

Agent-based ordinal classification for group decision making

Ons Nefla, Imène Brigui, Paolo Viappiani, Oussama Raboun

► **To cite this version:**

Ons Nefla, Imène Brigui, Paolo Viappiani, Oussama Raboun. Agent-based ordinal classification for group decision making. The 2020 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT'20), Dec 2020, Melbourne (en virtuel), Australia. hal-03037304

HAL Id: hal-03037304

<https://hal.sorbonne-universite.fr/hal-03037304>

Submitted on 3 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Agent-based ordinal classification for group decision making

Ons Nefla

Université de Paris Dauphine
PSL Research University CNRS, LAMSADE
75016 Paris, France
ons.nefla@dauphine.eu

Imène Brigui

emlyon business school
23 Avenue Guy de Collongue
Ecully, France
brigui-cthioui@em-lyon.com

Meltem Öztürk

Université de Paris Dauphine
PSL Research University CNRS, LAMSADE
75016 Paris, France
meltem.ozturk@dauphine.fr

Paolo Viappiani

LIP6, UMR7606
CNRS, Sorbonne Université
75005 Paris, France
paolo.viappiani@lip6.fr

Oussama Raboun

Keley Data
28 Rue du Dr Finlay
75015 Paris, France
oraboun@keley-data.com

Abstract—In this article we are interested in group decision aiding for an ordinal classification problem. Our approach is based on a multiagent system where each decision maker is represented by a user agent and the process is guided by a mediator agent. Each user agent has a personalized preference-based behavior defined by a utility function. The aim of the process is to converge to a group classification using a negotiation procedure. We present an experiment with real and simulated data in order to illustrate our approach and assess its performance (with respect to both user satisfaction and privacy) comparing it with two centralized methods.

Index Terms—Multiagent systems, negotiation, ordinal classification, group decision making, preferences, simulations

I. INTRODUCTION

Group decision making is a process in which multiple stakeholders, with individual preferences, act collectively to make a common decision. This decision could be a choice of an alternative among a set of possible alternatives or a classification of an object in a particular performance class. An ordinal classification problem [8], [11] consists in partitioning a set of objects into predefined ordered classes, called categories. The classification of hotels into one star, two stars, etc. is a common example. In internet applications, groups of people need to agree on a joint decision in a variety of situations, such as setting up meetings, planning vacations, watching movies together, classifying applications for a given job, to name but a few.

In this paper, we propose a negotiation-based ordinal classification process characterized by:

- ordinal agents' preferences,
- evolving preferences with respect to external factors (such as remaining time),
- the possibility of not revealing full preferences, therefore providing a certain degree of privacy.

In fact, our multiagent group decision making process is based on a multilateral negotiation to reach common ordinal classifications from individual preferences. Each user involved in the

classification process is represented by a user agent acting on behalf of her preferences, constraints and goals. We assume that all agents are cooperative and share the common goal of reaching an agreement. However, local preferences could be different or even contradictory and common classifications become hard to get to. In order to avoid such conflictual situations, a mediator agent is given the task of initiating and conducting the negotiation process. A key role of the mediator agent is to make an objective decision when the negotiation doesn't lead to an agreement.

Each user agent behaves dynamically all along the process; she makes her decisions based on an aggregation of several criteria representing her profile such as flexibility over time. The communication protocol is characterized by a some degree of privacy, as the agents decide whether to (partially) reveal their preferences.

The paper is organized as follows. In Section II we introduce our classification process; Section III is dedicated to detail agents algorithms and negotiation strategies. In Section IV we present our experimental results. We conclude with final remarks.

A. Related works

Group recommendation systems provide personalized recommendations to group of users. While most group recommender systems are based on aggregation procedures [1], [4], a number of recent works deal with negotiation techniques [5], [9], [15], [16] in order to make a single choice for the group. The main novelty of our work is therefore to tackle classification problems.

