TECHNICAL UNIVERSITY OF CRETE
ELECTRICAL AND COMPUTER ENGINEERING

MASTER THESIS

Machine Learning and Social Choice Theory for Personalized and Group Recommendations

Author:
ERRIKOS STREVINIOTIS

Committee:
GEORGIOS CHALKIADAKIS, ASSOCIATE PROFESSOR - SUPERVISOR
MICHAIL G. LAGOUDAKIS, PROFESSOR
VASILIS SAMOLADAS, ASSOCIATE PROFESSOR

A thesis submitted in fulfilment of the requirements for the degree of Master of Science in Electrical and Computer Engineering awarded by the School of Electrical and Computer Engineering at Technical University of Crete

Chania, November, 2022

ΠΟΛΥΤΕΧΝΕΙΟ ΚΡΗΤΗΣ
ΣΧΟΛΗ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ
ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ

Μεταπτυχιακή Διατριβή

Μηχανική Μάθηση και Θεωρία Κοινωνικής Επιλογής για Προσωποποιημένες Συστάσεις και Συστάσεις σε Ομάδες

Συγγραφέας:
Ερρίκος Στρεβινιώτης

Εξεταστική επιτροπή:
Γεώργιος Χαλκιαδάκης, Αναπληρωτής Καθηγητής - Επιβλέπων
Μιχαήλ Γ. Λαγουδάκης, Καθηγητής
Βασίλειος Σαμολαδάς, Αναπληρωτής Καθηγητής

Χανιά, Νοέμβριος, 2022

Abstract

ERRIKOS STREVINIOTIS

Machine Learning and Social Choice Theory for Personalized and Group Recommendations

Recommender systems are software tools that provide assistance to individuals who lack experience or knowledge in order to overcome the information overflow problem. In this work, we introduce two recommender systems that employ Bayesian learning and the Kalman Filter algorithm respectively, along with mechanisms derived from Social Choice Theory, in order to learn users preferences and provide efficient personalized recommendations and group recommendations.

To facilitate user preference learning, we propose a novel, lightweight preference elicitation process, during which the user is presented with and asked to rate a small number of generic images that are related with the items under recommendation. We then exploit these ratings to guide our approaches to generate *beliefs* regarding the user's preferences. In order to model the high uncertainty that exists in such settings, our system represents *both users and items as multivariate normal distributions*.

On top of these, we employ several *multiwinner voting rules* from the social choice literature to the personalized recommendations problem. Specifically, we equip with such mechanisms our recommenders, allowing for effective personalized recommendations while promoting diverse results with respect to several features. We then focus on the *group recommendation problem*, by extending our approach and employing various preference aggregation mechanisms alongside with a multiwinner voting rule, namely the Reweighted Approval Voting (RAV). The application of multiwinner voting rules for these problems is, to the best of our knowledge, done for the first time in the literature.

Thus, in this thesis we tackle both the personalized and group recommendations problems; and we do so focusing on *the tourism domain*. We conduct a systematic experimental evaluation of our approaches by applying them on a real-world dataset of Points of Interest (POIs) in the popular touristic destination of Agios Nikolaos, Crete, Greece. Interestingly, we study the effectiveness of our approaches when we equip our system with prior knowledge regarding the (average) preferences of specific user types (i.e., tourists belonging in specific age groups), given data we collected via questionnaires from actual tourists visiting the city of Agios Nikolaos.

Our experimental results *(i)* highlight the ability of our systems to successfully produce personalized recommendations that match the specific interests of a single user; *(ii)* confirm that the employment of prior knowledge regarding the preferences of tourists, based on their demographics, guides our recommender to avoid

the cold-start problem; *(iii)* demonstrate that the use of multiwinner mechanisms allows for diverse recommendations with respect to travel-related features, and increased system performance in the case of limited user-system interactions; and *(iv)* show that the use of multiwinner mechanisms allows for *fair group recommendations* with respect to the well-known m-PROPORTIONALITY and m-ENVY-FREENESS metrics. Last but not least, our personalized Bayesian recommendation algorithm is incorporated in a real-world mobile tour-planning application for Agios Nikolaos, Crete.

Περίληψη

Ερρίκος Στρεβινιώτης

Μηχανική Μάθηση και Θεωρία Κοινωνικής Επιλογής για
Προσωποποιημένες Συστάσεις και Συστάσεις σε Ομάδες

Τα συστήματα συστάσεων είναι εργαλεία λογισμικού τα οποία παρέχουν βοήθεια σε χρήστες οι οποίοι έχουν έλλειψη εμπειρίας ή γνώσης προκειμένου να ξεπεράσουν το πρόβλημα της πλεονάζουσας πληροφορίας. Σε αυτή την εργασία, παρουσιάζουμε δυο συστήματα συστάσεων τα οποία χρησιμοποιούν αντίστοιχα Μπαϋεσιανή μάθηση και τον αλγόριθμο του φίλτρου Κάλμαν, σε συνδυασμό με μηχανισμούς που προέρχονται από τη θεωρία κοινωνικής επιλογής προκειμένου να μάθουν τις προτιμήσεις των χρηστών ώστε να παρέχουν προσωποποιημένες συστάσεις και συστάσεις σε ομάδες χρηστών.

Η μάθηση των προτιμήσεων χρησιμοποιεί μια καινοτόμα και όχι απαιτητική προς τον χρήστη διαδικασία εκμαίευσης προτιμήσεων, κατά την οποία ο χρήστης παρέχει μια βαθμολογία για μερικές γενικές εικόνες, οι οποίες σχετίζονται με τα αντικείμενα τα οποία μπορούν να προταθούν, που του/της παρουσιάζουμε. Στη συνέχεια αξιοποιούμε αυτές τις βαθμολογίες ώστε οι προσεγγίσεις μας να παράγουν κάποιες πεποιθήσεις σχετικά με τις προτιμήσεις του χρήστη. Προκειμένου να μοντελοποιήσουμε την υψηλή αβεβαιότητα η οποία εμπεριέχεται σε τέτοια περιβάλλοντα, αναπαριστούμε τόσο τους χρήστες όσο και τα προς σύσταση αντικείμενα ως κανονικές κατανομές πολλών μεταβλητών.

Επιπρόσθετα, εφαρμόζουμε διάφορους κανόνες ψηφοφοριών πολλών νικητών από τη θεωρία κοινωνικής επιλογής στο πρόβλημα των προσωποποιημένων συστάσεων. Συγκεκριμένα, εξοπλίζοντας το σύστημά μας με τέτοιους μηχανισμούς μπορούμε να παράγουμε αποτελεσματικές προσωποποιημένες συστάσεις ενώ παράλληλα προάγουμε ποικίλες συστάσεις σε σχέση με διάφορα χαρακτηριστικά. Στη συνέχεια εστιάζουμε στο πρόβλημα των συστάσεων σε ομάδες επεκτείνοντας τις προσεγγίσεις μας και εφαρμόζοντας διάφορες στρατηγικές συνάθροισης μαζί με ένα κανόνα ψηφοφοριών πολλών νικητών, ο οποίος ονομάζεται Reweighted Approval voting (RAV). Εξ όσων γνωρίζουμε, η χρήση κανόνων ψηφοφοριών πολλών νικητών για τα παραπάνω προβλήματα έγινε για πρώτη φορά στην βιβλιογραφία.

Συνοψίζοντας, η μεταπτυχιακή μας εργασία εστιάζει στα προβλήματα των προσωποποιημένων συστάσεων και συστάσεων σε ομάδες. Έχει δε ως πεδίο εφαρμογής τον τομέα του τουρισμού. Διεξάγαμε μια συστηματική πειραματική αξιολόγηση των προσεγγίσεών μας τις οποίες εφαρμόσαμε σε πραγματικά δεδομένα από σημεία τουριστικού ενδιαφέροντος της πόλης του Αγίου Νικολάου Κρήτης, σημαντικού τουριστικού προορισμού. Επισημαίνουμε ότι μελετήσαμε την αποτελεσματικότητα των προσεγγίσεων μας όταν τις εξοπλίζουμε με πρότερη γνώση σχετικά με τις γενικές προτιμήσεις συγκεκριμένων κατηγοριών χρηστών (χρηστών οι οποίοι ανήκουν στην ίδια ηλικιακή ομάδα), με βάση δεδομένα τα οποία συλλέχθηκαν από ερωτηματολόγια πραγματικών τουριστών στην πόλη του Αγίου Νικολάου.

Στα πλαίσια της εργασίας μας διεξάγαμε μια συστηματική πειραματική αξιολόγηση των προσεγγίσεών μας τις οποίες εφαρμόσαμε σε πραγματικά δεδομένα της πόλης του Αγίου Νικολάου στη Κρήτη. Τα πειραματικά μας αποτελέσματα (i) δείχνουν την ικανότητα των προσεγγίσεών μας να παράγουν αποτελεσματικά προσωποποιημένες συστάσεις οι οποίες ταιριάζουν με τα ενδιαφέροντα ενός συγκεκριμένου χρήστη, (ii) επιβεβαιώνουν ότι η αξιοποίηση πρότερης γνώσης σχετικά με τις προτιμήσεις των τουριστών, με βάση τα δημογραφικά του χαρακτηριστικά, βοηθούν τα συστήματα συστάσεων μας να αποφύγουν το cold-start πρόβλημα, (iii) δεικνύουν ότι η χρήση κανόνων ψηφοφοριών πολλών νικητών επιτρέπει ποικιλόμορφες συστάσεις σε σχέση με διάφορα χαρακτηριστικά τα οποία σχετίζονται με το τουρισμό, ενώ αυξάνουν την απόδοση του συστήματός μας στη περίπτωση που η αλληλεπίδραση ενός χρήστη με το σύστημα είναι περιορισμένη, και (iv) επιδεικνύουν ότι η χρήση κανόνων ψηφοφοριών πολλών νικητών επιτρέπει δίκαιες συστάσεις σε ομάδες χρηστών σε σχέση με τις ευρέως γνωστές μετρικές m-PROPORTIONALITY και m-ENVY-FREENESS. Τέλος, ο Μπαϋεσιανός αλγόριθμος προσωποποιήμενων συστάσεων που αναπτύξαμε έχει ενσωματωθεί σε μια πραγματική εφαρμογή κινητών για τον προγραμματισμό τουριστικών περιηγήσεων στην πόλη του Αγίου Νικολάου Κρήτης.

Acknowledgements

First and foremost, I would like to deeply thank my supervisor, Professor Georgios Chalkiadakis, for all his guidance and trust throughout this master, and for motivating me to continue my studies on topics related to Artificial Intelligence. I also thank the members of the Intelligent Systems Laboratory, Dimitris, Iason, Giannis and Antonis, for our fruitful discussions. Second, I wish to thank the members of the committee, Professor Michail G. Lagoudakis and Professor Vasilis Samoladas, for their helpful comments and time. Finally, I wish to thank all of my friends and family for their unconditional love and support, not only throughout my studies, but also in every part of my life.

This research has been co-financed by the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH-CREATE-INNOVATE B cycle (project code: T2EDK-03135). Moreover, this M.Sc program was supported by the Onassis Foundation - Scholarship ID: G ZR 012-1/2021-2022.

ONASSIS
FOUNDATION

Contents

List of Figures

List of Tables

List of Algorithms

Chapter 1

Introduction

The explosive growth of information on the World Wide Web, as well as the rapid expansion of e-services, has presented users with a plethora of options, which frequently leads to more complex decision-making. For example, individuals often have to decide which music album they will listen to their road-trip, which movie they will watch on the theatre or even which travel destination they will visit during their next vacations. The vast array of options that is offered to users gives rise to the development of intelligent systems that can provide assistance to individuals who lack experience or knowledge in order to select the most appropriate option with respect to their interests or needs. This type of intelligent system are commonly known as *recommender systems*. In general, recommender systems, or recommendation systems, are tools that help users browse and retrieve relevant information or items of interest (e.g., products or services) from large collections [66]. Such systems find applications in various domains, e.g., movie domain, music domain, fashion domain and travel domain to name a few. LinkedIn, Netflix, YouTube, Spotify, TripAdvisor, Booking and Amazon are all examples of successful real-world recommender systems in use. There is plethora of reasons that drive service providers to employ such technologies. In particular recommender systems: *(i)* increase the number of items sold; *(ii)* sell more diverse items; *(iii)* increase the user satisfaction; *(iv)* increase user fidelity; and *(v)* better understand what the user wants [66].

In the literature the most established recommendation systems exploit user ratings over a large number of items in order to predict the ratings for unobserved items. As Figure 1.1 illustrates, the most common approaches used for recommendations can be categorized as: *(i)* Collaborative Filtering; *(ii)* Content-based; and *(iii)* Hybrid techniques. In more detail, Collaborative Filtering techniques produce recommendations based on user's social interaction and rankings provided by similar users, while Content-based methods recommend items based on the description of an item and a profile of user's preferences. However, both techniques have several limitations that reduce their effectiveness. For example, Collaborative Filtering technique suffers from the well-known *cold-start* problem while they face *scalability* issues [34]. In contrast, Content-based approaches suffer from *sparsity* of data and limited content analysis [34]. As a result, various Hybrid recommenders [59, 12] have been introduced in the literature in order to solve the aforementioned issues while helping to increase the recommendations accuracy. In such approaches the final recommendations are made based on an ensemble of different approaches.

Now, as mentioned earlier, recommender systems have been employed in various domains due to their beneficial properties [66]. One of the most common cases that such systems are applied in the tourism domain. Notably, there is a rising interest regarding the development and the study of tourism recommender systems. Figure 1.2 demonstrates the number of tourism-related recommendation system articles over a 30-year period [70]. Specifically, recommender systems in the tourism domain play the key role of digital guides for the various activities that a tourist destination might provide to visitors based on their preferences [46, 20, 18, 42]. Tourism recommenders can be broadly categorized as hotel recommenders, restaurant recommenders, tourism recommenders that are associated with group recommendations, tour planning (or travel packages); and tourist attraction recommenders (i.e., points of interests, museums, etc.) [14]. Generally, the majority of recommender systems focus on the *accuracy* of the recommended items with respect to the preferences of a targeted-user. However, the accuracy of the recommended items and the satisfaction of a user are not identical notions [92]. Specifically, many aspects influence users' perception regarding the recommended items, that usually are ignored by classic approaches [91]. One of these key-factors is the notion of *diversity* among the recommended items. For example, consider a user that loves the Italian cuisine and has visited a tourist destination in Italy. A recommender that constantly recommends pizzerias to the user, can be considered as an accurate one, but does not result to a satisfactory user experience, since this user would like to visit other kinds of POIs as well, e.g., a monument or a museum. Thus, it is obvious that this lack of diversity decreases the performance of the recommender, even if this system provides accurate recommendations with respect to user's preferences. As such, the development of more sophisticated systems that consider both accuracy and diversity-based criteria during the production of the final recommendations is crucial.

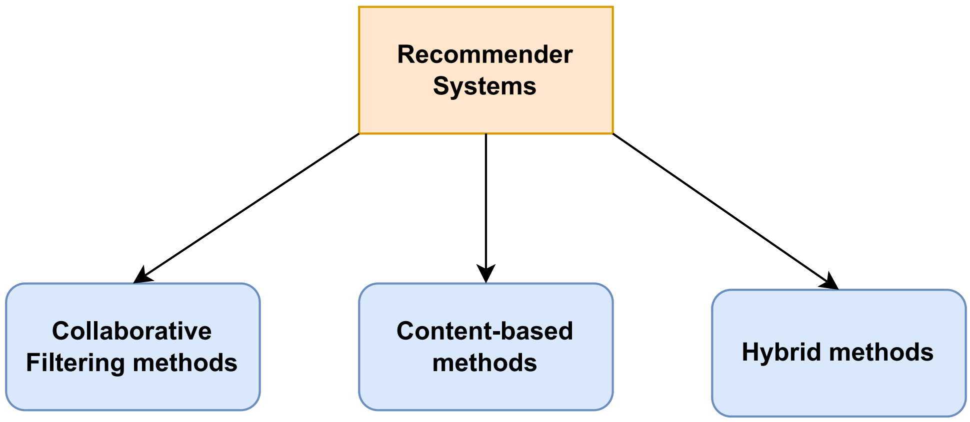

FIGURE 1.1: The main methods that recommender systems employ in the literature.

Moreover, most of recommender systems research focuses on recommendations for a single user, while in real life there is a plethora of scenarios where users form groups in order to experience some product or service, e.g., music in a car ride or a movie in the theater. In the tourism domain, in particular, group recommendations come in naturally since usually users choose to travel with company, e.g., friends or family. Thus, a recommender system in such a domain should take this aspect into account [21]; and, importantly, needs to ensure that recommendations are *fair* [73] with respect to the preferences of the individual group members. In general, the generation of group recommendations by a system is achieved by employing an *aggregation mechanism* that considers individuals' preferences [48]. However, there is a multitude of alternative ways to produce recommendations for groups of users [35], e.g., find the recommendations for each member of the group separately and produce the group recommendations by selecting items based on a preference aggregation technique; or build a group recommender model by merging group members into one [39].

On top of that, we should mention that in the complex domain of tourism, recommender systems have to face some challenging problems that do not exist in other fields. First, in the tourism product most of the times the user-items ratings are very sparse compared to other domains (e.g., the movies domain), and as such the employment of classic recommender system approaches, e.g., collaborative filtering techniques, can be a complicated task [21]. Thus, tourism recommender systems *suffer from a continuous state of "coldness"* [45], i.e., recommenders can not produce efficient personalized recommendations since no available data are provided [23]. Additionally, many tourists have limited time when they visit a travel destination, and as such a recommender system should be able to provide efficient recommendations in order to maximize their satisfaction with a light-weight user-system interaction. Finally, visitors commonly use their mobile phones as a tool in order to exploit an unknown destination. As such, the development of novel tourism recommender systems that employ *computationally efficient* algorithms that can run on the mobile devices of the users is of utmost significance.

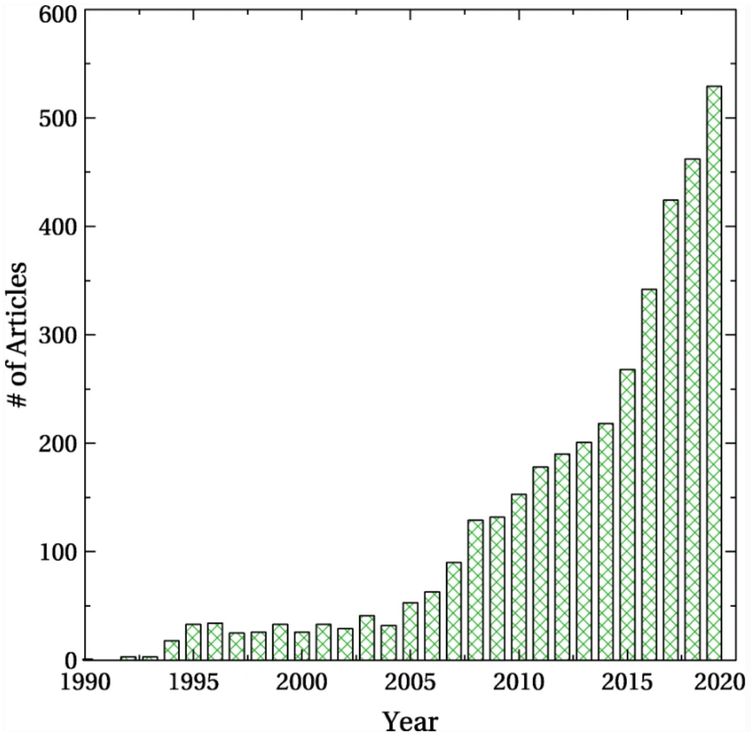

FIGURE 1.2: Distribution of tourism recommendation system articles by years (Figure taken from [70]).

