TOWARDS A MORE INTUITIVE CONTENT RECOMMENDER:
Utilizing Pairwise Comparison Strategies for Preference Ranking

Amer K. Mriziq
COMPSCI 299: Designing Algorithmic Media
December 2023
UC Berkeley

Preference inference is crucial to any system looking to distribute content to those who need it most. Most content recommenders today optimize for engagement (Thorburn, 2023) and infer revealed preferences rather than elicit the explicit preference from users. This paper explores pairwise comparison as a means of preference elicitation, and the underlying methods to generating a complete preference ranking or utility function. Recommenders based on explicit signals from pairwise comparison serve as an alternative to surveilling user behavior as a proxy for preference, which presents its own challenges (Milli et al. 2021).

To understand pairwise comparison as a means of producing a distribution of preferences, we must consider two design questions: what are the main methods for selecting pairs to ask the decision-maker, and how are the results of such voting methods processed to infer preferences or produce a utility function? Pairwise comparison, when used for preference elicitation, presents certain challenges. One significant issue is the cognitive burden it places on individuals, stemming from the numerous interactions required to establish a complete and comprehensive preference ranking. It must also contend with the challenges posed by the decision maker's potential indifference or uncertainty regarding two choices. The answers to these two research questions look to explore research focusing on addressing these issues by means of algorithmic pair selection and underlying probabilistic models to infer hidden preferences.

Strategies for Selecting Pairs for Comparison

The literature on pair selection emphasizes minimizing interactions to reduce the cognitive load on the decision maker. Ciomek et al., in their 2016 paper, introduced various heuristics to improve scoring while reducing the number of pairs required. These include Pair-wise Winning Indices (PWI) for assessing the likelihood of preferring one alternative over another, Rank Acceptability Indices (RAI) to estimate the probability of an alternative reaching a specific rank, along with methods for measuring uncertainty, evaluating information gain, and adjusting search depth. This framework aims to focus on the most critical characteristics of an item before suggesting it for comparison to the user/decision-maker. Notably, various approaches like Elo, probabilistic, and machine-learning can apply these heuristics to suggest optimal pairs. In 2017, Branke et al. gave a name to this system approach by terming the one-step look-ahead technique as a method to reduce interactions. This heuristic involves evaluating the potential information gain from each pair-wise comparison question before posing it to the decision maker (DM). It considers all possible preference structures aligning with the DM's previous answers to select the most informative question, thereby minimizing the number of questions needed to discern the DM's preferences.

The literature also explores different approaches to the “selection rule” for determining comparison pairs. Qian et al., in 2015, investigated additional techniques within the one-step look-ahead method, including a binary search strategy for selecting comparisons. Pair selection involves not just choosing items for comparison but also determining post-decision steps. For example, if a decision-maker prefers item A over B in a set containing A, B, C, and D, the next step must be decided: whether A continues for comparison with C or is replaced by another pair like C and D. This replacement policy is pivotal, with three options: retaining the chosen item, replacing it, or replacing both items. Particularly in preference ranking, the "winner remains" strategy, where the preferred item stays for further comparison, is often the most intuitive for decision makers. In 2019, Balog et al. at Google explored a novel pair-wise tag-interaction approach. Users compared and ranked content tags, and the resulting preferences were used to recommend content based on tag rankings. This method models semantically rich tag-based preferences, enabling transparent item recommendations and aligning with the transparency pillar of AI governance.

In the context of a framework that employs pairwise comparison for eliciting preferences, adaptive pair selection and preference ranking emerge as two distinct, but interconnected outcomes derived from a unified ranking model. This model functions by generating rankings for an extensive collection of items at each stage of decision-making. It utilizes a global rank to inform a one-step lookahead strategy. From this global ranking, the top 25 items are specifically chosen for recommendations or match-making purposes. The one-step lookahead strategy effectively leverages the global ranking as a guiding principle. It adaptively selects pairs of items for the decision-maker to evaluate, with the goal of optimizing the decision-making process. This optimization is achieved by focusing on explicit decisions that are based on the anticipated comparative strengths of the two items in question. This approach uses the global rankings to predict which pairs of items would yield the most informative and relevant comparisons for the decision-maker. This leads us to the next question; how are the results of such voting methods processed to infer preferences or produce a utility function?