Consensus Reaching Process (CRP) [3], [10], [13] is widely used to derive collective decisions from fuzzy agents' preferences [2], [12], [14]. It is based on a multi-round process where agents are allowed to discuss and negotiate in order to adapt their preferences at each round. While most of the CRP-based methods deal with ranking problems, Jabeur and

Martel [7] propose a CRP-based ordinal classification method based on the relative importance of agents. Jabeur and Martel’s method do not consider any dynamic external factor and agents’ privacy is not respected as agents need to share their preferences in order to reach a consensus.

II. CLASSIFICATION PROCESS

The classification process is based on an interaction between a mediator agent and several user agents. The user agents share a common goal that is to incrementally construct common classifications. Since agents could have different preferences, the mediator agent insures the convergence of the process.

We now present some useful preliminary notations :

- $\mathcal{A} = \{a_1, a_2, \dots, a_p\}$: the set of user agents representing the decision makers and m represents the mediator agent.
- $\mathcal{X} = \{x_1, x_2, \dots, x_n\}$: the set of alternatives considered for classification in $\mathcal{C} = \{c_1, \dots, c_q\}$, $q \geq 2$. We assume that, $\forall k \in \{1, \dots, q-1\}$: c_{k+1} is better than c_k .
- $cl : \mathcal{A} \times \mathcal{X} \rightarrow \mathcal{C}$: agents’ classifications. $cl(a_i, x_j) = c_k$ means that for agent a_i , the preferred classification of the alternative x_j is the category c_k .
- $c_* : \mathcal{X} \rightarrow \mathcal{C}$: the common/final classification; $c_*(x_j) = c_k$ means that the object x_j is classified in the category c_k at the end of the process.
- α , with $\alpha \in]0.5, 1]$: the majority threshold enabling a coalition of agents to be decisive.
- \hat{t} : the negotiation deadline.

The classification process is an iterative negotiation initiated by the mediator agent. At each round t , the mediator agent asks all user agents about how they would classify a given alternative. Each user agent has her own preference-based behavior; one agent may just communicate her preferred category while another user agent communicates a set of “possible” categories (depending on whether she is “selfish” or flexible, and on her sensitivity to the passing of time, etc.).

Once all user agents have sent their possible categories, the mediator agent makes one of the two following decisions :

- 1) either she classifies the alternative: when a qualified majority of user agents agreed with the classification,
- 2) or she postpones the classification: when there is no sufficient majority supporting any classification.

After having proposed all the alternatives to the user agents for classification, the mediator agent repeats the whole process again for the alternatives for which the classification has been postponed; until the process deadline is reached. When repeating the procedure, the preferences of some users may change, in particular displaying more flexibility. When time is over, the mediator agent uses plurality rule to classify the remaining alternatives. Our classification process is displayed in Figure 1.

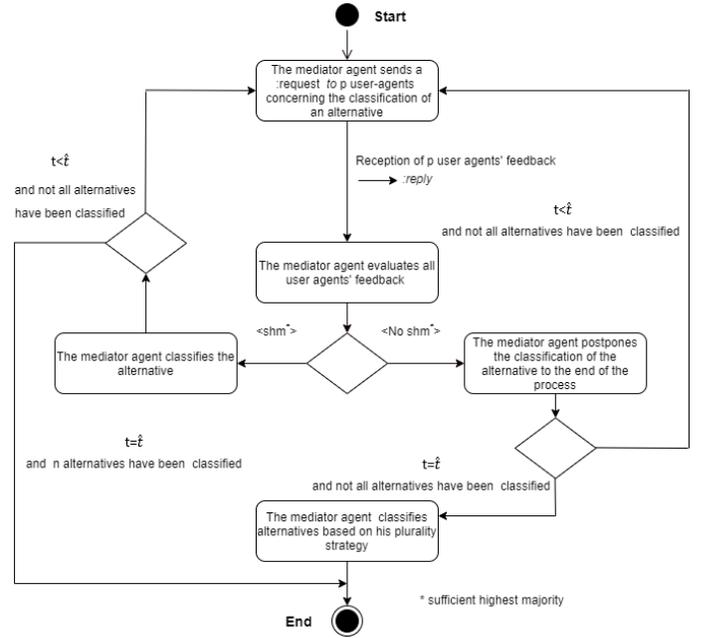


Fig. 1: The multiagent classification process.