1.1 Our Contributions

In this work we introduce two personalized recommender systems for the tourism domain that operate as travel guides for the activities that a tourist destination provides, helping visitors to discover POIs that are related to their interests. Specifically, our proposed systems employ *Bayesian learning* and the *Kalman Filter* algorithm accordingly. In more detail, real-world data on *points of interests (POIs)* and users' preferences were collected exploiting: *(i)* information and data provided by the Municipality of Agios Nikolaos; *(ii)* online sources; and *(iii)* questionnaires that were filled by tourists. We propose a novel picture-based process to elicit user's preferences, by efficiently updating system's beliefs regarding the user's interests. Furthermore, we equip our proposed recommenders with *prior knowledge*, obtained via questionnaires from real-world tourists, and experimentally study the effectiveness of each approach.

On top of that, we put forward a novel recommendation mechanism that can be used to increase the *diversity* of the final personalized recommendations (instead of simply "greedily" recommending the POIs so-far-perceived-as-best). Such diversity is important, and contributes to the overall quality of the recommendations—especially when these do not rely on ratings of other tourists/users, but are strictly personalized (in the sense that they are produced given only a few interactions with the user in question, as is the case in our system). Our mechanism is inspired by multiwinner election rules used in *social choice theory* [11, 26, 4], and which come complete with theoretical guarantees regarding properties satisfied by the set of the election winners (e.g., proportionality of the representation). We note that to the best of our knowledge such mechanisms have been studied as theoretical tools, as they have not been employed in real-world systems, except from the work in [30], where authors utilize such mechanisms in order to produce results (or recommendations) in a search engine. Specifically, our approach creates a *personalized election* in order to "elect" a set of k "winners" (corresponding to the k final recommendations made to the targeted tourist-user). We study several voting rules and show that such an approach can be useful since it provides *diverse* results with respect to travel related features—e.g., *Culture, History, Cuisine*, and other characteristics that a tourist attraction may provide. For instance, assume that a user has due to various reasons (e.g., could have been pressed with time) rated highly only cultural POIs, when visiting a tourist destination. A mechanism that produces *diverse* final recommendations (i.e., a set of "election winners"), would present to the user POIs that are related to *various* categories, and not just Culture-related ones—a fact that is in fact expected to enhance the tourist experience. As such, our approach can be thought of as aiming to tackle the manifestation of the classic *exploration vs exploitation problem* [80] in this domain.

Moreover, we extend our system on order to be able to solve the group recommendation problem by employing various aggregation strategies along with a multiwinner elcetion mechanism. Specifically, to the best of our knowledge, we employ for the first time in the literature, the *Reweighted Approval Voting (RAV)* multiwinner election mechanism in a real-world system in order to face the challenging problem of group recommendations in the tourism domain.

Finally, we conduct a systematic experimental evaluation of our recommender systems by applying them on a real-world dataset of Agios Nikolaos city. We note that this dataset was created for the needs of a real-world tour planning recommender system that is currently being developed in collaboration with an e-commerce company and the Municipality of Agios Nikolaos. Our results confirm: *(i)* the effectiveness of our approaches by highlighting their ability to provide top personalized recommendations with respect to user's preferences and interests; *(ii)* using multiwinner elections leads to improved system performance when the user-system interactions are limited, while the advantage thus provided decreases or evaporates with increased user-system interactions; and *(iii)* RAV multiwinner election mechanism outperforms several other well-known aggregation mechanisms for the problem of group recommendations with respect to standard fairness metrics derived from the recommender systems literature. Last but not least, our personalized Bayesian recommendation algorithm is incorporated in a real-world mobile tour-planning application for Agios Nikolaos city.

Parts of this research were published in two (2) international conferences [78, 77], while the constructed dataset was also used in [93] (as seen in Appendix A).

1.2 Thesis Structure

In Chapter 2 we provide the necessary background for this thesis. We present the basic notions and categories of recommender systems and we review several Bayesian, personalized and group recommendation systems. Moreover, we give an overview of Markov chains and hidden Markov models. Then we briefly describe the field of social choice theory and multiwinner elections. In Chapter 3 we describe our proposed recommender systems along with the mechanism that produce diverse final recommendations and the picture-based elicitation process. Moreover, we analyze how we extend our system in a way that is able to handle the challenging problem of group recommendations. In Chapter 4 we conduct a systematic experimental evaluation of our systems in the tourism domain and prove the effectiveness of our approaches with respect to several metrics. Finally, in Chapter 5 we summarize our contributions and the corresponding results and describe some interesting lines for future work.

Chapter 2

Background & Related Work

In this chapter, we provide the necessary background for this thesis. Specifically, we begin by discussing some baseline notions for recommender systems and presenting some Bayesian recommendation techniques. Furthermore, we briefly review some recommender systems for the problems of personalized and group recommendations. Then we provide an overview of Markov chains, hidden Markov models and the Kalman Filter algorithm. Finally, we describe some well-known notions of the social choice theory and present several rules for multiwinner elections.

2.1 Recommender Systems

Recommender systems are software tools that employ machine learning algorithms and techniques in order to help users find items or services that are closely related to their interests. Such systems provide recommendations by collecting information and data regarding the preferences of a targeted-user for a specific set of items [9]. This information, which is also called feedback, can be collected either *explicitly* or *implicitly* [43]. In the case of explicit user feedback, the system asks individuals to evaluate a specific item by providing a rating in a numeric scale or by providing a like/dislike functionality. On the other hand, systems collect information implicitly by monitoring user's activity, e.g., user's visiting time on a web page, items that they recently purchase, applications that they downloaded, etc.

Once the recommender has collected information regarding user's preferences it has to produce an estimation of the *utility* of a specific item for a given individual [89]. Formally, given a set of users U and a set of items I a recommender system predicts the utility of any user $u \in U$ for any item $i \in I$ and produces a list of recommended items $list_{rec}$ [1]:

$$f : U \times I \rightarrow list_{rec} \tag{2.1}$$

Traditionally, recommender systems select the items that maximise user's utility [1]:

$$\forall u \in U, \quad argmax_{i \in I} f(u, i) \tag{2.2}$$

As mentioned earlier, recommenders can predict the utility of a user for a specific item by employing any machine learning technique. e.g., Collaborative Filtering [71, 22, 2], Content-based [62, 44, 93], Neural Networks [83], etc. Notably, Zhang et al. [89] discuss how artificial intelligence can help to further improve the performance of recommender systems, while authors describe several methodologies that are employed by such systems.

Moreover, recommenders can be categorized as *personalized* and *non-personalized* based on the type of recommendations they provide [66]. In more detail, personalized recommenders exploit each user's prior purchase history in order to produce a list of recommended items that is adjusted on the interests of this specific individual. By contrast, non-personalized recommenders suggest items that are generally popular among the users. For example, TripAdvisor provides a list of the top-N attractions at any travel destination around the world. However, in this thesis we focus on the problem of personalized recommendations only.

2.1.1 Bayesian Recommender Systems

In real-world many features of recommender systems, including user preferences, are uncertain [32]. This uncertainty can be modelled either as *strict uncertainty* or as a *probabilistic* one. In the case of strict uncertainty, recommender systems combine knowledge from several sources by making a number of assumptions without considering the "strength" of those assumptions. By contrast, under probabilistic (or Bayesian) uncertainty, such systems have a distribution over a hypothesis. This results to more flexible systems which can mathematically model the uncertainty that exist in the setting. As such, over the years many *Bayesian recommender systems* have been proposed for various domains. Notably, Bayesian methods are able to provide efficient recommendations and, most importantly, such techniques can be applied for real-time mobile recommendation services [16].

To begin with, Rendle et al. [65] focused on the scenario where the recommender system is able to exploit implicit user feedback in order to generate a personalized ranking list over a set of available items. Their approach assumes that a user prefers a "clicked" item over the "unclicked" ones, and a ranked-list of items is produced after a number of user-system interactions. In more detail, an optimization criterion was introduced, called the $BPR - Opt$ criterion, which was derived from the maximum posterior estimator for optimal personalized ranking. Then, authors proposed the $LearnBPR$ algorithm for optimisation of models with respect to the aforementioned criterion. Specifically, this learning algorithm employs stochastic gradient descent techniques along with bootstrap sampling of training triples. Their approach was evaluated on two different datasets, i.e., related to online shopping and movies domain respectively, proving the effectiveness of their approach against other well-known learning algorithms. Later, some variants of their approach have been proposed in the literature [15, 58, 33, 63].

A Bayesian hierarchical model approach was proposed by Zhang and Koren [90] for personalized recommendations. The main idea of their model is that a system with a large number of users can effectively learn the preferences of a new user (that has limited interaction with the system) by exploiting preference-related data from other users via the hierarchical model. Intuitively, their system attempts to learn the most likely model parameters for the provided data. Additionally, a modified *Expectation-Maximization (EM)* algorithm [54] was employed in order to accelerate the learning procedure of many individuals' profiles. An experimental evaluation on real-world data from Netflix and MovieLens showed the efficiency and the effectiveness of their model.

Moreover, many researchers utilise a bipartite graph representation in order to capture the user-item interactions. However, many graph-based recommender models ignore the existence of uncertainty in such representations. Thus, Sun et al. [79] introduced a novel training framework based on Bayesian Graph Neural Networks (BGNNs), that addresses this issue by employing Bayesian Graph Convolutional Neural Networks. Additionally, their approach was evaluated on three well-known datasets. The results showed that BGNNs architectures: *(i)* achieve accurate and diverse recommendations; *(ii)* outperform state-of-the-art graph-based recommenders with respect to several well-known metrics derived from the recommender system domain.

Finally, the work that is most relevant to ours, in the sense that it employs *Bayesian learning* and uses a common representation among the users and the items is the one of [6]. In particular, their system provides personalized recommendations for the movie domain, while it learns the user preferences by presenting short films, i.e., movie trailers, to the user and performing Bayesian updating of its beliefs regarding the interests of the users. Notably, their approach does not suffer from the well-known problem of "cold-start" [41] that other techniques face due to data scarcity, e.g., Collaborative Filtering approaches.

2.1.2 Personalized Recommendations

In general, domains that contain various items of different types, which are highly connected with users' personal preferences and interests—for instance the tourism domain which requires recommendations for leisure and cultural POIs—call for the development of highly efficient recommender systems which can also cope with limited user interaction and feedback. As such, personalized recommenders have been employed in order to help users exploit the vast number of items that exist in a platform by presenting to them only few items that are highly related to user's interests. Thus, personalized recommenders for movies [88], news [82], music [61] and fashion [76] have been introduced among others in the literature. Regarding the tourism domain there is a plethora of tourism- or travel-oriented recommenders, potentially classified in different categories, as listed in [14]. Most of those systems recommend POIs that correspond to touristic attractions (e.g., restaurants, hotels, historical sites or museums), that are ideally highly connected with each individuals' preferences. Here we brief-review a few representative such systems focusing on the tourism domain.

Now, an extensive overview of tourism recommender system algorithms is provided by Borràs et al. [10]. The authors discuss alternative user interfaces, recommendation techniques and additional services that such system may provide. Sánchez and Bellogín [68] presented a survey of recommender systems that employ Location-Based Social Networks for POIs recommendations, while discussing open challenges for such approaches. A restaurant recommender system for mobile devices introduced by Zeng et al. [87]. Specifically, their approach adopts a user preference model based on the restaurants that the user has already visited in the past, and produce the final recommendations by exploiting the exact location of the user. Kbaier et al. [40] develop a Hybrid recommender system that generates personalized recommendation by combining three different techniques, i.e., collaborative filtering, content-based filtering and the demographic filtering. A personalized tour planner, called eCOMPASS, was proposed by Gavalas et al. [28]. In their work, the system recommended personalized tours, i.e., a route that contains several POIs, by also considering: *(i)* the weather forecast, and *(ii)* the possibility that the visitor can use the public transit to move among the POIs. PersTour [42] is another recommendation system that was developed for recommending several POIs to the visitors by considering *(i)* the popularity of a specific POI; and *(ii)* the preferences of a specific tourist. Additionally, PersTour modifies the tourist's visiting time for a specific POI, based on how relevant this POI is to her interests. Mishra et al. [53] introduced a recommender system that employs the Sentiment Intensity Analyzer (SIA) in order to find the most appropriate POIs for a specific user by exploring the available reviews of other visitors. In [74] authors introduced a picture-based recommender approach which suggests tourist destinations for a specific individual. Specifically, the user selects any set of pictures which is then used as input to computer vision models that generate a profile which describes the preferences of the tourist. Massimo and Ricci [46, 47] studied the problem of *next POIs recommendations*, i.e., given a set of POIs that a specific tourist has already visit provide recommendations for the not-yet-visited POIs. Specifically, Inverse Reinforcement Learning (IRL) techniques were employed in order to learn the reward function and the optimal POI selection policy. Finally, Gavalas et al. [29] provide a review of the state-of-the-art techniques for *mobile recommender systems* in the tourism domain.

2.1.3 Group Recommendations

Traditionally, most recommender systems research focuses on individuals by recommending items that maximise users' satisfaction. However, users often interact with one another, forming groups. This social aspect gives rise to the development of *group recommender systems*. In the literature, several recommenders have been proposed for groups in various domains, such as movies, music, TV programs, travel, etc. [19]. Briefly, in the movies domain such systems may operate by exploiting users' social and behavioral data [69], or by calculating the probability that a user likes a movie or not [31, 86]. Additionally, some Collaborative Filtering approaches have been introduced [7, 38]. MusicFX [50] is a group recommender system for music in shared environments, e.g., in gyms, which is also extended to other domains, i.e., restaurants [49].

In terms of work in our domain of interest, there is a number of tourism- or travel-oriented group recommender systems. McCarthy et al. [51] introduced the CATS travel recommender system for groups of users. Specifically, CATS is a collaborative advisory travel system that recommends travel destinations to groups of at most 4 individuals. An extension of CATS, was proposed in [67], where an extra module of negotiation was added. Moreover, *Bayesian methods*, which maintain and exploit probabilistic beliefs regarding user preferences, are able to provide high quality recommendations for both individual tourists and groups; and, most importantly, such techniques can be applied for real-time mobile recommendation services [16]. The travel recommender system in [16] exploits data from the "community-contributed" photos by giving tags to the groups. Moreover, users can be categorized based on some user specific features (i.e., age, gender, etc.); whereas groups are categorized based on the type of the formed group (i.e., group of friends, family, etc). Finally, [60] employs Bayesian networks to recommend restaurants to groups of people in mobile environments.

2.2 Gaussian or Normal Distributions

Here, we provide a background about the Gaussian distributions which are the basis for multivariate normal distributions that we extensively use in this work. The

Gaussian (or normal) distribution is a widely used model for the distribution of continuous variables. For a single variable x, the Gaussian distribution can be written in the form:

$$\mathcal{N}(x|\mu,\sigma^2) = \frac{1}{(2\pi\sigma^2)^{1/2}} \cdot exp\Big\{ -\frac{1}{2\sigma^2}(x-\mu)^2 \Big\}$$

where μ is the mean and σ^2 is the variance. Figure 2.1 illustrates a univariate normal distribution with $\mu = 0$ and $\sigma^2 = 1$.

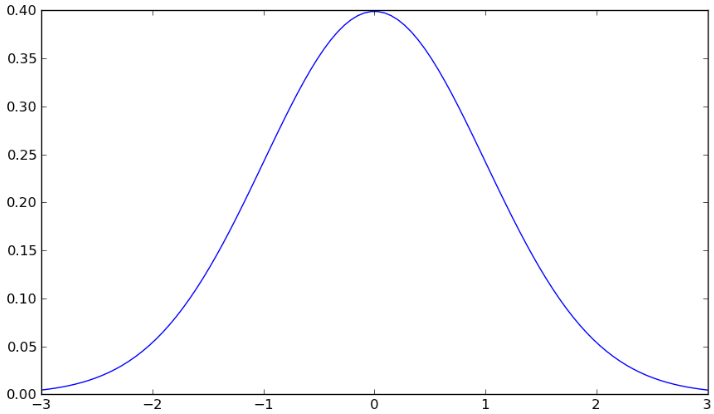

FIGURE 2.1: A univariate normal distribution with $\mu = 0$ and $\sigma^2 = 1$.

The *multivariate normal distribution* is a generalization of the one-dimensional normal distribution to higher dimensions. Formally, for a $D-$dimensional vector x, the multivariate normal distribution can be written as:

$$\mathcal{N}(x|\mu,\Sigma) = \frac{1}{(2\pi)^{D/2}} \frac{1}{|\Sigma|^{1/2}} \cdot exp\Big\{ -\frac{1}{2}(x-\mu)^T\Sigma^{-1}(x-\mu) \Big\}$$

where μ is a $D-$dimensional mean vector, Σ is a $D \times D$ covariance matrix, and $|\Sigma|$ is the determinant of Σ. Figure 2.2 illustrates a multivariate (or bivariate) normal distribution with $\mu = [0, 1]$ and $\sigma^2 = \begin{bmatrix} 1 & -0.5 \\ -0.5 & 1.5 \end{bmatrix}$.

2.3 Bayesian Inference

In many problems researchers employ probabilistic learning in order to estimate the parameters of an unknown distribution. One common technique that can be utilized for such a purpose is the Bayesian learning (or Bayesian updating). Specifically, Bayesian learning combines any prior distribution that represents our belief on the values of the unknown model's parameters with the likelihood of the newly observed data, in order to compute a posterior distribution, i.e., an updated belief. The prior and the posterior distributions are called *conjugate distributions* if they belong to the same probability distribution family, while the prior is called *conjugate prior*.