The Bradley-Terry Model for Pairwise Preference Ranking

Probabilistic models, like the popular Bradley-Terry (BT) model, posits a framework wherein the likelihood of one item being preferred over another is quantified, thus enabling a hierarchical scoring of the items based on preference data. (Vojnovic et al. 2023[23]). In its basic form, the BT model is for binary comparisons (A is preferred over B or vice versa). Extensions of the model can handle situations where ties are possible, or where choices among more than two items are made simultaneously (multinomial choice models). The model's effectiveness depends on the quantity and quality of pairwise comparison data. Sparse or biased data can significantly affect the model's performance. The BT model assumes that the probability of one item being preferred over another can be expressed solely in terms of their relative strengths or scores. This assumption implies a certain simplicity and transitivity in preferences. In the BT model, each item in a fixed set receives a score based on these comparisons. The probability of item A "beating" item B is expressed as the ratio of A's score (S_A) to the sum of scores of A and B ($S_A + S_B$):

$$P\ (A \mid B)\ =\ \frac{SA}{(SA\ +\ SB)}$$

The BT model assumes that the probability of one item being preferred over another can be expressed solely in terms of their relative strengths or scores; S values are always positive. The estimation of strength scores is dependent on whether you take a Bayesian approach or frequentist approach to parameter estimation. Each item is given a prior distribution (often with a 0 mean and standard deviation of 1). In the Bayesian framework, these priors are updated with observed data (the decisions or comparisons made) to form posterior distributions (Vitelli et al. 2018). Simply put, the model uses observed pairwise comparison data to update the parameters (scores). Bayesian methods are particularly useful for preference inference, and in involve initializing a prior distribution for each item and incorporating an update rule. Through iterative processes (like optimization algorithms in frequentist methods like MLE, Min-Max, or MCMC in Bayesian methods), these parameters converge to values that best reflect the observed preferences. The convergence is towards the scores that make the model's predicted probabilities of preferences most consistent with the observed data. The model essentially estimates the utility (or strength) of each item based on preferences. Bayesian methods are particularly useful in this context as they provide a distribution of possible utility values, reflecting the uncertainty inherent in the estimation process.

Approaches to parameter estimation, a key hyperparameter in probabilistic models like Bradley-Terry and Thurstone-Mosteller, aids in creating utility functions or preference rankings (Handley et al. 2001). These methods, especially Bayesian parameter estimation, consider prior distributions to rank items, even those not directly compared. It allows for indirect comparisons and can incorporate prior knowledge or assumptions. Bayesian approach is the more prevalent form of sampling since it allows for scoring items not compared and ultimately eases the decision maker's cognitive load by minimizing necessary comparisons for a comprehensive preference ranking.

In 2021, D. I. Mattos and É. M. S. Ramos open-sourced the BPCS (Bayesian Paired Comparison in Stan) package, using the optimization algorithm No-U-Turn (NUTS) Hamiltonian Monte Carlo sampler, a Bayesian method. This approach is coded in the Stan language and offers several advantages over the Gibbs sampler. Regarding the parameter estimation methods for generating priors, the BPCS package utilizes normal distributions centered around 0, with a variance of 3.0 for the priors. This approach allows for the modeling of probabilities in the context of the Bayesian Bradley-Terry model. In the paper WeBuildAI: Participatory Framework for Algorithmic Governance, the optimization algorithm employed is focused on learning linear utilities for random utility models. The final model used in this study is the TM utility model with linear mode utility. This approach was chosen due to its simplicity and effective performance, and it is noted for being straightforward to summarize and explain, as the utilities are linear with respect to features. For estimating the parameters of the TM (Thurstone-Mosteller) model, the approach involves assuming that each participant's mode utility for every potential allocation is a linear function of the feature vector corresponding to that allocation. This mode utility is essentially a weighted linear combination of the features.

There is still the challenge of incorporating items that lack direct comparisons. I have yet to find papers discussing implicit signals as additional features informing the priori distribution sampling within the Bradley-Terry model. These possible scoring functions can evaluate items that haven't been directly chosen by assigning scores based on the observed relationships between items that have been explicitly selected. This process involves preprocessing data by mapping relational weights between the items that were shown to the decision-maker, thereby allowing for a contextual assessment of the unseen items. Feinstein 2003 shows us how decision times can reveal relative preferences to alternative choices in pairwise comparison settings. Over time, this proposed model would optimize for soliciting explicit votes for items that have only been assessed through implicit signals, with the goal of establishing a more robust strength distribution, as mentioned by Ciomek (2016) with regards to a continuously adjusting search depth.

