

Title	<i>Bayesian Preference Elicitation for Group Decisions with the Plackett Luce Model</i>
Mission	<p>Group recommender systems [2,4,5] are automated tools that help groups of people to make a joint decision. A key task in any recommender system is the elicitation [1] of the users' preferences. In an interactive elicitation approach [1,5,6] questions (as "Do you prefer a to b?") are posed to the users in order to incrementally improve the knowledge about the users' preferences until a decision can be made with a certain degree of certainty.</p> <p>In this project we assume a Bayesian approach where a distribution over possible preferences is maintained; answers to questions are used to update this distribution over the preferences of the users. The uncertainty over the agents' preferences can be modelled using probabilistic ranking models [7], in particular we will focus on the Plackett-Luce model in this work.</p> <p>The preference elicitation is incremental: at each step, the system needs to decide the user to whom ask the question and which question to ask. The next question is chosen by picking the most informative question according to some measure of information value (generalizing the Bayesian elicitation framework for a single decision maker [6]). The procedure continues until the expected loss is below a given threshold.</p> <p>The candidate will develop a scalable Bayesian framework for preference elicitation for group decisions, building on recent works from the supervisors [3,6]. The main tasks will be:</p> <ul style="list-style-type: none">• Implementation of algorithms for learning the prior distributions from data• Design of strategies for choosing the next question to ask based on the Plackett Luce model and theoretical analysis (i.e. approximation guarantees)• Design and implementation of Bayesian inference• Adaptation of the framework to different preference aggregation methods [4]• Experimental evaluation to compare the effectiveness of the elicitation strategies in simulations using real datasets considering the tradeoff between cognitive cost (i.e. number of questions) and computational cost; comparison with previous work [3,5] <p><i>References</i></p> <p>[1] Dariusz Braziunas, Craig Boutilier. <i>Elicitation of Factored Utilities</i>. <i>AI Magazine</i> 29(4): 79-92 (2008)</p> <p>[2] Dara, S., Chowdary, C.R. & Kumar, C.. <i>A survey on group recommender systems</i>. <i>Journal of Intelligent Information Systems</i>, 2019, p 1-25.</p> <p>[3] Maeva Caillat, Nicolas Darcel, Cristina Manfredotti, Paolo Viappiani. <i>Bayesian Vote Elicitation for Group Recommendations</i>. <i>From Multiple Criteria Decision Aid to Preference Learning (DA2PL 2020)</i>,</p> <p>[4] Judith Masthoff. <i>Group Recommender Systems: Aggregation, Satisfaction and Group Attributes</i>. In <i>Recommender Systems Handbook</i>, pp 743-776 (2015)</p> <p>[5] Naamani-Dery, L., Golan, I., Kalech, M. et al. <i>Preference Elicitation for Group Decisions Using the Borda Voting Rule</i>. <i>Group Decision and Negotiation</i> (2015) 24: 1015.</p> <p>[6] Paolo Viappiani and Craig Boutilier. <i>On the Equivalence of Optimal Recommendation Sets and Myopically Optimal Query Sets</i>. <i>Artificial Intelligence Journal</i> (2020).</p> <p>[7] Lirong Xia. <i>Learning and Decision-Making from Rank Data</i>. <i>Synthesis Lectures on Artificial Intelligence and Machine Learning</i>, Morgan & Claypool Publishers 2019</p>
Prerequisites	Be enrolled in Master2 in computer science, applied mathematics or in a top engineering school. Some background in probability / statistics is helpful.
Place	UMR MIA-AgroParisTech. 16, rue Claude Bernard Paris 5th arrondissement.
Duration	5 à 6 mois de stage à partir de février / mars 2021, temps plein.
Remuneration	Gratification according to the INRA scale in force (approx. 570 Euros / month in 2019)
Contact	Cristina Manfredotti, AgroParisTech, cristina.manfredotti@agroparistech.fr , Paolo Viappiani, LIP6, Sorbonne Université, paolo.viappiani@lip6.fr