III. AGENTS’ DESCRIPTION

Our multiagent system is formed by user agents representing human decision makers and a mediator agent responsible of the conduct of the process. The process follows a communication protocol that specifies the valid actions of each agent.

A. User agents

Each user agent represents a different user and has its own preference-based behavior. A global utility function, kept private during the whole process, represents this behavior. When the mediator agent asks about the classification of an object x , the user agent calculates for each category the utility of classifying x in it; if she judges the utility sufficient then this category is possible for her. The user agent may change her opinion about the possible classification of x during the process. The utility function is defined as $U : \mathcal{X} \times \mathcal{C} \times \mathbb{N} \rightarrow [0, 1]$; $U(x, c_k, t)$ represents the utility of classifying x in c_k when $t < \hat{t}$. The threshold λ is used in order to decide whether or not the classification is possible: $v : \mathcal{X} \times \mathcal{C} \times \mathbb{N} \rightarrow \{0, 1\}$:

$$v(x, c_k, t) = \begin{cases} 1 & \text{If } U(x, c_k, t) \geq \lambda \\ 0 & \text{Otherwise} \end{cases}$$

There are different criteria that have an impact on the utility function U , namely: the difference between the decision maker’s preferred classification and the proposition of the mediator, the number of concessions already done and the remaining time. These three criteria are evaluated with respect to three subutility functions u , fl , and f , that we now discuss:

- $u : \mathbb{N} \rightarrow [0, 1]$ is used to model the tolerance in accepting a different classification based on the deviation of c_k from the agent’s preferred classification cl . The deviation is taken as the displacement, i.e. $d(c_i, c_j) = |i - j|$. Naturally, $u(0) = 1$ and u is not increasing.

- $fl : \mathbb{N} \rightarrow [0, 1]$ is used to model the flexibility of the agent with respect to the current time (measured in number of rounds), $t \in \mathbb{N}$, where $t \leq \hat{t}$: the closer we are to the deadline, the more the agent is open to collaborate.
- $f : \mathbb{N} \rightarrow [0, 1]$ models the agent's disposition towards making a new concession, when $r \in \mathbb{N}$ concessions have been made before. A concession occurs when the agent's classification is different from the group classification.

Each user agent can aggregate these three sub-utilities in different ways, in order to compute $U(x, c_k, t)$ for a given x, c_k, t . The choice of the aggregation operator depends on the decision-maker's behavior, dictated by a *strategy*. In principle there are many different aggregation methods that could be used. In this work we consider the following strategies:

- A *flexible* strategy makes use of a compensatory operator Δ , defined as follows: $\Delta(A_1, \dots, A_n) = \frac{\sum_{i=1}^n A_i}{n}$. A flexible agent uses $\Delta(u, fl, f)$ to compute $U(x, c_k, t)$; a low performance on one criterion u, fl or f , may be compensated by high performances on other criteria.
- An *inflexible* strategy makes use of a non-compensatory operator ∇ , defined as follows: $\nabla(A_1, \dots, A_n) = \sqrt[n]{\prod_{i=1}^n A_i}$. An agent following this strategy uses $\nabla(u, fl, f)$ to compute $U(x, c_k, t)$; all the performances on criteria u, fl and f should be good in order to get a good value in output.
- *Hybrid* strategies make use of combinations of the operators Δ and ∇ . The aggregation is performed in one of the following ways:

$$- \nabla(u, \Delta(fl, f)) \rightarrow \sqrt{u \times \left(\frac{fl+f}{2}\right)}.$$

An agent following this strategy allows compensation between fl and f . The result of such compensation is aggregated with u using ∇ .

$$- \Delta(u, \nabla(fl, f)) \rightarrow \frac{u + \sqrt{fl \times f}}{2}$$

An agent following this strategy combines the fl and f in a non-compensatory way, but allows compensation between u the pair fl and f .

