickly switch to another already recommended alternative without requiring more elicitation (Baldiga & Green, 2013; Lu & Boutilier, 2010). Consider, for example, a setting of 30 optional dinner locations for a group. If a fish restaurant is the winning item, but one of the group members dislikes fish,

* Corresponding author.
E-mail addresses: lihid@ariel.ac.il, nlihin@gmail.com (L. Naamani-Dery), kalech@bgu.ac.il (M. Kalech), liorrk@bgu.ac.il (L. Rokach), bshapira@bgu.ac.il (B. Shapira).

http://dx.doi.org/10.1016/j.eswa.2016.05.041
0957-4174/© 2016 Elsevier Ltd. All rights reserved.

the group might prefer to switch to a different alternative rather than to perform another elicitation round.

Approximation: a different tradeoff is the one that exists between the *accuracy* of the proposed winner item and the amount of preference elicitation required. We suggest outputting an item that approximately suits the group with some confidence level rather than outputting an item that definitely suits the group. As we later show, the confidence level is based on the items' winning probabilities. To reduce the elicitation even further, the two methods can be combined and top-k approximate items can be offered to the group. Consider, for example, a group that wishes to choose a movie to watch together out of movies available in the cinema. The members set the amount of options they wish to receive (k) and the level of confidence of the results. Thus, we define a new preference elicitation termination condition: *approximate k-winner termination*, namely where k items are found and one of them is the best item with a confidence level of $1-\alpha$ ($0 \leq \alpha \leq 1$).

Aggregation: Ideally, the preference aggregation strategy (i.e., the voting protocol) should be a fair one. In his well-known work, Arrow shows that there is no perfect aggregation system (Arrow, 1951). One of the major differences between aggregation strategies is the social environment in which they are used; in particular, the perspective in which fairness is viewed. The emphasis can be either towards the individual user or towards the majority of the group (Jameson & Smyth, 2007). Two aggregation strategies that differ in their emphasis and are used in group recommender systems are the *Majority Based Strategy* and the *Least Misery Strategy* (Masthoff, 2011). Similar concepts can be found in the social choice literature, termed utilitarianism and egalitarianism (Myerson, 1981). In the *Majority Based Strategy* the users' ratings of the different items are aggregated and the items with the highest total value are recommended. In the *Least Misery Strategy* the chosen items cannot be the least preferred by any of the users. The idea is that a group is as happy as its least happy member. One of the contributions of this paper is in proposing an efficient iterative preference elicitation algorithm which fits these strategies.

Overall, our goal is to reduce the communication cost in the preference elicitation process. We define the communication cost as the cost of querying one user for her preferences for one item. We allow users to submit the same rating for many items and do not request the users to hold a strict set of preferences over items. In this paper, we adopt the Range voting protocol which is adequate for this purpose; it requires users to submit a score within a certain range. Users are familiar with applications that ask for their score on an item, such as Amazon (www.amazon.com) or Netflix (www.netflix.com).

Preference elicitation becomes more challenging and interesting when a rating distribution of the voter-item preferences exists, i.e., a prior probability distribution of each voter's preferences for each item. For example, in the case of a group of users wish to watch a movie together, the distribution can be inferred from rankings of these movies by similar users using collaborative filtering methods (Koren & Sill, 2011). After each user-item query, new information is revealed. The necessary and possible winner sets are updated to check whether or not the termination condition has been reached.

In this paper, we offer three main innovations contributing to the minimization of preference elicitation:

1. **Selection**: we suggest terminating preference elicitation sooner by returning k alternatives to the group members rather than returning just one item.
2. **Approximate winners**: we suggest computing approximate winner or winners with some confidence level. This as well reduces the communication cost.
3. **Aggregation**: we suggest considering the *Least Misery* aggregation *Strategy* beyond the known Majority based strategy

We evaluated the approach on multiple datasets in different scenarios and application domains: (1) Two datasets that were collected using a group recommender system named "Lets Do It" which was built and operated in Ben-Gurion University. (2) Two real world datasets, the Netflix data (http://www.netflixprize.com) and Sushi data (Kamishima, Kazawa, & Akaho, 2005). (3) Simulated data which allow us to study the impact of the probability distribution. We show that selecting the suitable aggregation strategy and relaxing the termination condition can reduce communication up to 90%.

This paper is an extension of the authors' previous short paper (Naamani-Dery, Kalech, Rokach, & Shapira, 2014). In the previous paper, we shortly presented one preference elicitation algorithm (DIG) without approximation. In this paper we added an overview of the state of the art in the field of voting techniques (Section 2). We extended the model and definitions and added a model for approximation of the necessary winner (Section 3). We added another algorithm, ES (Section 4). This allows us to compare the algorithms' performance in different settings and show that each algorithm has an advantage in different scenarios. We have extended our evaluation to include a user-study, detailed experiments and a thorough analysis (Sections 5 and 6).

2. Related work

Group decision making consists of two phases: preference elicitation and preference aggregation. We start with describing the preference elicitation strategies considered in this paper, and move on to describe how preference elicitation.

2.1. Preference aggregation strategies

One of the contributions of this paper is to consider the Least Misery strategy, which, to our best knowledge, has not been studied in the context of *preference elicitation*. Throughout this paper we use "Majority" and "Least Misery" to refer to the Majority based strategy and the Least Misery based strategy (Masthoff, 2004).

Different studies have shown how different strategies affect group members (Masthoff, 2004; Senot, Kostadinov, Bouzid, Picault, & Aghasaryan, 2011). Masthoff studies how humans prefer to integrate personal recommendations. She concludes that users use the Majority Strategy, the Least Misery strategy and Majority without Misery strategy (Masthoff, 2004). Her findings motivate our research to focus on the Majority and the Least Misery strategies. These two strategies were also chosen by Baltrunas, Makcinskas, and Ricci (2010), in research focusing on the evaluation of the effectiveness of Group Recommender Systems obtained by aggregating user preferences.

In the *Majority Strategy* the users' ratings of the different items are aggregated and the item with the highest total value is the winner. Note that the result is similar to selecting the item with the highest average, thus this strategy is sometimes referred to as the Average Strategy or the Additive Strategy (Masthoff, 2011). The Majority strategy is used in numerous applications. For example: in the MusicFX system the square of the individual preferences are summed (McCarthy & Anagnost, 1998) and the Travel Decision Forum assists in planning a holiday (Jameson, 2004). Yet another example is of TV programs recommendation for a group (Masthoff, 2004; Yu, Zhou, Hao, & Gu, 2006), where the chosen program fits the wishes of the majority of the group. A disadvantage of this strategy is that it can be unfair towards users with the minority view. In fact, Yu et al. (2006) state that their system works well for a homogenous group but when the group is heterogeneous, dissatisfaction of the minority group occurs.

Download English Version:

https://daneshyari.com/en/article/6855649

Download Persian Version:

https://daneshyari.com/article/6855649

Daneshyari.com