Now, the *Normal-Inverse Wishart (NIW)* [72] is a multivariate four-parameter family of continuous probability distributions that has the desired property of *conjugacy* of a multivariate normal distribution (see Figure 2.2) with unknown mean vector and covariance matrix. Thus, NIW distribution can play the role of the prior representing

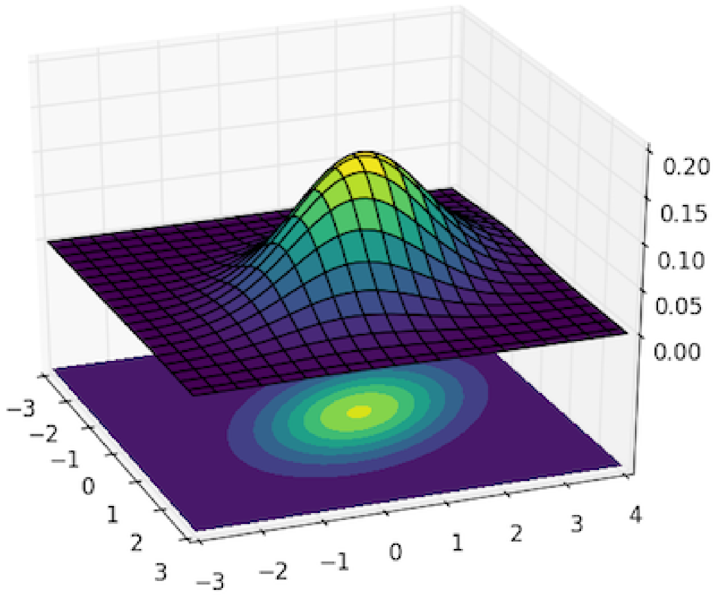

FIGURE 2.2: A multivariate normal distribution ($D = 2$) with $\mu = [0, 1]$ and $\sigma^2 = \begin{bmatrix} 1 & -0.5 \\ -0.5 & 1.5 \end{bmatrix}$.

recommender's belief over the unknown users model, as is done in [6], where both users and items are represented as multivariate normal distributions and the recommender is unaware of each user distribution. Generally, the use of conjugate priors offers a closed-form for the computation of the posterior distribution, resulting in a computationally efficient Bayesian updating procedure [6]. Formally, we can update the prior hyperparameters—κ_0, μ_0 (the mean vector), ν_0 (degrees of freedom), and Ψ_0 (the precision matrix)—using samples drawn from the data to get the posterior ones, as follows:

$$\mu_n = \frac{\kappa_0}{\kappa_0 + n} \cdot \mu_0 + \frac{n}{\kappa_0 + n} \cdot \bar{x} \tag{2.3}$$

$$\kappa_n = \kappa_0 + n \tag{2.4}$$

$$\nu_n = \nu_0 + n \tag{2.5}$$

$$\Psi_n = \Psi_0 + S + \frac{\kappa_0 \cdot n}{\kappa_0 + n} \cdot (\bar{x} - \mu_0) \cdot (\bar{x} - \mu_0)^T \tag{2.6}$$

$$S = \sum_{i=1}^{n} (x_i - \bar{x}) \cdot (x_i - \bar{x})^T \tag{2.7}$$

where $\bar{x}$ is the sample mean, n is the number of the samples, x_i are the samples drawn from the data, S is a scatter matrix.

Finally, we can use an *Inverse Wishart* and a *Normal* distribution to derive the *covariance matrix* Σ and the *mean vector* μ given the updated beliefs, as follows:

$$\Sigma \sim \mathcal{IW}(\Psi_n, \nu_n) \tag{2.8}$$

$$\mu|\Sigma \sim \mathcal{N}(\mu_n, \Sigma/\kappa_n) \tag{2.9}$$

2.4 Markov Chain

In probability theory a *stochastic process* is defined as a set of random variables $\mathbf{X}_t \in \mathcal{D}$, where $\mathbf{X}_t$ denotes the *state* of the process at time t and $\mathcal{D}$ is the set of all states. A Markov chain (or a Markov process) is a stochastic process, named after Russian mathematician Andrey Markov, which also satisfies the Markov property. Specifically, Markov chains make the assumption that the process's future state $\mathbf{X}_{t+1}$ depends only on its current state $\mathbf{X}_t$, and not on any of the previous $\mathbf{X}_{t-1}, \ldots, \mathbf{X}_0$ [52]. This property is known as *Markov property* and is also characterized as "memorylessness" property since the model has no memory. Formally:

$$Pr\{X_{t+1} = j | X_0 = i_0, X_1 = i_1, \cdots, X_{t-1} = i_{t-1}, X_t = i\} = Pr\{X_{t+1} = j | X_t = i\} = p_{ij} \tag{2.10}$$

At every transition of the Markov chain (from t to $t+1$), the process will visit state j with a probability p_{ij}, where i is the realization of $\mathbf{X}_t$ (i.e., the realization of the state at timestamp t). In general, it holds that:

$$0 \leq p_{ij} \leq 1 \tag{2.11}$$

$$\sum_j p_{ij} = 1 \tag{2.12}$$

As such we construct a *transition matrix* $\mathbf{P}$ that contains the transition probabilities from any state $i \in \mathcal{D}$ to any state $j \in \mathcal{D}$. Figure 2.3 illustrates an example of a Markov chain used for the prediction of weather showing the states and the probabilities of each transition.

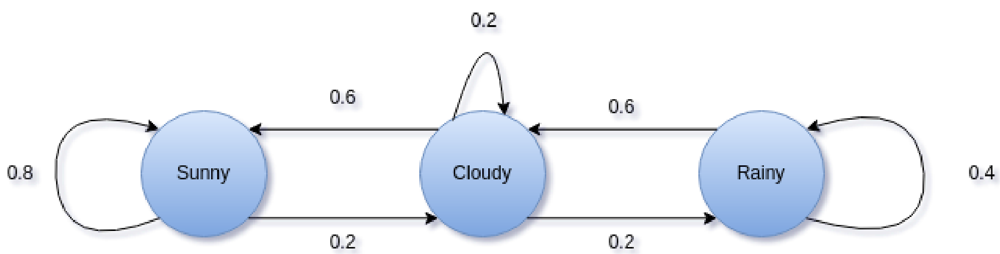

FIGURE 2.3: A Markov chain for weather prediction.

Generally, such models find application in various domains, such as computer science, biology and medicine, economics, speech recognition, etc., due to their desirable properties.

2.4.1 Hidden Markov Models

Hidden Markov models (HMM), like Markov chains, are graphical models which consist of two types of nodes (states), namely hidden and observable. HMMs have been widely employed in many applications, including voice recognition [64] and biological sequence analysis [85], and are well known for their efficiency in modeling dependencies between adjacent events due to their structure. More in detail, a state (or node) in HMM can be either *hidden* or *observable*. We denote as $\mathbf{X}_t$ the states that are hidden, i.e., we can not observe their values, and with $\mathbf{Y}_t$ the states that are observable to us. For instance, given the type of clothes that an individual is wearing each day (observations) our system tries to find the corresponding weather for this day (hidden state), i.e, cold or hot day. We also mention that HMMs satisfy the Markov property similar to the simple Markov chains. Figures 2.4 and 2.5 depict a directed and an undirected HMM respectively.

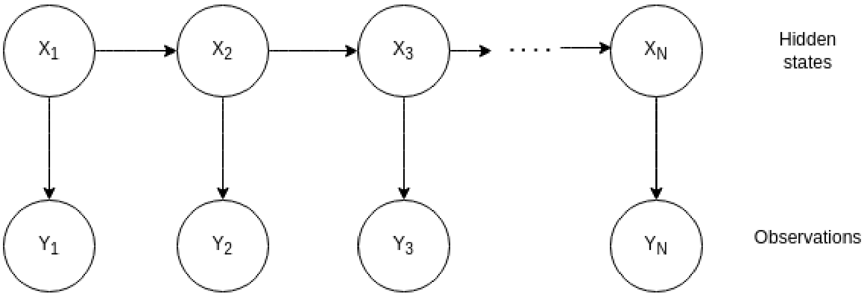

FIGURE 2.4: A directed Hidden Markov model.

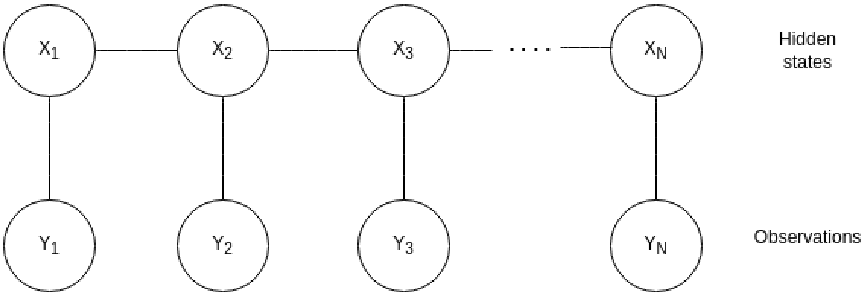

FIGURE 2.5: An undirected Hidden Markov model.

Intuitively, HMMs calculate the joint probability of the set of hidden states, given a set of observations. Once we infer this joint probability distribution, we select the sequence of hidden states that maximizes this probability. We mention that various HMM versions have been introduced in the literature, that alter and expand the fundamental concept to suit the requirements of different applications [85].

2.4.2 Bayes Filter

The Bayes Filter is a well-known technique that is employed for recursive state estimation. Generally, it finds application in many different scenarios such as robotics and autonomous driving. In more detail, given a set of observations (or measurements), $\mathbf{y}$, and a set of control commands, $\mathbf{z}$, its goal is to generate a belief, denoted as bel, regarding the current state of the system. Bayes Filter computes the $bel(\mathbf{x}_t)$ at timestamp t by combining its belief at the previous timestamp, i.e., $bel(\mathbf{x}_{t-1})$, along with the most recent control command $\mathbf{z}_t$ and the most recent measurement $\mathbf{y}_t$ [81]. Note, that commonly Bayes Filter assumes that the process can be represented as a Markov chain. Algorithm 1 illustrates the general algorithm of the Bayes Filter. In particular, this algorithm consists of two steps, namely the prediction step and the update step. In the prediction step, the algorithm calculates a belief over the current state, $\mathbf{x}_t$, based on the prior belief over the previous timestamp, $\mathbf{x}_{t-1}$, and the control command $\mathbf{z}_t$. Subsequently, in the update step, the algorithm multiplies the belief obtained in the prediction step, $\bar{bel}(\mathbf{x}_t)$, by the probability that the measurement $\mathbf{y}_t$ may have been observed and a normalization constant, η [81]. We note that a special-case of the Bayes Filter algorithm is the well-known Kalman Filter algorithm which makes the extra assumption that the system is linear Gaussian (see Section 2.4.3).

Algorithm 1: Bayes Filter

1 **For all $\mathbf{x}_t$ do:**
2 $\quad \bar{bel}(\mathbf{x}_t) = \int p(\mathbf{x}_t|\mathbf{z}_t, \mathbf{x}_{t-1}) \cdot bel(\mathbf{x}_{t-1}) dx$
3 $\quad bel(\mathbf{x}_t) = \eta \cdot p(\mathbf{y}_t|x_t) \cdot \bar{bel}(\mathbf{x}_t)$
4 **end**
5 return $bel(\mathbf{x}_t)$

2.4.3 Kalman Filter

Kalman Filter [37] is a well-known estimator that allows us to model any *linear system*[1], that exploits a series of observations that contain noise in order to produce noiseless estimates regarding the value of unknown variables. Additionally, we note that Kalman Filter is a special case of the Bayes filter where the dynamics and observation model is linear Gaussian. Generally, Kalman Filter was developed in the early 1960s for the Apollo program and have been employed in various real-world applications, e.g., robot localization, aircraft tracking etc. Intuitively, this algorithm produces estimates from inaccurate data, i.e., measurements or observations, by *filtering out* the existing noise.

More in detail, the states evolve based on linear dynamics:

$$x_{t+1} = Ax_t + v_t \tag{2.13}$$

where A is a constant matrix and $v_t \sim \mathcal{N}(0, Q)$ is the zero-mean Gaussian process noise. Additionally, Kalman Filter assumes that the observations depend on the state according to the linear measurements:

$$y_t = Cx_t + w_t \tag{2.14}$$

where C is a constant matrix and $w_t \sim \mathcal{N}(0, R)$ is the zero-mean Gaussian observation noise.

Now, Kalman Filter algorithm consists of two sub-steps, namely the *Prediction* and the *Update* sub-step. The goal of this algorithm is to produce the mean vector $\mu_{i|i}$ and the corresponding covariance matrix $\Sigma_{i|i}$ for each state.

Prediction Sub-Step The prediction sub-step for the corresponding marginals can be derived as follows:

$$\mu_{i+1|i} = A\mu_{i|i} \quad i = 0, 1, \cdots, t-1 \tag{2.15}$$

$$\Sigma_{i+1|i} = A\Sigma_{i|i}A^T + Q \quad i = 0, 1, \cdots, t-1 \tag{2.16}$$

[1] We note that for non-linear systems an extension of Kalman Filter has to be utilized, which is called Extended Kalman Filter [36].

Update Sub-Step Here we update the prediction by exploiting the new data, denoted as y_{i+1}. Specifically:

$$\mu_{i+1|i+1} = \mu_{i+1|i} + G_{i+1}(y_{i+1} - C\mu_{i+1|i}) \quad i = 0, 1, \cdots, t-1 \tag{2.17}$$

$$\Sigma_{i+1|i+1} = \Sigma_{i+1|i} - G_{i+1}C\Sigma_{i+1|i} \quad i = 0, 1, \cdots, t-1 \tag{2.18}$$

where G_{i+1} is the *Kalman gain* and can be computed as:

$$G_{i+1} = \Sigma_{i+1|i}C^T(C\Sigma_{i+1|i}C^T + R)^{-1} \tag{2.19}$$

The Kalman filter algorithm for linear Gaussian state transitions and measurements is illustrated in Algorithm 2.

Algorithm 2: Kalman Filter

Prediction Sub-Step:
  $\mu_{i+1|i} = A\mu_{i|i}$
  $\Sigma_{i+1|i} = A\Sigma_{i|i}A^T + Q$
Update Sub-Step:
  $G_{i+1} = \Sigma_{i+1|i}C^T(C\Sigma_{i+1|i}C^T + R)^{-1}$
  $\mu_{i+1|i+1} = \mu_{i+1|i} + G_{i+1}(y_{i+1} - C\mu_{i+1|i})$
  $\Sigma_{i+1|i+1} = \Sigma_{i+1|i} - G_{i+1}C\Sigma_{i+1|i}$
return $\mu_{i+1|i+1}, \Sigma_{i+1|i+1}$

2.5 Social Choice Theory

Social choice theory is a theoretical framework that applies to various domains such as economics, political science, computer science, etc. In general, social choice theory studies aggregation mechanisms of individual preferences in order to reach a collective choice or decision [11]. Over the years, particular emphasis has been given to the scenario of electing a single "winner" over a set of items or alternatives. In more detail, given a set of users or *voters* and a set of items or *alternatives*, social choice theory studies efficient mechanisms and rules in order to elect the best alternative with respect to voters' preferences. Formally, a *social choice function* utilize ranked *ballots* of users where [11]:

- U is a finite set of voters.
- I is a finite set of alternatives, with $|I| \geq 2$.
- Denote the set of linear orders on I by $\mathcal{L}(I)$.
- The ballot cast by voter u is a linear ordering $\succsim_u$ of I.
- A profile $P = \{\succsim_1, \succsim_2, \cdots\} \in \mathcal{L}(I)^U$ specifies such a ballot for each voter $u \in U$.
- A social choice function (SCF) or *voting rule* is a function $F : \mathcal{L}(I)^U \rightarrow 2^I \setminus \emptyset$, mapping any given profile to a non-empty set of winners.

Many single-winner voting rules have been proposed in the literature with most known *plurality, Borda count* and *Copeland* voting rules [11]:

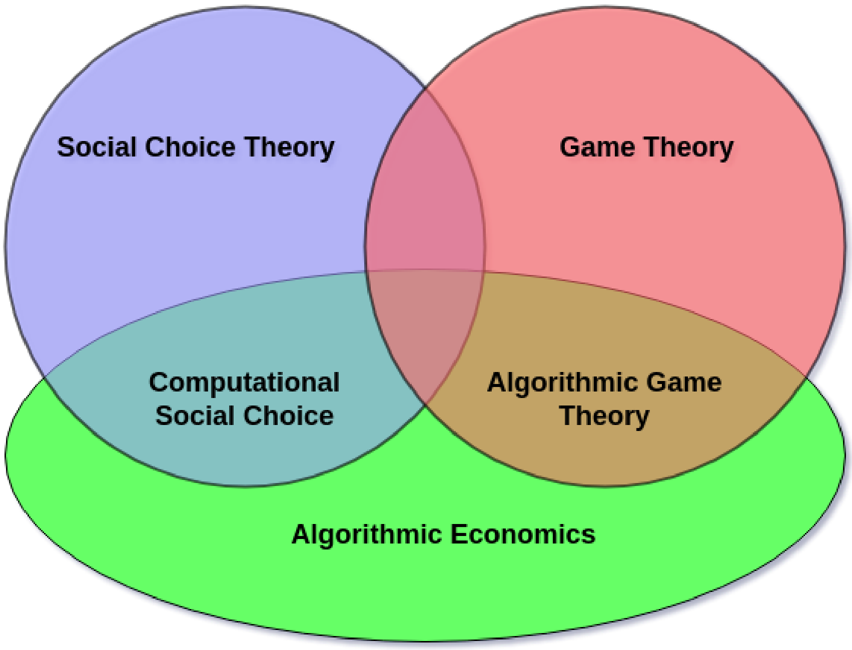

FIGURE 2.6: Relation of social choice, game theory and economics (figure constructed based on the scheme in the following link: https://www.algo.tu-berlin.de/menue/research/).

- Plurality: elects the candidate ranked first most often, i.e., each voter assigns one point to a candidate of her choice, and the candidate receiving the most votes wins.
- Borda: each voter gives $|I| - 1$ points to the candidate she ranks first, $|I| - 2$ to the candidate she ranks second, etc., where $|I|$ is the total number of alternatives. The candidate with the most points is elected.
- Copeland: award 1 point to a candidate for each pairwise majority contest won and $\frac{1}{2}$ for each draw. The candidate with the most points is elected.

2.5.1 Multiwinner Elections

However, there is another type of elections that has recently gained the interest of researchers. Specifically, in this type of elections, the purpose is to select a k-sized group of alternatives, i.e., a *committee* of size k, rather than just a single winner, i.e., a single alternative. This type of elections are also known as *Multiwinner elections* [27]; and can be categorized as: *(i)* Shortlisting, *(ii)* Diverse Committee, and *(iii)* Proportional Representation mechanisms [24], based on their type and their properties. Intuitively, a Shortlisting mechanism elects a committee consisting of the alternatives that have the best quality with respect to some feature(s), e.g., on a job interview scenario the application of a shortlisting mechanism would result to a committee consisting of candidates that have similar skills and characteristics [27]. By contrast, a Diverse Committee mechanism elects a committee consisting of alternatives that are diverse based on some feature(s). For instance, consider a travel agency that recommends travel destinations to a customer. The employment of such mechanisms in this case could produce a set of k travel destination that differ with each other based on their location (or any other feature). As such, a capital city in South America, an exotic island in the Pacific ocean and a traditional village in Asia could be elected as the recommendations of the agency. Finally, a Proportional Representation mechanism selects a committee that captures all the different preferences of the voters proportionally.