Algorithm 1 describes the negotiation process from the point of view of the user agent. The process starts by the request of the mediator. Once the user agent receives the `request` (line 11), she computes the U function in order to define the set of possible classification categories for x and to communicate it to the mediator. When the user agent receives the `inform` message (line 20) of the mediator agent precisising the final classification of x , she updates the number of concessions.

B. Mediator agent

The role of the mediator agent is to coordinate the classification process in order to reach common classifications; she has no knowledge about the user agents' strategies.

The mediator maintains a queue of alternatives that need to be classified, initialized as the set of alternatives. She starts by picking the alternative x at the top of the queue and sends a common request to all user agents concerning the classification of the alternative x . Then, she collects the

Algorithm 1: The UserAgent class

```

1 agent UserAgent:
2   properties
3     r ← 0          /* n. of concessions */
4     λ             /* acceptance threshold */
5   end
6   function main():
7     t ← 1
8     while true do
9       msg ← listenMessages()
10      if msg.getPerformative() == request
11        then
12          t ← t + 1
13          x ← msg.getContent()
14          C ← ∅
15          for c ∈ C do
16            if U(x, c, t) > λ then
17              C ← C ∪ c
18            end
19          end
20          reply(m, C) /* send reply */
21        else if
22          msg.getPerformative() == inform then
23            (x, c*) ← msg.getContent()
24            if c* ≠ cl(x) then
25              r ← r + 1 /* concession */
26            end
27          end
28        end
29      end
30    end

```

replies of all user agents, and verifies if there is a category supported by a qualified majority. If this is the case, then the mediator agent adopts the category as the classification of x (if more categories are acceptable to a majority, she chooses the one with the highest number) and informs the user agents. Otherwise, alternative x is added at the end of the queue of alternatives.

The process is repeated until all alternatives are classified or until we reach the end of the negotiation time \hat{t} . When time is over and some alternatives are still unclassified, the mediator agent makes a classification based on the plurality rule (an alternative is classified to the category supported by the highest number of user agents) and the process ends.

IV. EXPERIMENTAL APPROACH

In this section we present an application inspired from a real case. We consider that a group of persons (experts) have to decide in which category they classify different touristic destinations for a travel magazine. Their aim is to get a *collective* classification even if their personal preferences differ and may be contradictory. Our case study uses real data on the classifications and simulated data on the behaviors of experts.

We collect the classification of each expert to whom we assign a user agent a_i . We have 38 experts in total ($\mathcal{A} = \{a_1, \dots, a_{38}\}$), 50 destinations ($X = \{x_1, \dots, x_{50}\}$) and five categories (c_1, c_2, c_3, c_4 and c_5 representing “very bad”, “bad”, “medium”, “good”, “very good”).

We sample utility functions representing the behavior of the experts according to uniform distributions. We test our approach on three different groups, each following a different negotiation strategy:

- 1) Group 1 (the “flexible” agents): all agents are flexible and have the flexible strategy $\Delta(u, fl, f)$.
- 2) Group 2 (the “inflexible” agents) : all agents have the inflexible strategy $\nabla(u, fl, f)$.
- 3) Group 3 (“hybrid”): 10 agents with the flexible strategy, 10 agents with the non flexible strategy and 18 agents with the hybrid strategies; 9 agents have $\nabla(u, \Delta(fl, f))$ strategy while 9 agents have $\Delta(u, \nabla(fl, f))$ strategy.

Our process was designed and developed using Java Agent Development Environment (JADE)¹.

In our experiment below, we compare our multi agent approach to two simple centralized approaches (not based on negotiation) assuming that the experts give all their preferred classifications in the beginning of the process; thus with full revelation. We choose two intuitive centralized methods:

- the **median** method picks as a classification the category that is the statistical median of the distribution of expert classifications;
- the **plurality** method picks as a classification the category supported by the highest number of agents; i.e. the mode of the distribution of expert classifications.

