# Fair Sequential Group Recommendations in SQUIRREL Movies

Eva Chrysostomaki[1], Maria Stratigi[2], Vasilis Efthymiou[3], Kostas Stefanidis[2] and Dimitris Plexousakis[3]

[1]*University of Crete, Greece*
[2]*Tampere University, Finland*
[3]*Foundation for Research and Technology - Hellas, Greece*

**Abstract**
Group recommendations are becoming prominent, since there are several applications in which users form groups for activities (e.g., tourism, restaurants, movies). SQUIRREL is a framework for sequential group recommendations, i.e., it may be used by a group in several rounds, providing a different recommendation in each round. The goal is that, after some rounds, all group members will be fairly satisfied with the recommendations, even if, individually, some users may not be very satisfied in every round. In this paper, we present SQUIRREL Movies, an implementation of SQUIRREL in the movies domain.

## 1. Introduction

Group recommendations have become popular, mainly due to the effortless forming of groups in social media [1]. *Sequential group recommendations*, in turn, provide suggestions to groups for a series of recommendation rounds. A fair sequential group recommender should offer good results for the current round and also ensure the group members' satisfaction over a sequence of recommendations, i.e., the recommender should observe that: (*i*) the recommended items should be as relevant as possible to the group members at each round, and (*ii*) no group member is dissatisfied after a sequence of rounds. Standard group recommendation approaches fall short of achieving both sequential group recommendation objectives [2, 3]. Compared to those standard approaches, specifically designed sequential group recommendation approaches, which consider the previous recommendation rounds, show a significantly better performance [3].

The SQUIRREL framework [4] offers sequential group recommendations based on Reinforcement Learning (RL); the sequential nature of RL directly reflects the sequential nature of recommendations. In most group recommendation systems, the goal is to maximize the group members' satisfaction with the proposed data items, measured by how relevant are the items in the group recommendation list for each group member. In SQUIRREL, two alternative *reward* functions are employed for RL: (*i*) a simple aggregation of the individual satisfaction scores of each group member for all rounds so far, and (*ii*) a more sophisticated one considering a combination between users satisfaction and *users disagreement*, i.e., the difference between the most and the least satisfied group member.

In addition to rewards, SQUIRREL also consists of states and actions. A *state* describes the status of the group in the current round, e.g., how satisfied each group member is, and based on that, SQUIRREL selects an appropriate *action*, i.e., which group recommendation method to apply for the next round. Then, a recommendation list is returned to the group, a reward is calculated, and the group state is updated to reflect the latest change. Overall, SQUIRREL identifies the best possible action to apply, with respect to a selected goal. Any rank aggregation-based group recommendation method can be incorporated into SQUIRREL as a new action. To define the various components of the model, it exploits the notions of user satisfaction and user disagreement, and can combine different techniques to counter-balance the drawbacks inherently present in other recommendation solutions.

SQUIRREL is generalizable to any domain that considers sequential group recommendations, e.g., restaurants, tourism. However, SQUIRREL is not an end-to-end system; it runs once, for the last round only, with no user interaction. In this paper, we present SQUIRREL Movies, an implementation of SQUIRREL, applied in the movies domain. SQUIRREL Movies offers a graphical UI that allows users to interact with the system and receive recommendations for the current round, while providing movie information from external sources.

## 2. System Overview

In this section, we provide an overview of SQUIRREL [4], as well as its adaptation to the movies domain.

*Joint Workshops at 49th International Conference on Very Large Data Bases (VLDBW'23) — TaDA'23: Tabular Data Analysis Workshop, August 28 - September 1, 2023, Vancouver, Canada*
csd4361@csd.uoc.gr (E. Chrysostomaki); maria.stratigi@tuni.fi (M. Stratigi); vefthym@ics.forth.gr (V. Efthymiou); konstantinos.stefanidis@tuni.fi (K. Stefanidis); dp@ics.forth.gr (D. Plexousakis)

## 2.1. SQUIRREL Overview

Let $I$ denote a set of movies, $U$ a set of users, and $G \subseteq U$ a group of users. For each user $u \in G$, $B_u^j$ denotes an ordered list of recommended movies that have been generated by a single recommender system at a specific recommendation round $j$ for $u$. At round $j$, SQUIRREL chooses an appropriate aggregation function based on the current state of the group, to combine the individual members' recommendation lists $B_u^j$ into one list $GL_G^j$.

We apply SQUIRREL for a sequence of rounds to produce $\mu$ group recommendation lists for a group $G$, defined as $\mathcal{GR} = (GL_G^1, GL_G^2, ..., GL_G^\mu)$. At each round, we calculate each group member's individual satisfaction score, i.e., how relevant the items in the group list are for each member, as well as the aggregated group satisfaction score. Additionally, we calculate the disagreement between the group members, i.e., the variance between their satisfaction scores.