An important class of voting rules are the so-called *approval-based* rules, in which voters indicate the alternatives they "approve". Perhaps the most well-known such mechanism is the *Approval Voting (AV)*. Given a list of candidates, Approval Voting allows each voter to express her approval, i.e., her support, for many candidates. Then the mechanism elects the candidate who earned the greatest number of approval votes. In general, AV satisfies several desirable properties in the situation of a single winner, but it fails to attain proportionate representation in the case of multiwinner elections [3]. On the other hand, the *Proportional Approval Voting (PAV)* multiwinner election method is an approval-based rule that meets strong theoretical guarantees for election proportionality. In more detail, according to the PAV rule, the weight of each voter to the committee's final score is based on how many candidates from the voter's approval set were elected [3]. However, Skowron et al. [75] showed that winners determination under PAV is an NP-hard problem. To tackle this, a "sequential" PAV variant, namely the *Reweighted Approval Voting (RAV)*, which is essentially a "greedy" approximation of the PAV rule, has been introduced in the literature [27]:

Definition 1 (Reweighted Approval Voting - RAV [27]) *Consider an election with n voters where the i-th voter approves candidates in the set A_i. RAV starts with an empty committee S and executes k rounds. In each round it adds to S a candidate c with the maximal value of:*

$$\sum_{i:c\in A_i} \frac{1}{|S\cap A_i|+1} \tag{2.20}$$

That is, according to this definition, at each iteration RAV re-adjusts the weights of each voter's ballot in order to achieve a proportional representation in the final committee. Note that under RAV the determination of the set of winners can be computed in polynomial-time [5].

Another approval-based rule is the *Bloc* rule; according to Bloc, each voter selects her k favourite alternatives and the mechanism elects a committee consisting of the alternatives that were mentioned more frequently [25]. Finally, a well-known non-approval-based rule, the *k-Borda* selects the k alternatives with the highest Borda score [25, 11].

To the best of our knowledge, such mechanisms have mostly been researched as theoretical tools, as they have not been implemented in real-world recommender systems—with the notable exception of Gawron and Faliszewski [30]: in that paper, the authors introduced a system that exploits multiwinner election mechanisms in order to produce a set of resources (or items) that are similar to a given query. In more detail, by employing different mechanisms the system is able to control the degree of relation between the recommended resources and the given query. However, their proposed system operates more like a search engine rather than a classic recommender system, since it elects items that are somewhat related to the given query, instead of recommending items that can increase the satisfaction of a specific user (i.e., they do not take into account user-specific features in order to provide personalized recommendations).

Chapter 3

Our Approach

In this chapter, we introduce two recommendation systems that enhance the visitors' experience of a tourist destination by recommending POIs that "match" their preferences. In the first approach, inspired by the work of Babas et al. [6], we designed a Bayesian recommender system which performs *Bayesian learning* (or Bayesian updating) in order to learn user's (which corresponds to tourist's) preferences and provide her with efficient recommendations. Regarding the second recommender system we employ the *Kalman Filter* algorithm in order to "track" the position of the model that represents user's preferences in a multi-dimensional space of travel-related features in order to produce recommendations of items, i.e., POIs, that are "located" near to her interests in this multi-dimensional space. Specifically, we model the users and the items (i.e., the POIs of a tourist destination) of our setting, using a common representation—i.e., multivariate normal distributions over ranges of values, describing the degree that each feature describes a specific user or a specific item. Such density-based representation provides many advantages, e.g., better encoding of uncertainty for a representation [84]. We note that, our system is unaware of every user's *real distribution*, while it is fully informed about each item's distribution. The goal of our recommender systems, is to "predict" the distribution, i.e., to generate a belief, describing the interests of a specific user in order to recommend the POIs most appropriate for her, thus increasing her satisfaction.

To this purpose, our systems perform a series of "questions" in order to determine user's preferences by exploiting the data derived by this procedure. Specifically, we employ a novel picture-based user preference elicitation process, by presenting to the user a set of generic travel-related pictures and noting her "likes" in order to build a user model. Moreover, we equip our systems with some prior knowledge regarding the general preferences of a specific type of visitors, i.e., tourists that belong to the same age group.

On top of that, we design a social choice inspired mechanism that produces diverse personalized recommendations to the user with respect to the selected travel-related features. Specifically, given a specific user we create a "personalized election" based on her inferred model and employ various multiwinner election rules in order to produce diverse personalized recommendations.

Finally, we tackle the group recommendation problem in the tourism domain. To this end, we use "tools" derived from the social choice theory to effectively aggregate tourist group members' preferences in a fair manner. Specifically, we proceed to show how to tackle the group recommendation problem via employing various preference aggregation mechanisms—and, importantly, for the first time in tourism-oriented recommender systems, a multiwinner voting rule. Figure 3.1 depicts the architecture of our approach for the problems of personalized and group recommendations.

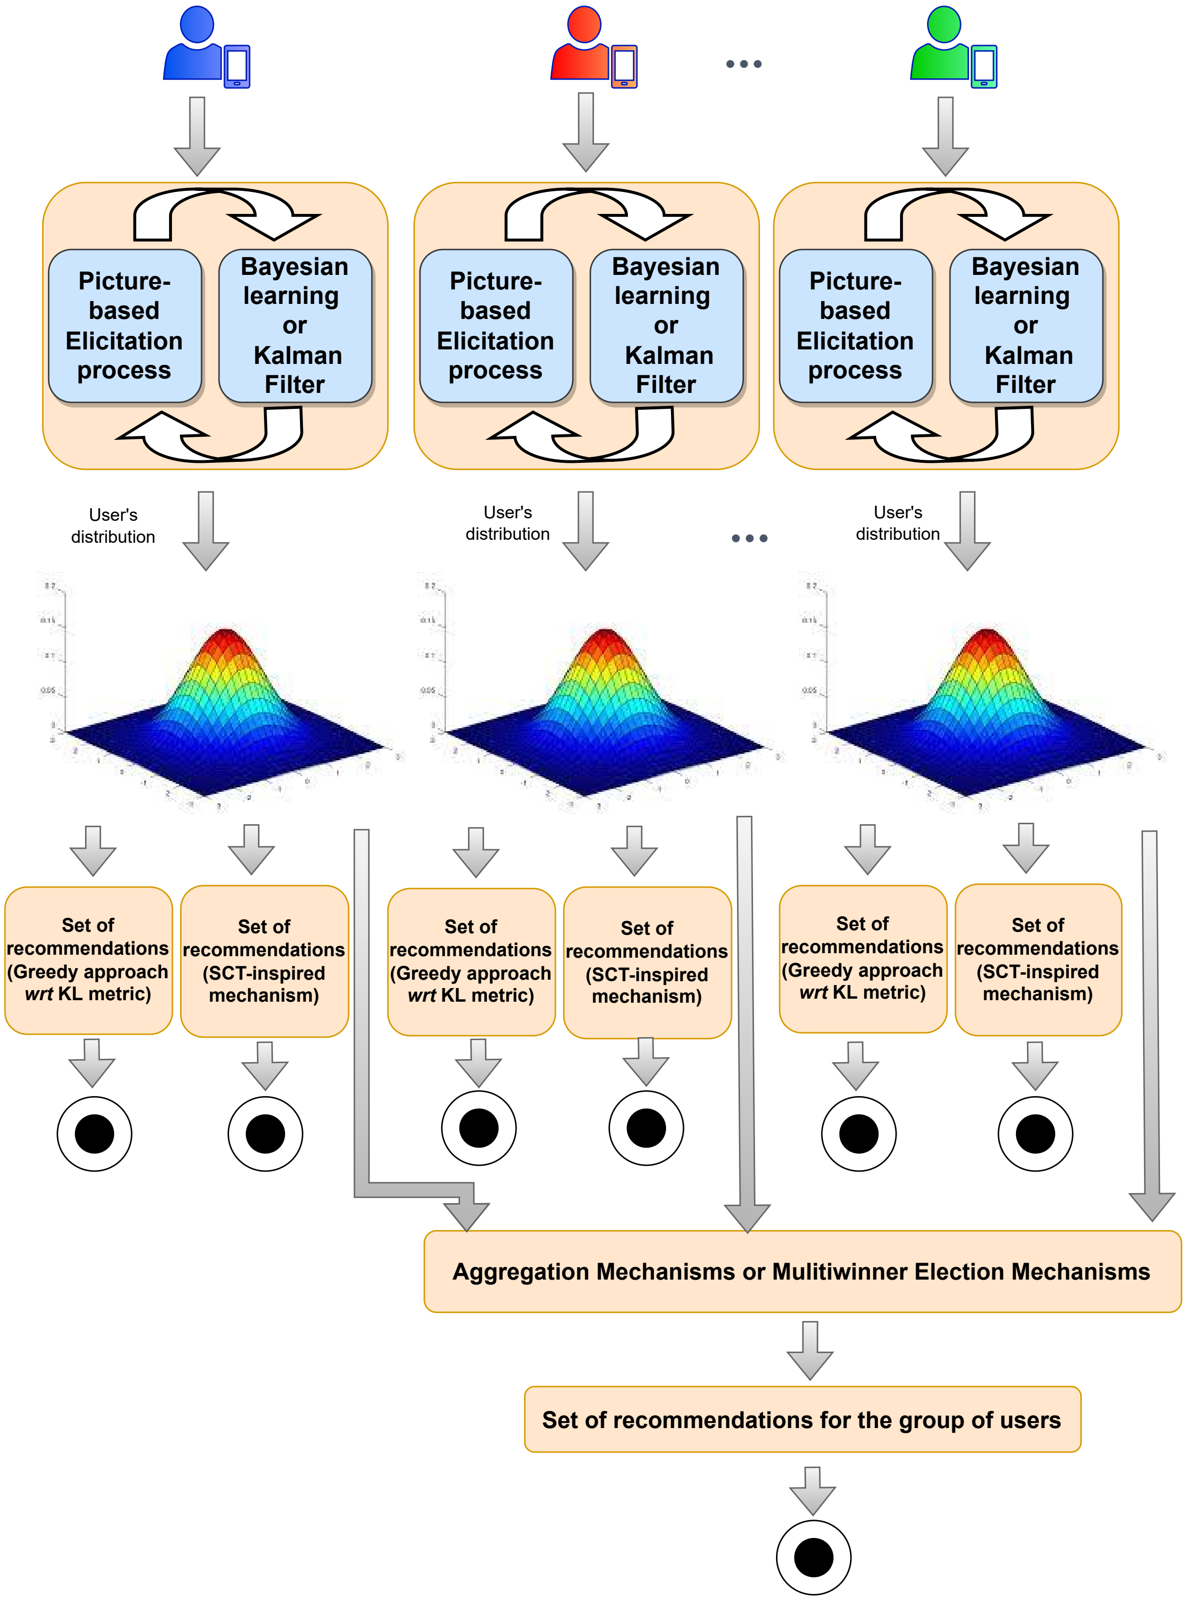

FIGURE 3.1: The architecture of our approach.

3.1 Elicitation Process

First, when a user registers in the system the *preference elicitation* procedure begins. In more detail, the preference elicitation procedure is an iterative process where in each iteration our system presents n alternative *generic images* to the user. We use the term "question" to refer to such a presentation of a set of generic images. These generic images, similar to POIs, are represented as multivariate Gaussians. Each generic image corresponds a specific *type* of POIs, i.e., a restaurant, a monument, a beach etc. We select this picture-based approach to elicit user's preferences, since there is a great complexity regarding the tourism product [21], and such approaches have been efficiently applied in order to elicit travel-related preferences [55, 56]. We highlight that the set of the generic images and the set of POIs do not share any common element (i.e., there are no items that belong to both lists).

Note that our agent selects which generic images to present to the user based on the available information (represented by the multivariate distribution), regarding her preferences (see Section 3.4). However, when a new user enters the system, most of the times we have no available information about her, so we have to *randomly* pick some generic images in the first iteration, unless our system has some prior knowledge at its disposal (e.g., a prior knowledge regarding the preferences of a specific age group, as we discuss later in Section 4.2.2). In such a case, our recommender is able to exploit the extra knowledge in order to pick the images that will present to the user in the first iteration. Once our algorithm selects the generic images that will present to the user, the user picks the image that is most "attractive" to her, with respect to her interests, and provides a rating on a 5-level Likert scale, where 5 implies that this image fits her preferences perfectly.

In our work, similarly to that of [6], we use the *Kullback-Leibler (KL) Divergence* criterion in order to compute the similarity between any item (generic image or POI) and any user based on her preferences and interests. In particular, knowing that both users and items share a common representation—since both are modelled as *multivariate normal distributions*— we can employ the KL-divergence criterion in order to find "how similar" their distributions are. Formally, the KL-divergence between a Gaussian x and a Gaussian y, of dimension D each, is given by:

$$KL(y||x) = \frac{1}{2}log|\Sigma_y^{-1}\Sigma_x| + \frac{1}{2}tr((\Sigma_y^{-1}\Sigma_x)^{-1}) - \frac{D}{2} + \frac{1}{2}(\mu_y - \mu_x)^T\Sigma_x^{-1}(\mu_y - \mu_x) \quad (3.1)$$

where $\Sigma_y, \mu_y, \Sigma_x$ and μ_x are the distributions' parameters, and $tr(\cdot)$ is the trace of the corresponding matrix [57]. In principle, a small KL-divergence between a Gaussian x and a Gaussian y means that they are similar, while a large KL-divergence means the distributions are not similar. Thus, in our work we make the natural assumption that the more similar the distributions of a user u and an item i are, the higher the rating (of user u for item i) would be. As such, the (predicted) rating of a user (represented as a Gaussian u) for an item (represented as a Gaussian i) can be calculated based on the following formula:

$$r_{u,i} = M - \frac{KL(u||i)}{M} \quad (3.2)$$

where M is the maximum rating the user can give to an item, i.e., $M = 5$. We highlight that up to this point *both* of our proposed approaches, i.e., the Bayesian learning and the Kalman Filter, operate identically. However each approach utilizes the provided ratings of a specific user differently. In what follows, we describe in detail how our systems operate in order to produce the final personalized recommendations for a targeted-user.

3.2 Bayesian Learning Approach

In this approach, our system draws an appropriate number of samples[1] from the distribution of the selected generic image based on the provided rating. Specifically, we use the *logistic function* [17] and the rating of the user in order to compute the exact number of samples that will be drawn from the generic image's distribution. Intuitively, the form of this function fits to our purposes, since a high rating (signifying the user likes a generic image), means that the distributions of the user and the image are similar, and as such a sufficiently large number of samples from the image's distribution will contribute to construct a good model regarding user's preferences; while a small rating means that the user does not feel that this image describes her preferences well. Thus, a small number of samples should be drawn since they are not representatives of user's interests.

Once the user enters her rating, our approach performs a *Bayesian updating* (based on the equations presented in Section 2.3) in order to produce an updated type of user (i.e., distribution) combining prior knowledge and the new data of this iteration, which corresponds to a user-system interaction. Note that the posterior derived in iteration t will be the prior for iteration $t + 1$, with which our system will choose which generic images to present to the user next (see Section 3.4). That is, the generic images at each iteration are chosen based on the beliefs updated on the previous iteration. This procedure terminates after m iterations. As such, our system estimates the parameters of user's (μ, Σ) based on the hyperparameters of the NIW distribution (see Section 2.3). Figure 3.2 depicts the overall process of elicitation and Bayesian learning. Finally, our (greedy) version of our system utilizes the estimated parameters of the user in order to produce its final recommendations, by applying the KL divergence—i.e., it greedily recommends the k POIs that are more similar to the tourist's inferred model. We note that, more user-system interactions, i.e., more m iterations, can provide our system better indications regarding the feature preferences of a specific user resulting a more representative user model.

[1] Note that depending on the rating the number of samples varies from 40 to 500.

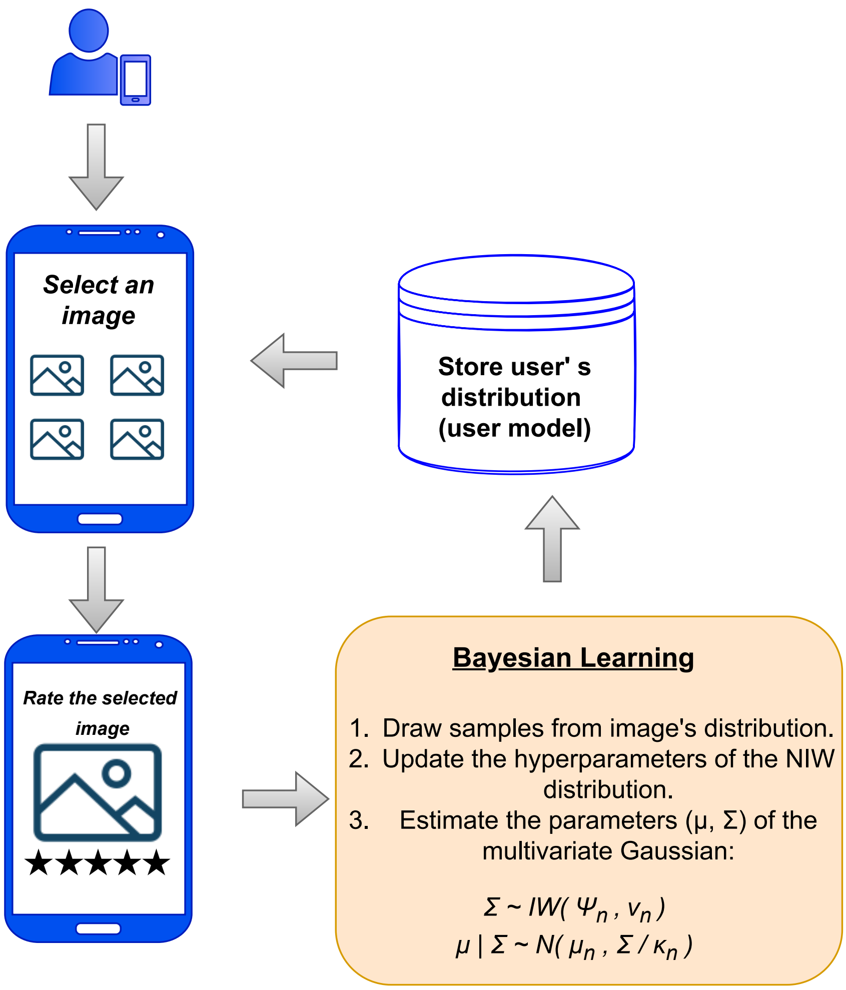

FIGURE 3.2: The preference elicitation process and the Bayesian learning approach.

3.3 Kalman Filter Approach

In this approach, user's provided ratings correspond to the observations that our system is able to exploit in order to update its beliefs regarding the "position" of user's model in the multi-dimensional space. For example, these ratings can be thought as measurements of a sensor in a robot that employs the Kalman Filter algorithm in order to determine its position in a room. Specifically, let us assume that a user u picked a generic image i and provided a rating $r_{u,i}$. However, in the Kalman Filter algorithm we have to use a belief, i.e., a Gaussian distribution, that represent the measurement and the corresponding uncertainty. As such, we have to exploit a Gaussian distribution in order to execute the Update sub-step (see Algorithm 2). One could consider the naive solution of exploiting i's (pure) multivariate Gaussian in order to perform the update. However, in such case, the recommender system is unable to exploit the user's provided rating, $r_{u,i}$. Thus, we chose to *generate* a distribution at each iteration that contains information regarding *(i)* item i features, and *(ii)* user rating $r_{u,i}$.