Conclusion

Pairwise comparison for preference elicitation in the context of content recommenders comes with a developing field of literature but lacks a unified framework for applying these systems together. Recent research indicates that Pairwise Comparison for Preference Elicitation (PCPE) can be effectively utilized as a participatory framework in AI and algorithmic governance. The key appeal of this approach lies in its human-in-the-loop design, which prioritizes participation. This participatory element makes PCPE-based recommenders a compelling alternative for developers. Instead of solely relying on tracking user behavior to discern revealed preferences and associated welfare outcomes, this method offers a more interactive and human-centric approach. The placement of such a tool within the user interface is another factor; for example, pairwise comparison during the onboarding phase of a software's user experience may be more useful for engagement purposes rather than its utilization in another corner of the user flow. My exploration leads me now to better understand these interconnected topics by implementing models like the Bradley-Terry using Bayesian optimization algorithms and including additional features (implicit signals) for parameter estimation. This system would score all items, informing a dynamic pair selection while producing an overall preference ranking. Such a participatory framework ensures that algorithms are not merely tools for engagement maximization but instruments that embody collective preferences, ethical considerations, and social values (M.K. Lee et al. 2019 [30]). This shift from a solely profit-driven model to a more democratized and transparent approach in algorithm design is crucial for mitigating the adverse impacts of current recommender systems. By embedding stakeholder participation at the core of algorithmic decision-making, we can create digital ecosystems that are not only more equitable and accountable but also more aligned with the diverse and evolving needs of society.

REFERENCES

[1]
A. Borodinov, A. Agafonov, and V. Myasnikov, “A Method of Preference and Utility Elicitation By Pairwise Comparisons and its Application to Intelligent Transportation Recommendation Systems,” *2020 10th International Conference on Information Science and Technology (ICIST)*, pp. 77–85, Sep. 2020, doi: 10.1109/ICIST49303.2020.9201952.

[2]
K. G. Jamieson and R. Nowak, “Active Ranking using Pairwise Comparisons,” presented at the Neural Information Processing Systems, Sep. 2011. Accessed: Dec. 09, 2023. [Online]. Available: https://www.semanticscholar.org/paper/Active-Ranking-using-Pairwise-Comparisons-Jamieson-Nowak/b5dfea70185c58f4e63f99f9ecc9924a201ed7a4

[3]
C. P. Whaley, “Collecting paired-comparison data with a sorting algorithm,” *Behavior Research Methods & Instrumentation*, vol. 11, no. 2, pp. 147–150, Mar. 1979, doi: 10.3758/BF03205636.

[4]
J. C. Handley, “Comparative Analysis of Bradley-Terry and Thurstone-Mosteller Paired Comparison Models for Image Quality Assessment,” 2001.

[5]
J. Branke, S. Corrente, S. Greco, and W. Gutjahr, “Efficient pairwise preference elicitation allowing for indifference,” *Computers & Operations Research*, vol. 88, pp. 175–186, Dec. 2017, doi: 10.1016/j.cor.2017.06.020.

[7]
Y. Hairu, J. Zhengyi, and Z. Haiyan, “Exploring a Method for Ranking Objects Based on Pairwise Comparison,” *2011 Fourth International Joint Conference on Computational Sciences and Optimization*, pp. 382–386, Apr. 2011, doi: 10.1109/CSO.2011.132.

[8]
A. Beutel *et al.*, “Fairness in Recommendation Ranking through Pairwise Comparisons.” arXiv, Mar. 02, 2019. Accessed: Dec. 08, 2023. [Online]. Available: http://arxiv.org/abs/1903.00780

[9]
K. Ciomek, M. Kadziński, and T. Tervonen, “Heuristics for prioritizing pair-wise elicitation questions with additive multi-attribute value models,” *Omega*, vol. 71, pp. 27–45, Sep. 2017, doi: 10.1016/j.omega.2016.08.012.

[10]
K. Ciomek, M. Kadziński, and T. Tervonen, “Heuristics for selecting pair-wise elicitation questions in multiple criteria choice problems,” *European Journal of Operational Research*, vol. 262, no. 2, pp. 693–707, Oct. 2017, doi: 10.1016/j.ejor.2017.04.021.

[11]
L. Qian, J. Gao, and H. V. Jagadish, “Learning user preferences by adaptive pairwise comparison,” *Proc. VLDB Endow.*, vol. 8, no. 11, pp. 1322–1333, Jul. 2015, doi: 10.14778/2809974.2809992.