For each group, we run 30 simulations using different utility functions of the user agents. We test our results for different values of the majority threshold α . To ensure that the different results are comparable, in all simulations the alternatives are placed in the queue in the same order. The acceptance threshold is fixed to $\lambda = 0.5$ for all the user agents and the negotiation deadline $\hat{t} = 50$.

We evaluate the performance of our process using two types of indicators. The indicators of the first type focus on the individual satisfaction of the agents, while those of the second type measure the amount of information revealed by each agent along the process. These indicators are discussed next.

A. Evaluation criteria

a) *Average Dissatisfaction (AvD)*: We compute the deviation between the final classification and the experts classifications. The results is averaged with respect to both the experts and to the travel destinations. The *AvD* from a group classification will be defined as:

$$AvD = \sum_{i=1}^p \sum_{j=1}^n \frac{|cl(i, x_j) - c_*(x_j)|}{n \times p}$$

where $c_*(x_j)$ is the final classification of destination x_j .

b) *Maximum Dissatisfaction (MaxD)* : We compute the highest value of dissatisfaction (averaged over alternatives) among the experts:

$$MaxD = \max_{i=1, \dots, p} \frac{1}{n} \sum_{j=1}^n |cl(i, x_j) - c_*(x_j)|.$$

¹<https://jade.tilab.com/>

c) *Dispersion of Dissatisfaction Index (DD)*: We measure the inequality with respect to the total dissatisfaction among experts; DD is inspired from the GINI index [6] which is used to measure inequality of income or wealth distribution of a population. In our domain, the fairest classification consists in the situation where all agents have the same value of dissatisfaction with respect to the final classifications. This index ranges between 0 in the case of perfect equality and 1 in the case of perfect inequality; it is computed comparing the cumulative sum of (sorted) individual values of dissatisfaction with the of an ideal case.

Assume, for instance, that we have three users where the preferred category of the first is, on average, 3 slots away from final classification, while the other two have their preferred category coinciding with the final one; the cumulative sums of dissatisfaction is $(0, 0, 3)$. In an perfectly fair situation the total dissatisfaction (3 slots) would be split equally between the users (1 each), with cumulative sum $(1, 2, 3)$. In this example the *DD* will be 1 as the 1-norm of the difference between cumulative sums is: $\|(0, 0, 3) - (1, 2, 3)\|_1 = 1 + 2 + 0 = 3$, we divide by the number of users (3), and obtain $DD = 3/3 = 1$.

In more details, *DD* corresponds to the difference between two Lorenz curves:

- the first one is represented by the identity function, displaying the theoretical fair classification.
- the second is represented by a normalized cumulative function displaying the ratio between the deviations² from a group classification, and the global deviation³.

Thus, DD is defined using the following steps:

- 1) We define the agent’s deviation *D* as a mapping assessing the average deviation for each agent. *D* is defined as:

$$D : \{1, \dots, p\} \longrightarrow [0, q-1] \\ i \longmapsto \sum_{j=1}^n \frac{|cl(i, x_j) - c_*(x_j)|}{n}$$

- 2) We define the global deviation *GD* as:

$$GD = \sum_{i=1}^p D(i).$$

- 3) To sort the deviation values, we use a permutation function σ , defined as on $[1, p]$, such that: $\forall i, j \in [1, p]$

$$D(i) < D(j) \Rightarrow \sigma(i) < \sigma(j)$$

Finally, the function *s* is a mapping defined as follows:

$$s : \{1, \dots, p\} \longrightarrow [0, 1] \\ i \longmapsto \sum_{j=1}^i \frac{D(\sigma^{-1}(j))}{GD}$$

where σ^{-1} is the inverse of σ .

Hence, the *DD* is defined as follows:

$$DD = 2 * \sum_{i=1}^p \frac{\frac{i}{p} - s(i)}{p}.$$

The term 2 is used in order to normalize the index between zero and one.