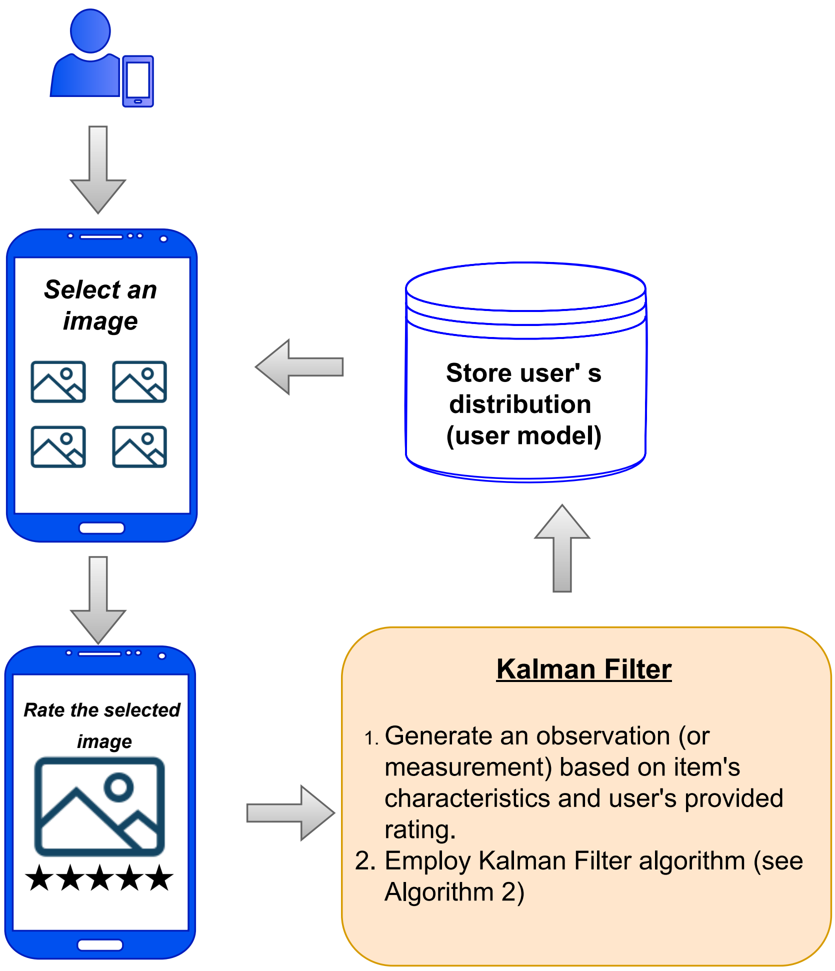

FIGURE 3.3: The preference elicitation process and the approach of Kalman Filter algorithm.

In particular, we create a (temporary) multivariate Gaussian distribution with mean vector μ_{tmp} and covariance matrix Σ_{tmp} that can be computed as follows:

- $y_i = \mu_{tmp} = \mu_i$
- $R = \Sigma_{tmp} = \begin{bmatrix} M - r_{u,i} & 0 & \cdots & 0 \\ 0 & M - r_{u,i} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & M - r_{u,i} \end{bmatrix}$

where μ_i is the mean vector of the selected generic image, M is the maximum rating the user can give to an item and $r_{u,i}$ is the rating of user u for item i. Intuitively, when a user enters a high rating it means that item's model is located near to user's model and a small uncertainty is added in the measurement, while a low score means that there is large distance between user and item in the multi-dimensional space, so we insert extra noise (or uncertainty) in the measurement. We note that the posterior derived in iteration t will be the prior for iteration $t + 1$, with which our system will choose which generic images to present to the user next (see Section 3.4). That is, the generic images at each iteration are chosen based on the beliefs updated on the previous iteration. This procedure terminates after m iterations. As such, our system estimates the parameters of user's (μ, Σ) based on the Kalman Filter algorithm (see

Algorithm 2). Figure 3.3 depicts the entire process of elicitation and Kalman Filter algorithm. Finally, similar to Bayesian learning approach, the (greedy) version of our system utilizes the estimated parameters of the user in order to produce its final recommendations, by applying the KL divergence—i.e., it greedily recommends the k POIs that are more similar to the tourist's inferred model.

3.4 Generic Images Selection

One crucial aspect of our recommendation systems is the selection of an effective method which decides the items, i.e., the generic images, that will be presented to the user at each iteration. The item selection method should handle in an efficient way the *exploration vs exploitation* dilemma in order to increase the performance of our system. Intuitively, we could summarise this dilemma as "Should I go for the decision that seems to be optimal, assuming that my current knowledge is reliable enough? Or should I go for a decision that seems to be sub-optimal for now, making the assumption that my knowledge could be inaccurate and that gathering new information could help me to improve it?".[2] In our work we studied three alternative item selection methods, namely: a greedy mechanism based on the KL-divergence, the VPI exploration [6], and the Boltzmann selection. The Boltzmann selection mechanism performed better and as such is the method we detail here. Notably, in the experimental evaluation the results corresponds only to this mechanism.

Intuitively, Boltzmann exploration [13] tells us to pick an action with a probability that is proportional to its average reward. As such, actions with greater average rewards are picked with higher probability. Formally, at each time step t, our agent assigns a selection probability to each item i using the following formula:

$$Pr(i) = \frac{e^{r_{u,i}/T}}{\sum_{j=1}^{n} e^{r_{u,i}/T}} \tag{3.3}$$

where $T = c \cdot a^t$, with c be a constant value, $r_{u,i}$ is the quantity computed in Equation 3.2, and $a < 1$. In this method, the exploration vs exploitation trade-off is controlled via a temperature parameter T that decreases over time to progressively reduce exploration. Specifically, for very small values of T the action with the highest average reward (in our case, the highest predicted $r_{u,i}$ rating) is more likely to be selected. On the other hand, in initial stages where the value of T is large, the Boltzmann method effectively corresponds to a random policy. As such, in our work, similar to [6], we choose to decrease the value of T over time, i.e., the exploration is progressively reduced at each iteration.

3.5 Personalized Recommendations using Multiwinner Election Mechanisms

As described earlier, our system performs either a Bayesian inference procedure or the Kalman Filter algorithm in order to learn the interests of a targeted-user, i.e., to create a distribution that describes the user with respect to some travel-related features. Once the user's model has been constructed, both our recommenders greedily recommend the k items that are more similar to the tourist's inferred distribution

[2]Phrase taken from:
https://towardsdatascience.com/the-exploration-exploitation-dilemma-f5622fbe1e82

based on the KL divergence criterion. However, such an approach can lead to several undesirable properties regarding the final recommendations that the user will receive. For example, let us assume that a tourist during the preference elicitation process chose to give a high rating only to restaurants. In this case, a "greedy" approach will produce recommendations which will only be related to restaurants. Such a property might be desirable in other domains such as movies (e.g., action movies, etc.) but in the tourism domain, a tourist would like to have a variety of different POIs that she could visit in a travel destination, e.g. a restaurant, a monument, a bar, and so on. In other words, in order for the recommender algorithm to be able to provide a multitude of choices to the tourist, it should guarantee that there will be some diversity between its final k recommendations. Moreover, we remind the reader that the computation of the recommendations' quality is based on the user model constructed so far. However, such a model is constructed using only limited user-system interactions resulting to a rather naive estimation of the actual user model; thus, we conjecture that providing diversity in the final recommendations is helpful. Thus, in our approach we take inspiration from social choice theory and in particular from multiwinner election mechanisms in order to provide a novel method aiming to ensure diverse results with respect to any subset of features. Note that this method is refferd to as a social choice theory one ("SCT-inspired mechanism" in Figure 3.1) and is depicted graphically in Figure 3.4. We now proceed to provide the details of this method.

In general, given a set of users (or voters), N, and their preferences, multiwinner election mechanisms select a k-sized set of alternatives (i.e., "a committee"). Such mechanisms satisfy specific properties based on their type (see Section 2.5). Given this, we now detail our approach to provide diverse recommendations to a tourist with respect to some travel-related features, instead of greedily selecting the POIs that are more similar to user's inferred model with respect to the KL metric.

To this end, given a specific tourist that is described by a multivariate Gaussian distribution over a set of features, we create a "personalized election" in order to produce our final recommendations. In more detail, given a user u, we exploit her (inferred) mean vector, μ_u, in order to create a set of voters based on u's values over the selected travel-related features, i.e., the values on the μ_u for each (selected) feature. As such, we generate a set of voters, V, that provides proportional representation of user's preferences over the features, i.e., a feature that has a high score will be represented by more voters than a feature that has a low value for u. Specifically, for any travel-related feature f, with value f_v, we generate $\lceil f_v \cdot 10 \rceil$ voters. Moreover, we make the natural assumptions that any voter that has been generated from feature f: *(i)* approves an item i, i.e., a POI, that has a value that is greater than 3 on feature f, i.e., $f_i \geq 3$; and *(ii)* prefers an item i over an item j if and only if i has a greater value than j on feature f, i.e., $f_{i,v} > f_{j,v}$,[3]. On the other hand the set of alternatives, A, and the set of POIs are identical. Therefore, we can apply *any*[4] multiwinner rule on this election in order to be able to pick *several* "personalized" winners that also satisfy specific properties (e.g., proportionality of the preferences' representation) guaranteed by the rule in question, based on the travel-related features and the preferences of the tourist. As such, by not simply "greedily" recommending the so-far-perceived-as-best POIs, the proposed approach provides diverse

[3]We note that if $f_{i,v} = f_{j,v}$ we randomly select which item is preferred ($i \succ j$ or $j \succ i$) for each voter independently.

[4]In our experiments below, we test several multiwinner election rules—specifically, AV, RAV, k-Borda and Bloc (see Section 2.5).

recommendations. This fact is evident in our experimental results, where we see that using multiwinner elections leads to improved system performance when the user-system interactions are limited—with this advantage decreasing or evaporating with increased user-system interactions.

We also point out that such an approach allows us to provide diverse recommendations with respect to *any selected subset* of travel-related features. In more detail, our approach can maintain such property by creating a "personalized election", where the voters have been produced *solely* from the selected subset of features following the procedure that has been already described in this section. (Of course, the set of the alternatives remains the same—i.e., it comprises all the available POIs.)

3.6 Recommendation Process for Groups

In this section we describe how to apply preference aggregation methods and mechanisms on top of any single user recommendation technique. In more detail, we consider a set of items (or alternatives) and a set of users (or agents), denoted as $\mathcal{I}$ and $\mathcal{U}$ correspondingly. We can employ *any* recommender system technique of choice (e.g., collaborative filtering, content-based, Bayesian, or any other), in order to generate a predicted score, $r_{u,i}$, for any possible combination of i and u, where $i \in \mathcal{I}$ and $u \in \mathcal{U}$—i.e., the (single-user) recommender system predicts that user u will rate item i with a score of $r_{u,i}$. Moreover, for each individual u, we can make the natural assumption that u prefers item i over j if the predicted score of i is larger than the one of j, i.e., $i \succ_u j$ if and only if $r_{u,i} > r_{u,j}$. Thus, our system is able to create the preference list for any individual u.

Now, assume that a group, denoted as g, consists of $|g|$ members corresponding to tourists, i.e., individual users. Our system is able to exploit the aforementioned preference list of every member of the group, derived after the employment of the single-user recommendation technique of choice during the (independent) interaction of each user with the system. Then, our group recommender system is able to produce group recommendations to any group g, by exploiting the preference lists of the group members and by applying *any* aggregation mechanism of choice.

In the experimental evaluation of our system, we explore several well-known aggregation strategies, such as the *Least Misery (LM)* strategy, *Most Pleasure (MP)* strategy and the *Additive Utilitarian (AU)* strategy. Formally, the LM mechanism provides recommendations (or items) that maximize the minimum individual rating among the members of the group, i.e., for each item $i \in \mathcal{I}$, LM mechanism assigns a score equal to $min_{u \in g,i}\{r_{u,i}\}$ and recommends the item(s) with the highest score(s). Similarly, the MP strategy elects the items that maximize the maximum individual rating among the members of the group, i.e., for each item $i \in \mathcal{I}$, MP assigns a score equal to $max_{u \in g,i}\{r_{u,i}\}$ and recommends the item(s) with the highest score(s); while the AU mechanism recommends the items that maximize the average individuals ratings among the members, , i.e., for each item $i \in \mathcal{I}$, AU assigns a score equal to $avg_{u \in g,i}\{r_{u,i}\}$ and recommends the item(s) with the highest score(s). Note that the LM, MP and AU aggregation strategies have been in the past employed for tackling the problem of group recommendations in the tourism domain [19].

In addition, we use the RAV multiwinner election mechanism (see Section 2.5)—for the first time in the recommender systems' literature. We remind the reader that RAV is a greedy multiwinner voting rule that elects a committee (i.e., a set of items),

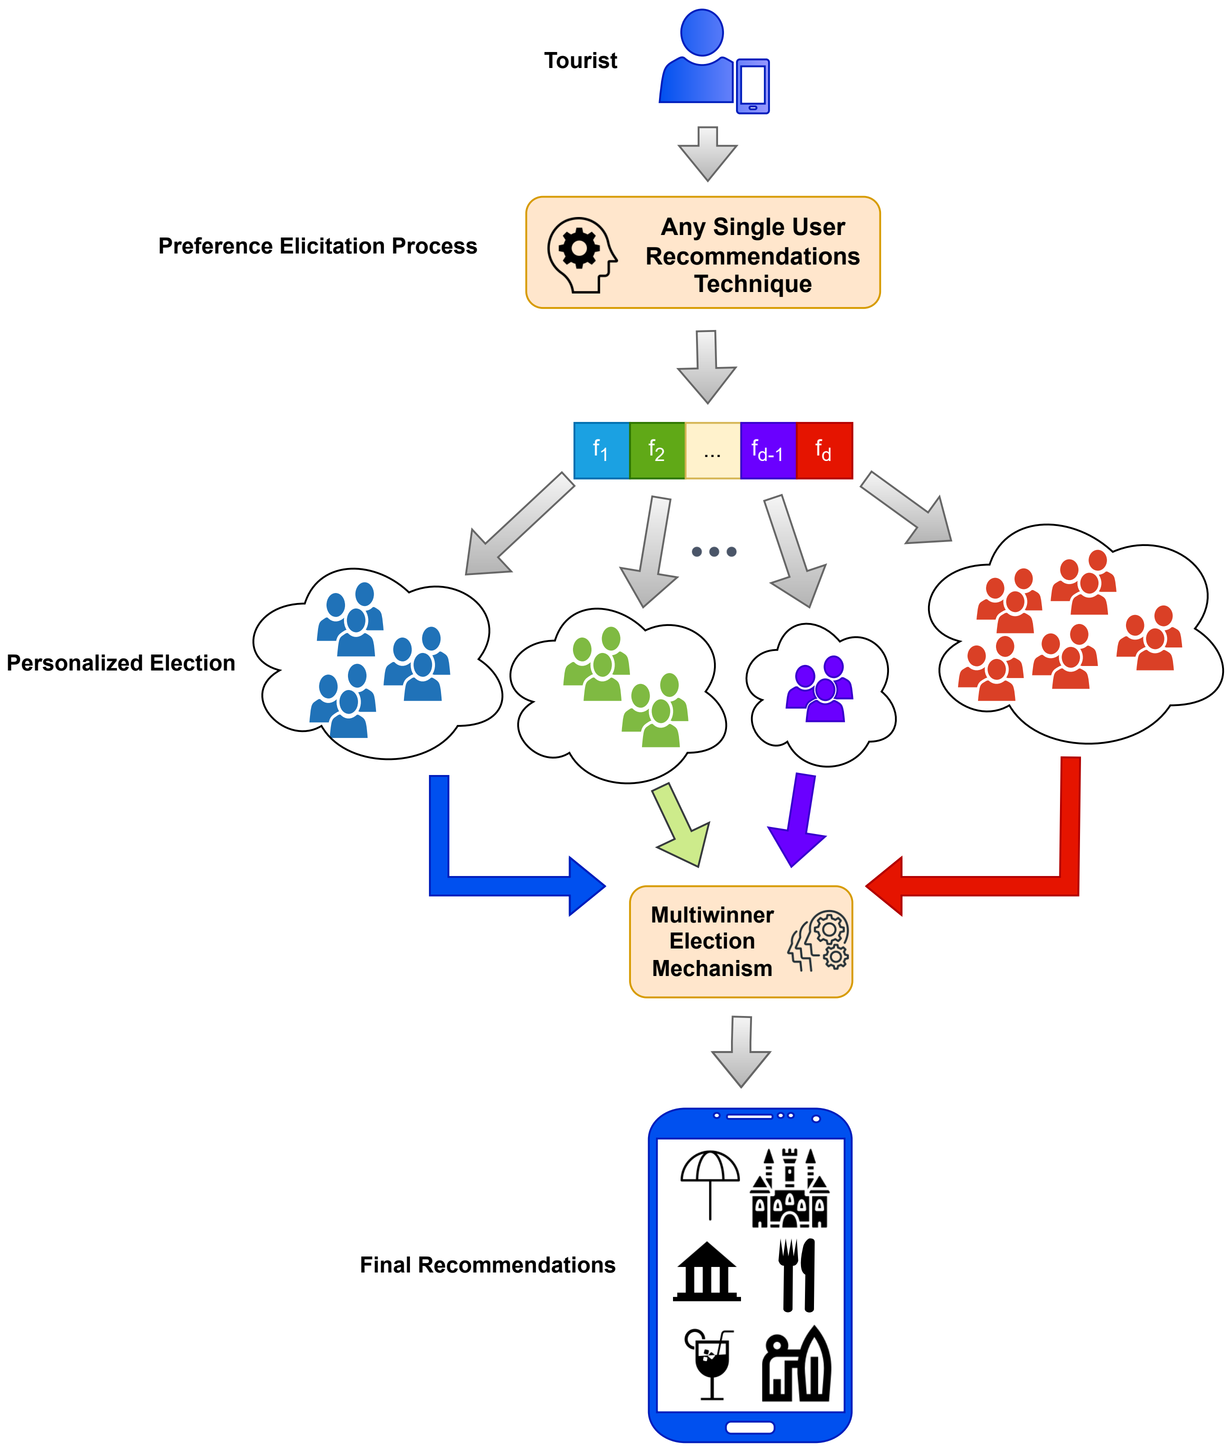

FIGURE 3.4: The social choice inspired mechanism for generating the final personalized recommendations.

that represent proportionally the preferences of the voters, i.e., the members of the group. Generally, in the problem of group recommendations, such mechanisms can be very useful due to their properties, since the proportionality that is achieved by the elected committee provides a notion of fairness among the members of the group