SQUIRREL can be defined as a Markov decision process $(S, \mathrm{A}, P, R)$, aiming to maximize the accumulative reward $R$ after each recommendation round $j$, where:

• $S$ describes the *group state*, expressed via the satisfaction scores of the group members. We keep an individual state for each group member $u$ at each round $j$, defined as $s_u^j = satO(u, j)$, where $satO$ is the overall satisfaction of user $u$ at round $j$. Briefly, it describes how relevant are the data items in the group recommendation list $GL_G^j$, compared to the best case scenario for the user $u$, which is their individual recommendations $B_u^j$.

• A is a set of distinct *actions*, i.e., different aggregation functions employed by the SQUIRREL model. SQUIRREL employs 6 aggregation functions (see [4]), ranging from a simple average to a far more complex SDAA [3].

• $P_a(s, s') = Pr(s_{j+1} = s' | s_j = s, a_j = a)$ is the probability to transition from state $s$ to state $s'$ during round $j$, under the action $a$.

• $R_a(s, s')$ is the reward gained from transitioning from state $s$ to state $s'$. We define two reward functions. First, we examine only the overall satisfaction of the group by averaging the individual satisfaction of all group members. Second, in addition to the overall satisfaction, we also consider the disagreement between group member, defined as the difference in the satisfaction scores between the most and least satisfied group member.

The goal of the model is to find a policy $\pi(a|s)$ that takes action $a \in A$ during state $s \in S$ to maximize the expected discounted cumulative reward after $\mu$ rounds:

$$max \mathbb{E}[R(\mu)], \tag{1}$$

where

$$R(\mu) = \sum_{t=0}^{\mu} \gamma R_\alpha(s, s'), \tag{2}$$

with $0 \le \gamma \le 1$.

At the beginning of round $j$, the group is given to a single-user recommender system that produces a recommendation list $B_u^j$ for each group member $u$. These lists are then given to the SQUIRREL model, where the agent observes the state of the environment $S_j$ and selects an appropriate action $\alpha_j$ to aggregate the lists $B_u^j$. This results in the transitioning of the model to the next state $S_{j+1}$, where we update the overall satisfaction of the users and calculate the reward $R_{j+1}$. Finally, SQUIRREL returns to the group the generated group recommendation list $GL_G^j$.

## 2.2. SQUIRREL Movies Features

In addition to the features described above and a Web UI, we have added some new functionalities to SQUIRREL Movies, such as recording the state of each group at each recommendation round, retrieving external movie data from MovieLens using MovieLens id, and providing user-friendly visual and textual explanations regarding the group recommendations. We cover only the latter next.

**Providing explanations.** SQUIRREL Movies provides both textual and visual (through graphs) explanations of the recommendations returned both for each round, as well as across rounds. For example, a user can change the reward function and get a brief, intuitive explanation of each aggregation function. In addition to textual explanations, the user can be presented a line chart with all the individual group members satisfaction not only for the current round, but also for the previous rounds. Further visual explanations include several bar charts with different group fairness measures (e.g., MaxMin) [4], justifying the recommendations of SQUIRREL Movies.

**Dataset and code.** We have utilized the 20M MovieLens Dataset [5], containing 20M ratings across 27,3K movies given by 138,5K users between January 1995 and March 2015 in tabular format (<user, item, rating, timestamp>), as processed in [4] to form diverse groups of 5 users. The data are available on https://github.com/mariaStratigi/SQUIRREL and the code of SQUIRREL Movies is publicly available on https://github.com/Eva-Chris/SQUIRREL-Web-App.

# 3. Conclusion

We have presented SQUIRREL Movies, demonstrating the adaptation of the SQUIRREL sequential group recommendations using reinforcement learning in the movies domain. SQUIRREL Movies provides visual and textual explanations, justifying its recommendations, and retrieves movie data from external sources, to help users decide on the movie to watch next. SQUIRREL Movies is just one of the many possible domain adaptations of SQUIRREL.

**Acknowledgement.** This work has received funding from the Hellenic Foundation for Research and Innovation (HFRI) and the General Secretariat for Research and Technology (GSRT), under grant agreement No 969.

## References

[1] E. Pitoura, K. Stefanidis, G. Koutrika, Fairness in rankings and recommendations: an overview, VLDB J. 31 (2022) 431–458.

[2] M. Stratigi, J. Nummenmaa, E. Pitoura, K. Stefanidis, Fair sequential group recommendations, in: SAC, 2020, pp. 1443–1452.

[3] M. Stratigi, E. Pitoura, J. Nummenmaa, K. Stefanidis, Sequential group recommendations based on satisfaction and disagreement scores, J. Intell. Inf. Syst. 58 (2022) 227–254.

[4] M. Stratigi, E. Pitoura, K. Stefanidis, SQUIRREL: A framework for sequential group recommendations through reinforcement learning, Inf. Syst. 112 (2023) 102128.

[5] F. M. Harper, J. A. Konstan, The movielens datasets: History and context, ACM Trans. Interact. Intell. Syst. 5 (2016) 19:1–19:19.