[12]
Z. Hoffman, “Method of pair-selection in content addressable systems,” *Nuclear Instruments and Methods*, vol. 51, no. 1, pp. 141–152, May 1967, doi: 10.1016/0029-554X(67)90375-8.

[13]
Z. Li, Q. Xu, Y. Jiang, K. Ma, X. Cao, and Q. Huang, “Neural Collaborative Preference Learning With Pairwise Comparisons,” *IEEE Trans. Multimedia*, vol. 23, pp. 1977–1989, 2021, doi: 10.1109/TMM.2020.3006373.

[14]
H. Narasimhan, A. Cotter, M. Gupta, and S. Wang, "Pairwise Fairness for Ranking and Regression." arXiv, Jan. 07, 2020. Accessed: Dec. 06, 2023. [Online]. Available: http://arxiv.org/abs/1906.05330

[15]
N. M. Tran, "Pairwise ranking: Choice of method can produce arbitrarily different rank order," *Linear Algebra and its Applications*, vol. 438, no. 3, pp. 1012–1024, Feb. 2013, doi: 10.1016/j.laa.2012.08.028.

[16]
"[PDF] Pairwise ranking: choice of method can produce arbitrarily different rank order | Semantic Scholar", doi: 10.1016/J.LAA.2012.08.028.

[17]
N. Dandekar, "Pointwise vs. Pairwise vs. Listwise Learning to Rank," Medium. Accessed: Dec. 07, 2023. [Online]. Available: https://medium.com/@nikhilbd/pointwise-vs-pairwise-vs-listwise-learning-to-rank-80a8fe8fadfd

[19]
J. Fürnkranz and E. Hüllermeier, "Preference Learning and Ranking by Pairwise Comparison," in *Preference Learning*, J. Fürnkranz and E. Hüllermeier, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 65–82. doi: 10.1007/978-3-642-14125-6_4.

[21]
V. Vitelli, Ø. Sørensen, M. Crispino, A. Frigessi, and E. Arjas, "Probabilistic Preference Learning with the Mallows Rank Model".

[22]
J. Burrell and M. Fourcade, "The Society of Algorithms," *Annu. Rev. Sociol.*, vol. 47, no. 1, pp. 213–237, Jul. 2021, doi: 10.1146/annurev-soc-090820-020800

[23]
L. Qian, J. Gao, and H. V. Jagadish, "Learning user preferences by adaptive pairwise comparison," *Proc. VLDB Endow.*, vol. 8, no. 11, pp. 1322–1333, Jul. 2015, doi: 10.14778/2809974.2809992.

[24]
"Accelerated MM Algorithms for Ranking Scores Inference from Comparison Data," ar5iv. Accessed: Dec. 07, 2023. [Online]. Available: https://ar5iv.labs.arxiv.org/html/1901.00150

[25]
K. Balog, F. Radlinski, and S. Arakelyan, "Transparent, Scrutable and Explainable User Models for Personalized Recommendation," in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, Paris France: ACM, Jul. 2019, pp. 265–274. doi: 10.1145/3331184.3331211.

[26]
D. I. Mattos and É. M. S. Ramos, "Bayesian Paired-Comparison with the bpcs Package." arXiv, Sep. 20, 2021. [Online]. Available: http://arxiv.org/abs/2101.11227

[27]
L. Thorburn, "How Platform Recommenders Work," Understanding Recommenders. Accessed: Dec. 09, 2023. [Online]. Available: https://medium.com/understanding-recommenders/how-platform-recommenders-work-15e260d9a15a

[28]
S. Milli, L. Belli, and M. Hardt, "From Optimizing Engagement to Measuring Value," in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Mar. 2021, pp. 714–722. doi: 10.1145/3442188.3445933.

[29]

J. Feinstein, “A Web-Centric Preference Acquisition and Decision Support System Employing Decision Times to express Relative Preferences,” Oct. 2003. Accessed: Dec. 30, 2023. [Online]. Available: https://www.semanticscholar.org/paper/A-Web-Centric-Preference-Acquisition-and-Decision-Feinstein/93a6b5aefdced070d8e9d2327b35c257743bbc35

[30]
M. K. Lee *et al.*, “WeBuildAI: Participatory Framework for Algorithmic Governance,” *Proc. ACM Hum.-Comput. Interact.*, vol. 3, no. CSCW, p. 181:1-181:35, Nov. 2019, doi: 10.1145/3359283.