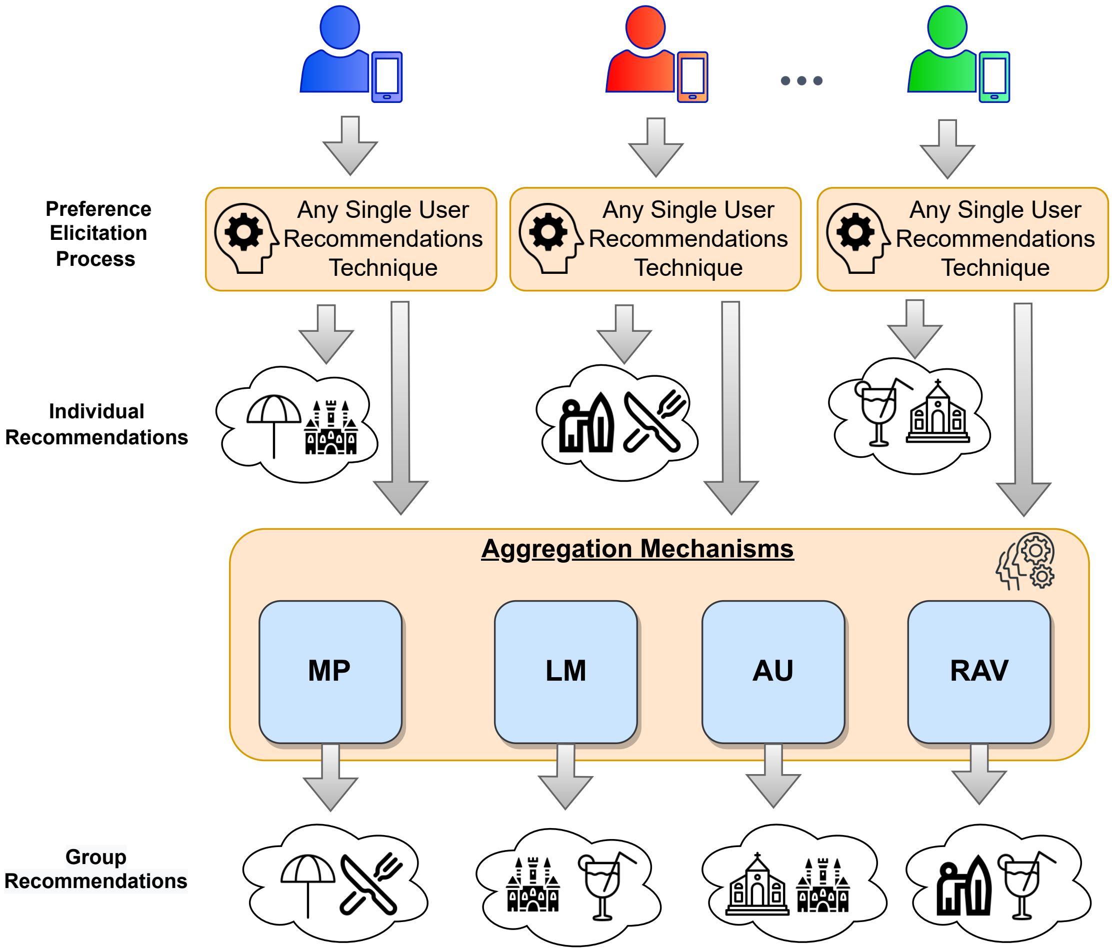

FIGURE 3.5: The group recommendation process.

with respect to their preferences.[5] We note that, the RAV mechanism was selected since our purpose is to find a computationally efficient mechanism that is able to satisfy fairness in a real world application, where the group size can be large and the recommended items have to be displayed quickly in order to ensure that visitors would not waste their time waiting for the results, i.e., the final group recommendations. Moreover, RAV mechanism can be considered as a good greedy approximation algorithm for the PAV rule [27], which is the only w-AV rule which satisfies the property of *Extended Justified Representation (EJR)* [27, 3]. Note that all aforementioned aggregation mechanisms can be employed for any possible group size.

To the best of our knowledge, this work applies for the first time multiwinner elections in a real-world recommender system. As such, in order to evaluate our mechanism we chose to focus on the metrics of m-*PROPORTIONALITY* and m-*ENVY-FREENESS*, which are derived from the recommender systems literature; and exploit them for the first time for the evaluation of multiwinner rules (see Section 4.3).

[5]Of course, proportionality is only one notion of fairness provided by some multiwinner election mechanims. In other cases, one may want to define a notion of fairness based on the properties of shortlisting or diversity that other multiwinner election mechanisms satisfy.

Chapter 4

Experimental Evaluation

In this chapter we present a series of experiments to evaluate: *(i)* the ability of our approaches to learn the preferences of a user, i.e., a tourist, and provide high quality recommendations *wrt* her preferences; *(ii)* the performance of several multiwinner election mechanisms with respect to the quality of the final recommendations; and *(iii)* the effectiveness of the RAV multiwinner election mechanism for the problem of group recommendations with respect to the fairness metrics of *m-PROPORTIONALITY* and *m-ENVY-FREENESS* derived from the recommender system literature. We tested our proposed approaches on a real-world dataset including 430 POIs of a popular tourist resort on a Greek island. Alongside the available POIs, we used 90 generic images for the preference elicitation process. Note that for this series of experiments we assume that each user, POI, and generic image can be described as a multivariate normal distribution, where $D = 12$. Specifically, we consider the following travel-related features: *Culture, Sun & Sea, History/Archaeology, Adventure/Sports, Affordable Prices, Family-friendly activities & facilities, Rural Tourism, Luxury Accommodation, Nightlife, Gastronomy/Cuisine, General Shopping* and *Shopping Local Products*. We note that the Boltzmann exploration parameters for all series of experiments were set to: $c = 1$, $\alpha = 0.5$ and $t \leq 3$, with $t_0 = 0$ and $t_{i+1} = t_i + 1$.

4.1 Datasets

First of all, we constructed a real-world dataset that contains 430 real-world POIs that are located in the region of Agios Nikolaos, Crete. We highlight that this dataset was designed with the cooperation of Technical University of Crete, Hellenic Mediterranean University and the Municipality of Agios Nikolaos city for the needs of "*ViP, Visit Planner: Integrated Information and Tour Planning Service for Cruise Tourism based on Hybrid Recommender Systems*" research project (project code: T2EDK-03135). In more detail, this dataset contains museums, archaeological sites, monuments, restaurants, cafe-bars, islands, beaches etc. that are located in the city of Agios Nikolaos. Finally, we created a generic-image dataset that contains 90 generic images that capture general types of POIs that can be found in Agios Nikolaos, e.g., an image of a restaurant, a beach, a museum, a monument etc. This dataset is exploited by our system in order to learn the preferences and the interests of each individual in order to provide personalized recommendations of POIs (that are stored in the aforementioned POIs dataset) for a targeted-user. Once again, we mention that the dataset of the generic images and the dataset of POIs of Agios Nikolaos do not share any common element (i.e., there are no items that belong to both datasets).

4.2 Personalized Recommendations

First we present a series of experiments performed to evaluate the greedy version of our approaches for personalized recommendations, using synthetic tourists. In more detail, using preference-related data collected from actual tourists via questionnaires, we generate 500 synthetic tourists for the age groups of 18-25, 26-35, 36-45, 46-55, 56-67 and 67 plus respectively, i.e., 3000 synthetic users in total. The average of the answers (i.e., values in the scale of 1...5 recorded on the questionnaires) of each age group for the 12 features provided by the questionnaires is calculated for each corresponding feature. In order to generate the synthetic users for each age group, we take 500 samples from a multivariate normal distribution, whose mean vector is the average values of the 12 features; the covariance matrix is diagonal and each diagonal element takes a value equal to 1. We run a series of user-system simulations with varying slate size—i.e., number of generic images $n = \{1, 2, 3, 4, 5\}$ that are presented to a user in each question asked, and number of questions $m = \{1, 2, 3, 4, 5\}$ asked. Finally, for each $\langle n, m \rangle$ combination we ran 1000 simulations: in every simulation, we randomly pick a tourist out of our 3000 synthetic users. As an evaluation metric, we compare the real distribution that represents user's actual preferences with her inferred one—i.e., the distribution that our algorithms constructed via the preference elicitation method (see Section 3.1)—by using Equation 3.2. In what follows, we denote this resulting "predicted rating" metric by r for short. Notice that had the inferred distribution been *identical* to the real one, the value of this metric would be equal to 5, i.e., the maximum rating that an item can receive by a user.

4.2.1 No Prior Knowledge Employed

Table 4.1 illustrates the results of the "greedy" Bayesian learning approach on a first set of experiments that does not exploit any prior knowledge regarding the user—i.e., the system does not have any information regarding the user's age group, but uses an uninformative prior (i.e., one with a mean vector that contains in each dimension a value of 1, and a diagonal covariance matrix where each diagonal element takes a value equal to 2). Note that the presented results are average values over 1000 simulations for each $\langle n, m \rangle$ combination setting tested. We can see that for n fixed across different settings, the r metric achieved by our algorithm increases as the number of questions, i.e., m, increases. Such a result is expected, since when the system makes more questions to a user, more information regarding her preferences is revealed, and as such our approach builds a better model with respect to the user's preferences. Also, for a fixed number of m questions, we observe that as n (i.e., the number of alternative pictures shown per question) increases, the r achieved increases as well. This is due to the fact that when the system provides more options to a user, then a picture that best captures her preferences is easier to be found. Thus, by increasing the number of options n, our system is able to build a better user model, since it exploits better quality information regarding user interests.

Similarly, Table 4.2 demonstrates the corresponding results of the "greedy" Kalman Filter approach on the aforementioned set of experiments. In general, our results showed that this approach performs slightly better comparing to the Bayesian learning recommender in two cases: *(i)* when $m = 1$, i.e., when there is a single user-system interaction; and *(ii)* for the combination $m = 2$ and $n = 1$. However, we

	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
$m = 1$	2.6	2.85	3.02	3.12	3.24
$m = 2$	2.74	3.06	3.24	3.39	3.5
$m = 3$	2.81	3.15	3.4	3.54	3.67
$m = 4$	2.86	3.23	3.5	3.67	3.82
$m = 5$	2.9	3.36	3.61	3.78	3.92

TABLE 4.1: Bayesian learning: The r results using uninformative priors (no prior knowledge exploited). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
$m = 1$	2.95	3.07	3.21	3.26	3.27
$m = 2$	2.89	2.98	3.15	3.24	3.29
$m = 3$	2.75	2.91	3.11	3.16	3.23
$m = 4$	2.61	2.85	2.98	3.11	3.19
$m = 5$	2.52	2.75	2.93	3.02	3.05

TABLE 4.2: Kalman Filter: The r results using uninformative priors (no prior knowledge exploited). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

notice that for most of $\langle n, m \rangle$ combinations the Bayesian learning approach outperforms the system that applies the Kalman Filter algorithm with respect to the r metric. Indicatively, for $n = 5$ and $m = 5$ the first approach achieves a score of 3.92 (out of 5), while Kalman Filter has a score of 3.05 (out of 5) (i.e., 17.4%) better performance, while for $n = 5$ and $m = 3$ (which is the actual combination that the real-world mobile application utilizes) the first approach achieves a score of 3.67, while Kalman Filter has a score of 3.23 (i.e., 8.8%) better system performance. Additionally, Table 4.2 indicates that as the number of questions, i.e., m, increases, the performance of the Kalman approach slightly drops. Such an effect can be justified with the fact that the Kalman Filter approach applies a *heuristic* mechanism which simply modifies the parameters of the user model and then, rather naively, assumes that the produced user model instantiation constitutes an actual observation for the Kalman Filter to use. By contrast, the Bayesian method uses a large number of samples from the item distribution in order to update its belief regarding user model. Apparently, this method exploits better the rich information incorporated in the item distribution, leading to more accurate user model. The Kalman Filter method, on the other hand, appears to be unable to treat the ratings of the user and the corresponding uncertainty in order to update its belief regarding user's interests. This results to a decrease of its performance with respect to the r metric.

Following that, we used another metric that captures the effectiveness of our approach based on the *accuracy* (or quality) of the recommended POIs. Specifically, for each individual that interacts with the system, we use her "real" distribution in order to create the list l_{real}, which contains the *top-N* POIs for this specific user (i.e., given a user u, the N POIs from the dataset that score the highest $r_{u,i}$ scores). At the same time, we use the inferred model for this user (i.e., the distribution that our algorithms created via the preference elicitation procedure), in order to create the list l_{inf}, which contains the *top-20* POIs for the inferred user, i.e., the final personalized recommendations that each approach produced. Thus, we can measure the similarity between the lists l_{real} and l_{inf} by finding the number of *common elements* (*com*), i.e., POIs, that these two lists share: that is, we count how many of the (real) user's *top-N* items coincide with ones in the recommended *top-20* list of items (given the inferred user model). Formally:

$$com = \frac{|l_{real} \cap l_{inf}|}{20} \tag{4.1}$$

For this set of experiments we set $N = \{20, 43, 86\}$. Note that, the "top-20" corresponds to only the 4.5% of our dataset. Similarly, the "top-43" constitutes only the 10% of our dataset, while "top-86" the 20%.

Table 4.3 shows the Bayesian learning system's average score of 1000 simulations for each combination of $\langle n, m \rangle$. As seen, for a given n (or a given m), the percentage of common elements that lists l_{real} and l_{inf} share, is rising as m (or n) rises. This is natural, since for larger n (or m) our system collects more information regarding user's preferences and as such is able to provide better recommendations. Indicatively, in settings with $n = 5, m = 5$ our agent recommends POIs with 38.82% of them being among the best 20 POIs of each user. Accordingly, 57.81% of them being among the best 43 POIs and 75.19% of them being among the best 86 POIs of each user. Thus, after only a small number of interactions with each user, our approach is able to provide recommendations that match the user interests to a large extent.

Table 4.4 presents the Kalman Filter system's average score of 1000 simulations for each combination of $\langle n, m \rangle$. As the results show this approach has similar behaviour with the Bayesian learning approach, i.e., as n (or m) increase the percentage of common elements rises, i.e., the recommendations' quality that our system provides to the user is better with respect to the interests of the user. However, we notice that for each $\langle n, m \rangle$ combination the Bayesian learning approach outperforms the system that applies the Kalman Filter algorithm. Indicatively, for $n = 5$ and $m = 5$ the Bayesian learning approach achieves a score of 38.82% (in the case of top-20), while Kalman Filter has a score of 31.07% , while for $n = 5$ and $m = 3$ (which is the actual combination that the real-world mobile application utilizes) the first approach achieves a score of 35%, while Kalman Filter has a score of 22.27%. This result is natural since the Bayesian learning approach creates a better model than the one that utilizes the Kalman Filter algorithm, with respect to the r metric, as described earlier in this section. Thus, the former approach is able to produce more efficient personalized recommendations when no prior knowledge is available to our system.

Moreover, as mentioned earlier, we notice that in Table 4.2 the performance of the Kalman Filter approach drops when the number of questions increases, while it achieves higher scores compared to the Bayesian learning approach with respect to the r metric (essentially only) when $m = 1$. Now, Tables 4.3 and 4.4 indicate that for all $\langle n, m \rangle$ the approach of Bayesian learning is the one that *constantly* produces higher quality recommendations (even when $m = 1$) with respect to the top-20, top-43 and top-86 metrics. This phenomenon can be justified by the fact that Kalman Filter is unable to treat appropriately the uncertainty that exists in the setting, since it uses a heuristic mechanism, and may produce user models that are extremely confident and general. Thus, even when Kalman Filter scores better than Bayesian learning with respect to the r metric, it is conceivable that it will fail to produce top quality recommendations (as, in fact, is shown in our results). For instance, if our system is confident that the user is "neutral" for every available travel-related feature—i.e., the user mean vector has a value of 2.5 for each feature—then the inferred distribution would have a relatively high r score irrespective of the true user model, but it would be unable to produce high quality recommendations for the user since the system's belief would not be able to focus on POIs that the user prefers, since system's belief is that the user is "neutral" for every travel-related feature.

		$n=1$	$n=2$	$n=3$	$n=4$	$n=5$
$m=1$	$top-20$	13.75%	17.9%	21.52%	24.11%	27.3%
	$top-43$	25%	32.66%	35.91%	40.92%	43.08%
	$top-86$	38.43%	47.49%	54.43%	55.95%	59.53%
$m=2$	$top-20$	13.86%	21.17%	23.94%	27.15%	30.32%
	$top-43$	25.27%	32.78%	41.19%	44.67%	47.33%
	$top-86$	38.49%	49.37%	54.93%	60.73%	65.37%
$m=3$	$top-20$	14.06%	21.94%	27.5%	30.29%	35%
	$top-43$	25.58%	34.67%	42.94%	46.79%	52%
	$top-86$	38.63%	53.92%	59.94%	65.89%	70.41%
$m=4$	$top-20$	14.31%	23.3%	29.28%	33.36%	36.16%
	$top-43$	27.01%	36.52%	43.66%	51.63%	55.78%
	$top-86$	42.09%	54.15%	63.48%	69.67%	72.67%
$m=5$	$top-20$	14.48%	23.46%	30.37%	35.74%	38.82%
	$top-43$	28.56%	37.92%	47.4%	53.58%	57.81%
	$top-86$	42.73%	56.14%	65.44%	72.08%	75.19%

TABLE 4.3: Bayesian learning: Similarity of l_{real} and l_{inf} for "top-20", "top-43" and "top-86" (uninformative priors). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

		$n=1$	$n=2$	$n=3$	$n=4$	$n=5$
$m=1$	$top-20$	4.2%	5.91%	7.83%	9.47%	9.71%
	$top-43$	9.12%	12.51%	14.57%	16.28%	17.18%
	$top-86$	16.19%	21.23%	23.96%	26.07%	29.95%
$m=2$	$top-20$	7.51%	9.18%	13.59%	15.74%	16.09%
	$top-43$	14.7%	17.86%	22.14%	26.24%	29.53%
	$top-86$	24.16%	32.08%	38.54%	44.73%	49.18%
$m=3$	$top-20$	7.6%	12.05%	16.17%	18.42%	22.27%
	$top-43$	14.92%	24.63%	28.22%	33.55%	37%
	$top-86$	25.83%	37.21%	46.85%	50.28%	57.02%
$m=4$	$top-20$	7.38%	13.14%	19.21%	24.83%	29.01%
	$top-43$	15.88%	24.85%	32.2%	40.8%	45.29%
	$top-86$	26.75%	40.36%	51%	60.23%	66.02%
$m=5$	$top-20$	7.94%	15.83%	23%	29.47%	31.07%
	$top-43$	16.78%	27.62%	36.15%	46.8%	51.04%
	$top-86$	30.15%	44.74%	55.81%	63.5%	68.32%

TABLE 4.4: Kalman Filter: Similarity of l_{real} and l_{inf} for "top-20", "top-43" and "top-86" (uninformative priors). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

4.2.2 Employing Prior Knowledge

In this set of experiments, we study the scenario in which our system has at its disposal some prior knowledge regarding general user types preferences. Specifically, we construct some age-related priors—i.e., prior distributions regarding the general preferences of tourists that belong to the same age group, by exploiting data collected via questionnaires from actual visitors. These priors are constructed in the exact same way as the ones from which we draw the synthetic users from, with the only difference being that we insert higher uncertainty, i.e. the diagonal covariance matrix elements have a value equal to 2. In this scenario, we study the performance of our recommender systems for each age group and we compare it with the case where no such extra information is available to us. For each age group, we generated 1000 users via the process already described in the beginning of this section, along with their corresponding priors.

Tables 4.5, 4.9, 4.7 and 4.11 depicts the average results of the approach that employs Bayesian learning over 1000 simulations for the age group of 18-25, 26-35, 36-45, 46-55, 56-67 and 67 plus, while Tables 4.6, 4.10, 4.8 and 4.12 depicts the corresponding average results of the approach that employs the Kalman Filter. In more detail, Tables 4.5, 4.6, 4.7 and 4.8 capture the performance of our approaches for each age group when no prior information is available. Again, we use the r metric in order to evaluate the inference of our system and the number of recommended POIs that belongs to the "top-20" of the user. Tables 4.9, 4.10, 4.11 and 4.12 present the corresponding results of our approaches when our system knows the user's age group and can thus exploit its prior knowledge regarding this age group's preferences.