²Sorted from the less deviating to the most deviating user agent.

³The sum over all user agents’ deviations.

d) *Information Privacy Index (IPI)* : It is related to the amount of information shared between user agents and the mediator agent. When we use centralized methods (median and plurality), experts divulge all their classifications. When we use our approach, the preference based behavior of experts (utility functions, aggregations, thresholds) are only known by their user agent. Moreover, a user agent does not communicate about preferred classifications of her expert. She only communicates about “possible” classifications (calculated using the preferred classification of the expert and her utility function, etc.). For instance, it may happen that the user agent says that all the categories are possible, which means that the mediator did not obtain any clear information about the preferred classification of the expert.

Thus, the basic idea of our indicator is the following: the more “possible categories” are communicated to the mediator, the less information is shared. The index $IPI(i, x_j)$ measures the information shared with the mediator (if the user agent reports a single category, this has to be the preferred category, that is revealed to the mediator). We define $IPI(i, x_j)$ as:

$$IPI = \frac{1}{(q-1)} \left(\sum_{C \in \mathcal{C}} v_i(x_j, C) - 1 \right).$$

B. Experimental results

Our approach is based on a majoritarian rule (the default value of majority threshold α is 0.5). However one can use greater values of α in order to be more exigent about final classifications (for instance $\alpha = 0.7$ means that the category c_k can be chosen by the mediator if and only if 70% of user agents define it as possible). Naturally when α gets greater, a common classification becomes difficult to be reached and there are more postponed classifications (that will be decided using plurality). Our simulations show us that for $\alpha = 0.5$, in Group 1 (resp. Group 2 and Group 3) 100% (resp. 98% and 45%) of alternatives are classified before the deadline, as the agents reach an agreement (hence without the use of plurality). This proportions get smaller when we increase α ; for $\alpha = 0.7$, in Group 1 (resp. Group 2 and Group 3) 87% (resp. 19% and 8%) of alternatives are classified before the deadline.

a) *Average Dissatisfaction (AvD)*: Figure 2 shows average dissatisfaction values depending on the value of majority threshold α . Five methods are tested: our approach (on Group1, Group2 and Group3), the plurality, and the median methods. As expected, when we use our negotiation-based approach the average dissatisfaction of Group1 is greater than the one of Group2, which is greater than the one of Group3.

In general we remark that the results are convincing for all the methods, considering 5 categories, the classification of a travel destination given by an expert has an average deviation around 0.7-0.8 which is less than one category (note that, as there are 5 categories, the theoretical maximum is 4).

We remark also that the plurality and median methods have very similar values of AvD and our approach gets closer to them when experts become less flexible (Group 3).

Not surprisingly, the bigger the value of α , the closer our approach to the plurality method. This is due to the increase of postponed classifications which may not be treated in time \hat{t} . Thus, the mediator agent takes an objective decision based on plurality over the user agents’ feedback.

b) *Maximum Dissatisfaction Index (MaxD)* : Figure 3 presents the maximum dissatisfaction values with respect to different values of the majority threshold α . *MaxD* values of our simulations are between 2.5 and 2.9 (by definition the range of the index is obviously between 0 and 4). It means that in our group of experts there exists at least one expert having contradictory preferences with the majority of experts and his average deviation from the final classifications is between 2 and 3 categories. A quick look to our data shows that there are two experts who were very strict and classified a big number of destinations in the *very bad* category. These two individuals, whose preferences set them apart from the rest of the group, may have more influence towards the final result when the other experts are flexible or the median method is used. This observation is coherent with the fact that the plurality approach “protects” less the minorities.

Unlike the average dissatisfaction, the maximum dissatisfaction of Group1 is smaller than the one of Group2, which is smaller than the one of Group3. It shows that when the experts have flexible preferences, even if the average dissatisfaction of the group is higher than the one of a inflexible group, the unhappiest expert of a flexible group is more satisfied than the one of a mixed or inflexible group.

The values of *MaxD* are bounded by the values obtained with the plurality (worst case) and the median (best case). Like with the *AvD*, not surprisingly, the bigger the value of α , the closer our approach to the plurality method.

c) *Dispersion of Dissatisfaction Index (DD)*: Figure 4 shows the dispersion of dissatisfaction values with respect to different values of α . While the theoretical range of the index is between 0 and 1, in our simulations, we observe DD values between 0.16 and 0.25. This shows that all the methods perform quite fairly (as they distribute the global dissatisfaction in a somewhat “egalitarian” way). One explanation of this may be related to the nature of our case study. Even though our experts have different preferences over the destinations, our classification problem is not that controversial.⁴

Notice that for α between 0.5 and 0.8, our approach provides more fair results than centralized methods and as in the case of *MaxD* and , the greater α , the closer we are to the plurality. With respect to DD, our method performs better when agents are of Group1 than with Group2, and even worse with agents of Group3.

d) *Information Privacy Index (IPI)* : Figure 5 illustrates information privacy index with respect to different values of α . As explained in the previous section, IPI is zero for centralized methods since the preferred classifications of all experts are known by the one who controls the process. Using

⁴As an example of a domain where the classification is probably more controversial, consider the case of classifying nuclear energies as either *acceptable*, *neutral* or *not acceptable*.

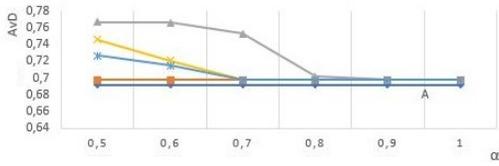


Fig. 2: AVD

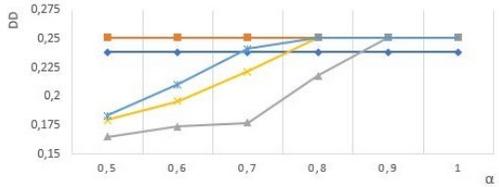


Fig. 4: DD

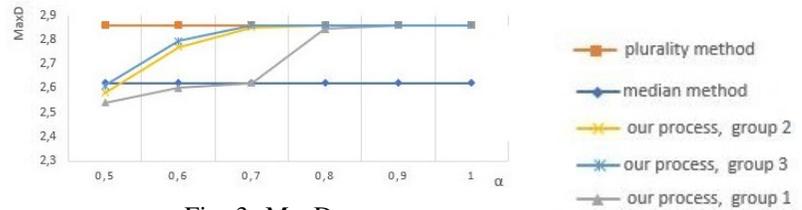


Fig. 3: MaxD

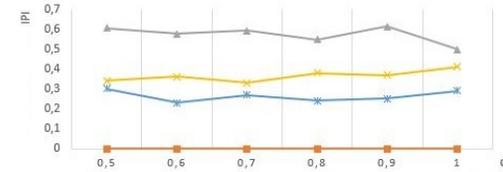


Fig. 5: IPI

our negotiation-based approach we observe values between 0.23 and 0.612 (the theoretical maximum is 1).

Not surprisingly, the more flexible are experts' preferences, the less information is shared (Group1 is "better" than Group2, which is better than Group3). Contrarily to the previous indicators, the difference between the values of our three groups are quite different. Remark that when our experts are all flexible, only less than 30% of classification information is shared with the mediator.

Unlike the other indices, the value of IPI is not dependent on α . According to IPI, our approach performs significantly better than centralized methods.

V. CONCLUSION

We have proposed a multiagent process for deriving a common ordinal classification for a group of agents from their individual classifications. In our approach each user is represented by an agent whose actions are based on her utility function. Our experiments showed that our process, when compared to centralized approaches, is fairer (the worst-off agent is more satisfied) and better protects user privacy.