First, we notice that with such prior knowledge in their disposal, both of our recommender systems significantly and consistently outperform their (corresponding) versions that has no prior knowledge available. Moreover, the recommenders are then able to provide high-scoring recommendations with only limited interaction among the user and the system, i.e. when the values of n and m are small, vastly outperforming the uninformed version. However, as the n and m values increase, the margins between the performance of the two versions decrease. Such a result is expected since the prior knowledge gives insights to our systems regarding the user's preferences, when we have limited information about a user's interests. In other words, the prior knowledge helps our systems to efficiently deal with the cold-start problem. Notably, the results illustrated in Tables 4.9 and 4.10 indicate that there is no clear winner regarding the performance of our approaches with respect to the r metric when prior knowledge is available to our system

Furthermore, the results presented in Tables 4.11 and 4.12 indicate that the approach that utilizes the Kalman Filter algorithm is consistently better than the approach that employs Bayesian learning for every examined age group. As such, the recommender system that employs the Kalman Filter algorithm is able to provide more efficient recommendations with respect to the interest of a targeted-user when (high quality) prior knowledge is in our disposal. However, given the way the generic age-related priors are constructed and the fact that the corresponding synthetic users' distributions are closely related to the priors, this result has to be taken with a grain of salt: Kalman Filter constructs and exploits a single Gaussian "observation" which is very close to the very high quality prior provided, thus calling off the need for and diminishing the value of sampling the distribution (as Bayesian learning does).

	Age Group	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
$m = 1$	**18 − 25**	2.64	2.95	3.17	3.29	3.39
	26 − 35	2.82	3.13	3.24	3.39	3.46
	36 − 45	2.93	3.27	3.43	3.54	3.64
	46 − 55	2.98	3.25	3.44	3.53	3.6
	56 − 67	2.72	2.95	3.14	3.26	3.36
	67+	2.49	2.68	2.83	2.95	3.06
$m = 2$	**18 − 25**	2.81	3.15	3.4	3.54	3.68
	26 − 35	2.97	3.3	3.49	3.65	3.73
	36 − 45	3.1	3.44	3.64	3.78	3.87
	46 − 55	3.19	3.44	3.64	3.75	3.83
	56 − 67	2.88	3.17	3.36	3.49	3.58
	67+	2.69	2.9	3.08	3.09	3.31
$m = 3$	**18 − 25**	2.85	3.3	3.57	3.74	3.9
	26 − 35	3.03	3.43	3.71	3.82	3.89
	36 − 45	3.13	3.57	3.81	3.96	4.06
	46 − 55	3.21	3.57	3.79	3.92	3.99
	56 − 67	2.93	3.3	3.52	3.66	3.81
	67+	2.71	3.01	3.22	3.38	3.51

TABLE 4.5: Bayesian learning: Performance with no available prior knowledge (r metric). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

	Age Group	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
$m = 1$	**18 − 25**	2.27	2.43	2.57	2.64	2.96
	26 − 35	2.82	3.05	3.13	3.21	3.26
	36 − 45	2.7	3.01	3.11	3.21	3.23
	46 − 55	3.54	3.73	3.77	3.81	3.84
	56 − 67	3.33	3.46	3.52	3.54	3.63
	67+	3.4	3.46	3.56	3.6	3.61
$m = 2$	**18 − 25**	2.29	2.56	2.72	2.82	2.92
	26 − 35	2.83	3.06	3.21	3.28	3.32
	36 − 45	2.78	3.02	3.18	3.28	3.33
	46 − 55	3.45	3.59	3.67	3.75	3.78
	56 − 67	3.21	3.41	3.49	3.54	3.6
	67+	3.28	3.4	3.47	3.52	3.58
$m = 3$	**18 − 25**	2.34	2.59	2.7	2.87	2.95
	26 − 35	2.77	3.05	3.16	3.25	3.29
	36 − 45	2.73	3.01	3.14	3.25	3.28
	46 − 55	3.28	3.48	3.55	3.62	3.65
	56 − 67	3.06	3.24	3.38	3.43	3.49
	67+	3.08	3.24	3.35	3.4	3.48

TABLE 4.6: Kalman Filter: Performance with no available prior knowledge (r metric). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

	Age Group	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
$m = 1$	**18 − 25**	15.01%	20.21%	23.95%	26.45%	27.5%
	26 − 35	12.71%	18.56%	21.56%	23.45%	25.06%
	36 − 45	16.64%	23.01%	26.54%	30.32%	31.31%
	46 − 55	14.78%	21.02%	24.28%	26.39%	28.76%
	56 − 67	13.99%	18.81%	22.84%	26.23%	27.9%
	67+	14.8%	17.95%	20.82%	23.75%	25.75%
$m = 2$	**18 − 25**	14.76%	20.38%	25.9%	29.12%	32.17%
	26 − 35	12.76%	20.27%	23.4%	27.4%	28.34%
	36 − 45	17.55%	25%	29.8%	35.29%	36.76%
	46 − 55	16.78%	22.59%	29.14%	30.64%	33.35%
	56 − 67	15.53%	21.32%	25.64%	28.51%	30.5%
	67+	17.9%	19.64%	23.93%	24.05%	28.4%
$m = 3$	**18 − 25**	14.9%	23.97%	28.53%	31.47%	35.18%
	26 − 35	12.8%	21.49%	27.34%	29.15%	30.86%
	36 − 45	17.18%	27.96%	34.41%	39.09%	41.28%
	46 − 55	16.79%	25.65%	29.98%	32.9%	36.99%
	56 − 67	14.46%	22.93%	28.49%	31.56%	35.6%
	67+	15.32%	21.08%	25.37%	28.08%	32.16%

TABLE 4.7: Bayesian learning: Performance with no available prior knowledge (Similarity of l_{real} and l_{inf} for "top-20"). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

	Age Group	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
$m = 1$	**18 − 25**	2.14%	3.18%	4.45%	4.97%	13.52%
	26 − 35	2.8%	4.4%	5.99%	6.81%	7.63%
	36 − 45	2.92%	5.79%	6.88%	8.28%	8.34%
	46 − 55	6.45%	11%	13.11%	13.94%	15.49%
	56 − 67	4.51%	6.28%	8.12%	9.28%	11.07%
	67+	3.94%	5.32%	7.64%	8.15%	8.99%
$m = 2$	**18 − 25**	2.98%	5.79%	7.6%	9.02%	9.78%
	26 − 35	5.93%	9%	11.43%	12.16%	12.96%
	36 − 45	5.45%	9.27%	11.41%	13.92%	17.57%
	46 − 55	11.61%	16.42%	19.58%	22.5%	22.79%
	56 − 67	6.85%	13.99%	15.01%	17.61%	19.39%
	67+	6.85%	12.2%	13.46%	16.71%	18.27%
$m = 3$	**18 − 25**	4%	6.89%	8.89%	12.6%	15.03 %
	26 − 35	6.28%	10.97%	13.68%	16.82%	19.14%
	36 − 45	6.32%	11.95%	15.97%	21.74%	24.33%
	46 − 55	11.04%	18.65%	24.08%	29.11%	32.11%
	56 − 67	7.49%	14.01%	20.69%	22.73%	27.5%
	67+	7.77%	13.43%	16.82%	20.57%	26.98%

TABLE 4.8: Kalman Filter: Performance with no available prior knowledge (Similarity of l_{real} and l_{inf} for "top-20"). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

	Age Group	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$
$m=1$	**18 − 25**	4.11	4.12	4.13	4.13	4.13
	26 − 35	3.94	3.98	3.99	4	4.01
	36 − 45	4.04	4.1	4.12	4.13	4.14
	46 − 55	3.94	4	4.01	4.02	4.02
	56 − 67	3.97	4	3.98	3.98	3.98
	67+	3.68	3.71	3.72	3.74	3.73
$m=2$	**18 − 25**	4.21	4.15	4.16	4.18	4.19
	26 − 35	3.84	3.93	3.96	4.06	4.07
	36 − 45	4.03	4.14	4.15	4.18	4.21
	46 − 55	3.86	4.05	4.07	4.06	4.13
	56 − 67	3.95	3.96	3.97	3.99	3.99
	67+	3.78	3.79	3.77	3.76	3.77
$m=3$	**18 − 25**	4.19	4.2	4.17	4.2	4.21
	26 − 35	3.75	3.97	4	4.09	4.11
	36 − 45	3.97	4.18	4.19	4.22	4.25
	46 − 55	3.87	4.07	4.1	4.1	4.12
	56 − 67	3.9	3.96	3.97	3.98	4.01
	67+	3.75	3.78	3.78	3.76	3.78

TABLE 4.9: Bayesian learning: Performance with available prior knowledge (r metric). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

	Age Group	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$
$m=1$	**18 − 25**	4.39	4.4	4.4	4.41	4.41
	26 − 35	4.44	4.44	4.45	4.46	4.46
	36 − 45	4.41	4.4	4.41	4.41	4.41
	46 − 55	4.43	4.44	4.45	4.45	4.45
	56 − 67	4.43	4.43	4.43	4.43	4.43
	67+	4.48	4.5	4.5	4.5	4.5
$m=2$	**18 − 25**	4.08	4.12	4.14	4.14	4.14
	26 − 35	4.11	4.12	4.13	4.14	4.14
	36 − 45	4.06	4.09	4.09	4.09	4.1
	46 − 55	4.09	4.09	4.09	4.11	4.12
	56 − 67	4.09	4.09	4.1	4.1	4.13
	67+	4.13	4.12	4.13	4.16	4.18
$m=3$	**18 − 25**	3.85	3.88	3.89	3.9	3.92
	26 − 35	3.83	3.85	3.86	3.87	3.89
	36 − 45	3.79	3.84	3.83	3.85	3.86
	46 − 55	3.8	3.82	3.84	3.86	3.87
	56 − 67	3.79	3.83	3.83	3.87	3.88
	67+	3.82	3.85	3.87	3.9	3.94

TABLE 4.10: Kalman Filter: Performance with available prior knowledge (r metric). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

	Age Group	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
$m = 1$	**18 − 25**	45.86%	45.06%	45.1%	46.62%	46.68%
	26 − 35	29.38%	32.03%	30.9%	34.83%	34.54%
	36 − 45	48.52%	53.03%	52.36%	49.44%	50.3%
	46 − 55	41.15%	36.68%	39.32%	39.42%	40.43%
	56 − 67	44%	43.78%	42.96%	43.15%	42.72%
	67+	41.53%	42.84%	43.09%	43.94%	43.43%
$m = 2$	**18 − 25**	42.4%	37.05%	38%	38.68%	38.51%
	26 − 35	21.62%	25.1%	30.87%	40.14%	37.26%
	36 − 45	47.74%	50.4%	47.86%	49.95%	48.95%
	46 − 55	27.19%	36.33%	38.9%	38.07%	41.75%
	56 − 67	31.02%	34.51%	40.23%	39.43%	44.23%
	67+	33.27%	39.75%	39.6%	44.64%	44.14%
$m = 3$	**18 − 25**	37.25%	34.71%	35.2%	37.02%	38.25%
	26 − 35	19.28%	30.45%	32.77%	37.65%	36.63%
	36 − 45	46.86%	49.83%	47.89%	48.97%	49.83%
	46 − 55	27.2%	39.98%	43.74%	41.08%	41.19%
	56 − 67	30.72%	39.96%	41.44%	41.2%	45.31%
	67+	33.95%	41.82%	42.09%	45.14%	45.64%

TABLE 4.11: Bayesian learning: Performance with available prior knowledge (Similarity of l_{real} and l_{inf} for "top-20"). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

	Age Group	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
$m = 1$	**18 − 25**	49.71%	50.77%	50.59%	51.68%	53.14%
	26 − 35	48.88%	45.92%	47.37%	49.48%	48.66%
	36 − 45	56.58%	57.55%	60.14%	59.95%	59.8%
	46 − 55	55.76%	55.7%	57.31%	55.45%	56.93%
	56 − 67	58.16%	57.14%	58.34%	56.64%	56.87%
	67+	57.66%	54.71%	55.84%	57.31%	56.76%
$m = 2$	**18 − 25**	50.06%	50.72%	50.87%	51.89%	51.96%
	26 − 35	44.75%	42.48%	41.93%	45.89%	46.38%
	36 − 45	57.12%	58.59%	58.48%	56.07%	56.94%
	46 − 55	55.22%	54.76%	54.71%	52.31%	52.55%
	56 − 67	55.91%	55.41%	56.07%	54.89%	52.8%
	67+	54.44%	53.46%	52.65%	52.67%	53.9%
$m = 3$	**18 − 25**	45.82%	50.03%	49.59%	49.65%	51.05%
	26 − 35	44.57%	39.25%	38.26%	43.92%	47%
	36 − 45	53.9%	57.2%	54.05%	55.96%	60.88%
	46 − 55	47.08%	51.43%	52.11%	47.77%	48.35%
	56 − 67	46.77%	48.35%	47.27%	53.35%	51.7%
	67+	47.37%	50.2%	50.08%	49.55%	52.91%

TABLE 4.12: Kalman Filter: Performance with available prior knowledge (Similarity of l_{real} and l_{inf} for "top-20"). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

Nevertheless, these results do show that high-quality priors can indeed be exploited by our systems. Moreover, the experiment provides further insights to the methods' behaviour, given the following interesting observation: Notice that, with prior knowledge available, there are cases where the performance of our approaches in terms of recommendations made using KL-divergence, drops with more questions asked or images displayed (cf. Table 4.11 and Table 4.12). We attribute this behaviour to the fact that apparently our age-related prior is of a very high quality (indeed, it is very similar to the synthetic user's distribution), and is thus able to capture the preferences of a user to a very large extent—i.e., it describes her interests for every travel-related feature that we have used in our approach. By contrast, when the preference elicitation process kicks in, we present generic images that cannot but represent a very specific type of POIs (e.g., a beach, a restaurant etc.). Thus, when the visitor selects and rates a picture, our approaches are able to exploit data that are related to the selected generic image in order to update their *belief* regarding user's preferences. In such a case, the belief regarding user's interests will be improved only for specific features (that are representatives of the selected generic image), causing a small drop in recommendations' quality.

4.2.3 Multiwinner Elections for Personalized Recommendations

Here we evaluate our social choice-inspired mechanism for generating our final recommendations to a tourist. We note that since no prior knowledge is exploited in this set of experiments, we chose to employ *only* the Bayesian learning approach since it constantly outperformed the version of Kalman Filter algorithm (see Section 4.2.1). To this end, we created 1000 synthetic tourists following the exact same procedure already described earlier (see Section 4.2), and applied the following well-known aggregation strategies: *(i)* Approval Voting (AV), *(ii)* Reweighted Approval Voting (RAV), *(iii)* Bloc, and *(iv)* k-Borda. We chose to evaluate our approach by measuring the similarity between the lists l_{real} and l_{inf} for the *top-20* POIs of each user in order to study the trade-off between the *accuracy* and the *diversity* of the final recommendations.

Table 4.13 illustrates the results of our approach on this set of experiments. Note that the presented results are the average values over 1000 simulations of experiments for each $\langle n, m \rangle$ combination. First, we observe that when the values of n and m are small (i.e. 1 or 2), the greedy approach usually ranks lower in terms of similarity score among most multiwinner election mechanisms methods—while for $n = 1$ and $n = 2$ *all* multiwinner election mechanisms perform better than "greedy" with respect to our metric. These performance results are natural, since for smaller n and m we collect limited information regarding the interests of the user. As a result, our inferred model does not describe the user preferences very accurately. Thus, the selection of the POIs based on the model inferred so far—given only the very limited information provided by the user—is not the optimal policy. On the other hand, the employment of multiwinner election mechanisms allows us to move beyond this suboptimal greedy approach, while it still exploits our inferred model (see Section 3.5). However, as n and m increase, and our system is able to collect more information regarding the interests of the tourist, we notice that the *greedy* performs better than most of the voting rules, but its advantage is in most cases not significant—in contrast to what was the case for low n and m values, where the advantage of the social choice-inspired methods was broad. Additionally, the employment of the voting rules results to a diverse "elected committee", corresponding to

	Rule	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
$m = 1$	Greedy	14.07%	19.83%	22.29%	24.42%	26.44%
	AV	22.36%	**26.15%**	**27.76%**	**28.7%**	**29.63%**
	RAV	**22.75%**	24.53%	25.78%	26.12%	26.65%
	Bloc	19.13%	23.48%	24.66%	23.73%	25.64%
	k-Borda	21.59%	24.88%	25.65%	26.66%	27.61%
$m = 2$	Greedy	15.34%	20.56%	24.39%	27.86%	30.31%
	AV	22.97%	**26.24%**	**28.21%**	**29.93%**	**30.85%**
	RAV	**23.25%**	24.99%	25.67%	26.6%	27.2%
	Bloc	20.83%	21.93%	23.68%	26.11%	25.57%
	k-Borda	22.08%	25.49%	25.89%	27.37%	28.97%
$m = 3$	Greedy	15.05%	22%	26.92%	**30%**	**34.17%**
	AV	22.81%	**26.03%**	**28.62%**	29.94%	31.59%
	RAV	**23.26%**	24.82%	25.85%	26.84%	27.56%
	Bloc	19.91%	22.64%	24.84%	26.76%	28.86%
	k-Borda	21.23%	24.54%	27.16%	28.6%	28.83%

TABLE 4.13: Performance of aggregation mechanisms with no available prior knowledge (Similarity of l_{real} and l_{inf} for "top-20"). The results are average values over 1000 simulations for each $\langle n, m \rangle$ combination.

diverse final recommendations with respect to the travel-related features. Thus, we see that such mechanisms can help us *(i)* achieve better performance regarding the quality of our recommendations when the user-system interaction is limited; and *(ii)* produce various POIs of different types, by simply paying a small quality penalty compared to the performance of the greedy approach.

Finally, we notice that when $n = 1$ the RAV rule performs better with respect to our evaluation metric, showing that a proportional representation of the travel-related features (our "voters"), is a rational choice when limited information is in our disposal. As the user-system interaction increase, and more information regarding her interests is available to our system, the AV approach achieves the highest performance.

4.3 Group Recommendations

In our experiments we evaluate the performance of different aggregation mechanisms for group recommendations, in terms of the standard *m-PROPORTIONALITY* and *m-ENVY-FREENESS* fairness metrics [73]. These respectively signify the percentages of users that consider m items in a recommended set to be either in their top Δ_p% of preferred items, or for which the user belongs in some top Δ_e% of users that are favored by the recommendation of these items (these will be clarified more below). We created synthetic groups g of users of various sizes—specifically, $|g| = \{5, 10, 15, 20\}$—and applied the following aggregation strategies: *(i)* Least Misery (LM), *(ii)* Most Pleasure (MP), *(iii)* Additive Utilitarian (AU), and *(iv)* Reweighted Approval Voting (RAV).