In future works we plan to test our approach in larger datasets (increasing the number of alternatives and decision makers) and with more sophisticated decision strategies (for instance, including the possibility for agents to vetoing a proposed classification). We would also like to consider incomplete or imprecise classification data, such as interval of categories, that can facilitate the task of expressing preference information. Moreover, we could relax the assumption that alternatives have the same importance. Similarly, we could model the different degree of expertise in classifying the alternatives; in particular an expert may have a high confidence in her preferred classification for an alternative, while she may have low confidence in other alternatives.

REFERENCES

- [1] Linas Baltrunas, Tadas Makcinskas, and Francesco Ricci. Group recommendations with rank aggregation and collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 119–126. ACM, 2010.
- [2] Nicola Capuano, Francisco Chiclana, Hamido Fujita, Enrique Herrera-Viedma, and Vincenzo Loia. Fuzzy group decision making with incomplete information guided by social influence. *IEEE Transactions on Fuzzy Systems*, 26(3):1704–1718, 2017.
- [3] Yucheng Dong, Quanbo Zha, Hengjie Zhang, Gang Kou, Hamido Fujita, Francisco Chiclana, and Enrique Herrera-Viedma. Consensus reaching in social network group decision making: Re-search paradigms and challenges. *Knowledge-Based Systems*, 162:3–13, 2018.
- [4] Inma Garcia, Sergio Pajares, Laura Sebastia, and Eva Onaindia. Preference elicitation techniques for group recommender systems. *Information Sciences*, 189:155–175, 2012.
- [5] Inma Garcia and Laura Sebastia. A negotiation framework for heterogeneous group recommendation. *Expert Systems with Applications*, 41(4):1245–1261, 2014.
- [6] Corrado Giugni. *Variabilità e mutabilità: Contributo allo studio delle distribuzioni e delle relazioni statistiche*. Cuppini, 1912.
- [7] Khaled Jabeur and Jean-Marc Martel. An ordinal sorting method for group decision-making. *European Journal of Operational Research*, 180(3):1272–1289, 2007.
- [8] Vincent Mousseau, Roman Slowinski, and Piotr Zielniewicz. A user-oriented implementation of the ELECTRE-TRI method integrating preference elicitation support. *Computers & OR*, 27(7-8):757–777, 2000.
- [9] Rachel Perez Palha. Negotiation throughout flexible and interactive tradeoffs applied to construction procurement. *Automation in Construction*, 99:39–51, 2019.
- [10] Ignacio J Pérez, Francisco Javier Cabrerizo, Sergio Alonso, YC Dong, Francisco Chiclana, and Enrique Herrera-Viedma. On dynamic consensus processes in group decision making problems. *Information Sciences*, 459:20–35, 2018.
- [11] Bernard Roy. The outranking approach and the foundations of electre methods. In *Readings in multiple criteria decision aid*, pages 155–183. Springer, 1990.
- [12] Raquel Ureña, Francisco Javier Cabrerizo, Juan Antonio Morente-Molinera, and Enrique Herrera-Viedma. Gdm-r: A new framework in r to support fuzzy group decision making processes. *Information Sciences*, 357:161–181, 2016.
- [13] Raquel Ureña, Francisco Chiclana, Guy Melancon, and Enrique Herrera-Viedma. A social network based approach for consensus achievement in multi-person decision making. *Information Fusion*, 47:72–87, 2019.
- [14] Raquel Ureña, Francisco Chiclana, Juan Antonio Morente-Molinera, and Enrique Herrera-Viedma. Managing incomplete preference relations in decision making: a re-view and future trends. *Information Sciences*, 302:14–32, 2015.
- [15] Christian Villavicencio, Silvia Schiaffino, J Andres Diaz-Pace, and Ariel Monteserin. Group recommender systems: A multi-agent solution. *Knowledge-Based Systems*, 164:436–458, 2019.
- [16] Christian Paulo Villavicencio, Silvia Noemi Schiaffino, Jorge Andres Diaz Pace, and Ariel José Monteserin. A group recommendation system for movies based on MAS. 2016.