In more detail, for the RAV mechanism, we assume that a user *approves* an item, i.e., a POI, if and only if her expected score (computed from Equation 3.2) for this item is larger than 3. We note that in order to compute the expected user's rating for an item we use her inferred model (i.e. the model that our approach constructed via our preference elicitation process) and not her real one (since we want to assess our system's ability to provide fair recommendations, and the real user model is not known to the system). For the *m-PROPORTIONALITY* fairness metric we assume that Δ_p for all users is set to 0.1—i.e., a user *likes* a POI, if this POI is ranked in the top-10% of the user's preferences over all available POIs in the dataset. Additionally, for the *m-ENVY-FREENESS* fairness metric we assume that $\Delta_e = 0.4$—i.e., a user is *envy-free* for a POI, if for this POI the user is in the favored top-40% of the group (i.e, in the 40% of the group members that prefer the POI more than the rest 60%). Finally, for both *m-PROPORTIONALITY* and *m-ENVY-FREENESS* metrics we consider that the m items for which the corresponding property is required is set to:

$$m = \frac{\textit{\# of recommended POIs}}{|g|} \tag{4.2}$$

where *# of recommended POIs* = 20.

Group size	Metrics	RAV	MP	LM	AU
$\|g\| = 5$	m-PROPORTIONALITY	0.93	0.79	0.72	0.83
	m-ENVY-FREENESS	0.79	0.62	0.65	0.62
$\|g\| = 10$	m-PROPORTIONALITY	0.96	0.89	0.87	0.88
	m-ENVY-FREENESS	0.92	0.79	0.75	0.78
$\|g\| = 15$	m-PROPORTIONALITY	0.98	0.97	0.95	0.94
	m-ENVY-FREENESS	0.97	0.92	0.86	0.92
$\|g\| = 20$	m-PROPORTIONALITY	0.98	0.97	0.95	0.94
	m-ENVY-FREENESS	0.97	0.92	0.85	0.92

TABLE 4.14: Fairness results (averages over 500 simulations per $|g|$).

Again, we note that since no prior knowledge is exploited in this set of experiments, we chose to employ *only* the Bayesian learning approach since it constantly outperformed the version of Kalman Filter algorithm (see Section 4.2.1). Table 4.14 illustrates the results of our approach on this set of experiments. Note that the presented results are the average values over 500 simulations of experiments on settings with the same properties—i.e., we randomly generated 500 groups for each group size $|g|$ value, and ran one such simulation per generated group (i.e., we ran 2000 simulations in total). We can see that the RAV mechanism is able to provide very efficient group recommendations for every setting with respect to our fairness metrics, achieving consistently better performance compared to the other aggregation mechanisms. In fact, the RAV mechanism usually achieves a score over 92%, i.e., the proportionality and envy-freeness properties are achieved for almost all members of the group, irrespective of group size (with only one exception, for $|g| = 5$ when a 4-*ENVY-FREENESS* of 79% is achieved). Notice also that RAV's performance is significantly better than that of the other mechanisms for the smaller group sizes. In the case of larger groups, however, every aggregation mechanism achieves very high scores (signifying fair group recommendations), which are also comparable to each other. Such a result is expected, since for larger groups the m parameter for both metrics decreases (see Equation 4.2). Thus, we need fewer items (i.e., m items) in the final set of recommendations, in order to consider this set of POIs fair for a member of the group with respect to our metrics. Additionally, for larger groups it is easier to find members that share similar interests—e.g., it is easier to recommend POIs that satisfy more than one members belonging to the group.

4.3.1 Discussion

In general, many researchers have focused on fairness notions in the recommender systems literature. Note also that a recommender can use only the inferred model for generating recommendations. As such, as is natural and common in the literature, we exploit the *inferred* user model instead of the real one for producing group recommendations.

Of course, the real user model can be exploited, for the evaluation of the *elicitation* procedure—i.e., for answering the question "How well our system has learned the (real) preferences of the user, via the selected elicitation process?". (We tackle this question when evaluating single-user recommendations, using the real model of the user and the corresponding metrics (see Section 4.2).) However, we believe that it is *not* appropriate to exploit the real user model to evaluate the mechanisms for group recommendations in real-world systems. To explain this, assume that we employ any multiwinner mechanism that is able to perform *perfectly* with respect to our metrics. If our elicitation process cannot learn the users' preferences efficiently (e.g., due to a small number of interactions), then when evaluating our multiwinner mechanism with respect to m-*PROPORTIONALITY* and m-*ENVY-FREENESS*, the results would show that the mechanism does not perform well, when evaluated with respect to the real users' models. This effect however would be due to the fact that our inferred model does not describe efficiently the real users' preferences, and not because our mechanism is unable to provide a committee that satisfies to a large extent our selected metrics. Hence, we do not consider it appropriate to use the real user models for evaluating the multiwinner election mechanisms, as by doing so we would not be able to draw clear conclusions as to which component is responsible for potential poor performance of the real-world system, as our example indicates.

On the other hand, as discussed in Chapter 5, one could add to the pipeline a step in which the real users evaluate the final group recommendations (indeed, our real-world application includes such a step). This would result in an update of the inferred users' models (using any technique of choice), leading to improved recommendations in future interactions with the system.

Chapter 5

Conclusions & Future Work

In this thesis, we designed two personalized recommender systems for the complex domain of tourism. Specifically, we introduced a recommender system that employs Bayesian learning in order to model the preferences of a targeted-tourist, while our second approach applies the well-known Kalman Filter algorithm in order to "locate the position" of user's model in a multidimensional space of travel-related features. Moreover, we provide a novel light-weight picture-based elicitation process that enables our recommenders to exploit data regarding the interests of a visitor in order to generate an efficient model that describes accurately the preferences of a specific user. Then, we equipped our systems with prior knowledge exploited via questionnaires from real-world tourists, and studied how this extra information affects the effectiveness of our approaches.

On top of that, inspired from social choice theory we designed a novel recommendation mechanism that provides diverse personalized recommendations since one can apply any multiwinner election mechanism in order to produce a set of recommendations, i.e., a committee, that satisfies any desired property. We also extend our recommender system in order to tackle the challenging problem of group recommendations in such domain. In particular, our approach utilizes a number of well-known preference aggregation mechanisms alongside a multiwinner voting rule, namely the *Reweighted Approval Voting (RAV)* mechanism. To the best of our knowledge, such mechanism has been employed for the first time in a real-world system in this work.

Finally, we conducted a systematic experimental evaluation of our approach using a real-world dataset that contains various POIs that are located in the city of Agios Nikolaos, Crete. Our results: *(i)* confirm that our approaches are able to provide effective personalized recommendations; *(ii)* show that prior knowledge regarding the general preferences of a specific category of visitors, i.e., visitors that belong to same age-group, enables our systems to face efficiently the cold-start problem; *(iii)* indicate that the Bayesian learning approach produces more accurate recommendations compared to the Kalman Filter approach when no prior knowledge is available, while Kalman Filter approach outperforms the Bayesian learning approach when high quality prior knowledge is available to us; *(iv)* demonstrate that the use of multiwinner mechanisms for personalized recommendations allows for diverse recommendations in terms of travel-related attributes, which are nevertheless of high quality and clearly outperform the simple "greedy" approach when information provided by the users is limited (and thus our confidence in the inferred user model is low); and *(v)* the employment of the RAV multiwinner election mechanism results to *fair* group recommendations with respect to the fairness metrics of *m-PROPORTIONALITY* and *m-ENVY-FREENESS* derived from the recommender system literature.

In future work, we intend to further evaluate our approaches in scenarios in which different types of prior knowledge is available—i.e., when we have and can exploit information regarding the general preferences of a type of visitors not only based on the age group that they belong to, but also their cultural background, their gender, etc. We also plan to employ other multi-winner election mechanisms, as well as equip our systems with an additional *negotiations module* that helps the members of the group to decide fairly among the recommended POIs. Moreover, one can investigate the potential of providing theoretical guarantees regarding the quality of the final recommendations that a targeted-user will receive based on the sampling method that is employed in the Bayesian learning approach. Another interesting line of work would be the employment of other multiwinner election mechanisms that provide more theoretical guarantees derived from the social choice literature, e.g., *Justified Representation (JR)* and *Extended Justified Representation (EJR)*, in order to *(i)* evaluate them with respect to fairness metrics from other domains; *(ii)* find a brake-down point in terms of computational complexity. It is also worth studying the performance of multiwinner mechanisms in the group recommendation problem for groups that share similar and dissimilar interests.

We also plan to study via experimental evaluation the effect of a *feedback component* (that can easily be inserted in the architecture's pipeline), where tourists would be able to provide a rating of their experience from a recommended POI, resulting to an update of recommender's belief regarding the interests of this visitor.[1]. Specifically, once a visitor has completed her visit of a specific POI, she would be able to provide a feedback capturing how good was this recommendation with respect to her preferences. The effective exploitation of such information could benefit our system to boost its performance by providing personalized recommendations that increase even more the satisfaction of a specific tourist. Moreover, the development of a Hybrid recommender system that combines the techniques of the proposed approaches is a promising direction of work in order to combine the beneficial properties of each technique in a single system. For example, one could employ the Bayesian learning approach during the elicitation process (when no available information is available to us) in order to construct a representative user model and utilize the Kalman Filter approach when a user provides feedback regarding the POIs that she selected to visit. Additionally, the Kalman Filter approach can be utilized in order to handle scenarios where user's preferences change over time, i.e., mood alterations, by inserting noise in the prediction sub-step of the algorithm. Furthermore, our approaches could also exploit aspects such as weather conditions in order to produce the final recommendations, i.e., to adjust its recommendations based on some context. Finally, we plan to extensively test our recommendations approaches with actual tourists, via employing our real-world mobile application for short-term visits planning to this purpose; and try these ideas in other related domains of practical interest, such as road-trips planning [8].

[1]Such a component already exists in the real-world mobile applications in which our algorithms are incorporated.

Appendix A

Publications

This master resulted in the following three (3) publications [77, 78, 93] in peer-reviewed conferences:

- Errikos Streviniotis and Georgios Chalkiadakis: *Multiwinner Election Mechanisms for Diverse Personalized Bayesian Recommendations for the Tourism Domain*. In Proceedings of RecSys Workshop on Recommenders in Tourism (RecTour 2022), co-located with the 16th ACM Conference on Recommender Systems, Seattle WA and online, September 2022. [77]
- Errikos Streviniotis and Georgios Chalkiadakis: *Preference Aggregation Mechanisms for a Tourism-Oriented Bayesian Recommender*. In Proceedings of the 24th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA-2022), Valencia, Spain, November 2022. [78]
- Ioannis Panagiotis Ziogas, Errikos Streviniotis, Harris Papadakis, and Georgios Chalkiadakis: *Content-Based Recommendations Using Similarity Distance Measures with Application in the Tourism Domain*. In Proceedings of the 12th EETN Conference on Artificial Intelligence (SETN 2022), Corfu, Greece, September 2022. [93]

Appendix B

Questionnaires

As mentioned earlier, for the needs of this research, questionnaires were created and distributed to real-world visitors of the city of Agios Nikolaos during the summer season of 2021.

The main purpose of the questionnaires was to collect information about the preferences of tourists who choose the city of Agios Nikolaos, i.e., activities that they did during their visit, approximate expenditures on purchases of products and services. etc.

We highlight that the design of these questionnaires was completed by the cooperation of Technical University of Crete and Hellenic Mediterranean University.

Moreover, we attach the exact form of the specific questionnaires that were used for this research. Note that an online version of these questionnaires was used and can be found in the folowing link:

https://docs.google.com/forms/d/e/1FAIpQLSdoXzq3QXF67akkWTOKB_4wa1M32pw8ccBcRe3eAYZb0BkpbA/viewform

Agios Nikolaos Tourist Holiday Activities Survey

We at the Municipality of Agios Nikolaos strive to upgrade our hospitality services. Help us succeed in our mission by taking this survey. Thank you in advance.

1. Please state your nationality: ☐

2. What is your gender? Male ☐ Female ☐ Prefer not to say ☐

3. Would you like to tell us what age group you belong to?
   18-25 ☐ 26-35 ☐ 36-45 ☐ 46-55 ☐ 56 -67 ☐ 67+ ☐

4. Is Agios Nikolaos…
   …your final holiday destination, as part of an organized tour? ☐
   …a port that is part of a cruise? ☐
   …a place you chose for a daily excursion while vacationing elsewhere? ☐
   Other (Please specify): ☐

5. Are you travelling:
   Alone ☐ As a couple ☐ With friends ☐ As a family (at least one child) ☐
   With a group ☐

6. Please rate (1-5, with 5=most important) the following criteria describing your ideal vacation:

Culture		Luxury accommodation and leisure	
Sun & Sea		Nightlife	
History / Archaeology		Gastronomy / Cuisine	
Adventure / Sports		Sea Sports	
Affordable prices		General shopping	
Family-friendly activities & facilities		Shopping for local products	
Rural tourism		Hiking / Trekking	

Other (Please specify): ☐

7. Duration of stay in Agios Nikolaos:
   1 day ☐ 2 days ☐ 3-5 days ☐ 6-9 days ☐ 9 days or more ☐

8. Activities that you chose/will choose during your stay in Agios Nikolaos:

Visits to archaeological sites, museums, churches	☐	Visits to beaches	☐
Visits to Restaurants, Bars, Clubs	☐	Visit to Spinalonga	☐
Daily tours	☐	Shopping (souvenirs, clothes, traditional local products, etc.)	☐
Extreme Sports	☐	Relaxing by the lake	☐

Other (Please specify): ☐

9. Approximately how much money did you/will you spend on purchases of products and services during your stay in Agios Nikolaos (<u>per day per person</u>)?

   Up to €50 ☐ €51-100 ☐ €101-150 ☐ €151-200 ☐ €200 or more ☐

10. How much money did you/will you spend per category of expenses (<u>per day per person)</u> ?

    (A) Up to €10, **(B)** €11-€20, **(C)** €21-€50, **(D)** €51-€100, **(E)** €101+

Food & Beverages		General Shopping	
Nightlife		Local Products	
Tours & Driving		Culture and Archaeological Sites	
Accommodation			

Appendix C

Mobile Application Prototype

Our Bayesian learning algorithm is incorporated in a real-world mobile tour-planning application for the city of Agios Nikolaos in Crete along with other recommendation techniques, namely a *Collaborative Filtering* approach and a *Content-based* approach. We mention that this application, called "ViP: Visit Planner: Integrated Information and Tour Planning Service for Cruise Tourism based on Hybrid Recommender Systems" has been co-financed by the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH -CREATE-INNOVATE B cycle (project code: T2EDK-03135).[1] Finally, the application and user-interface was developed by NET-MECHANICS e-commerce company. In this chapter we present some screenshots of the application for mobile devices.

Figures C.1 and C.2 depict the elicitation process for the Bayesian recommender system. In particular, Figure C.1 illustrates the process of selecting the most satisfying generic image with respect to the tourist's preferences; while Figure C.2 presents the procedure of inserting a score for the selected generic image.

Figure C.3 shows the final recommendations of our recommender to a targeted-user. Note that in the up-right side of each picture the "matching percentage" of each POI is presented to the user based on his or her interests. Figure C.4 demonstrates the personalized route and the exact position of each POI in a map of Agios Nikolaos city.

Finally, we present in Figure C.5 screenshots regarding the *(i)* welcome page of the application; *(ii)* the registration of a new user; *(iii)* the settings menu for the personalized route; and *(iv)* the menu of the final recommendations for the personalized route.

[1] Specifically, the Collaborative Filtering approach was developed by the Hellenic Mediterranean University, while the Bayesian and Content-based recommenders were implemented by the Technical University of Crete.

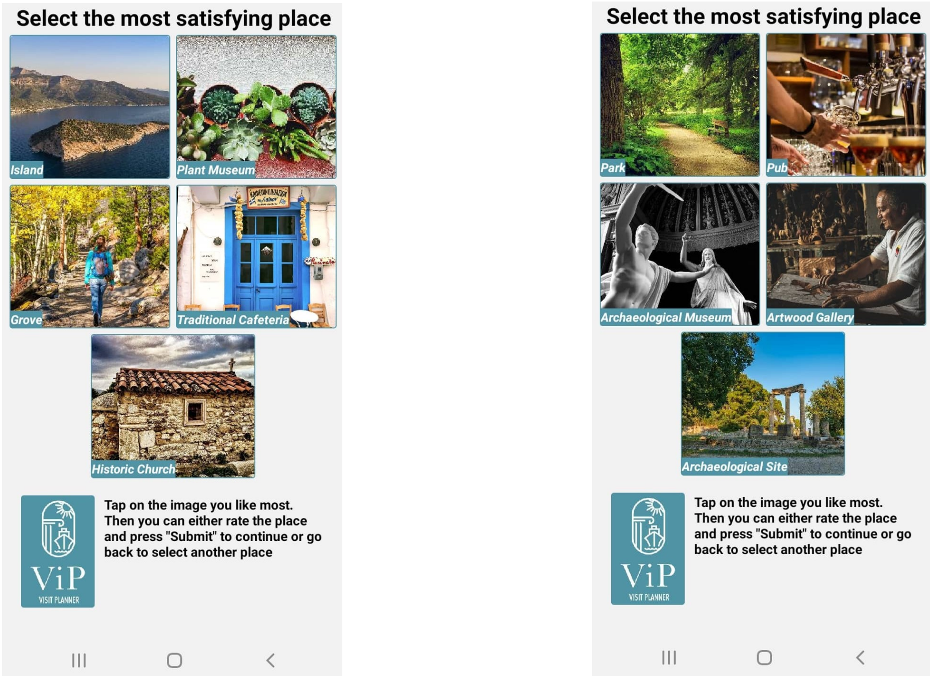

(A) (B)

FIGURE C.1: Screenshots: Selection of the most satisfying generic image.

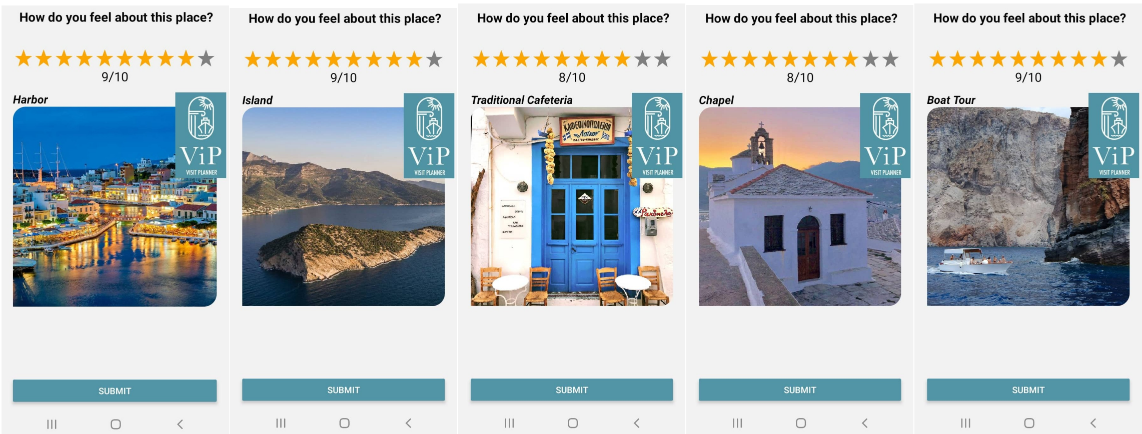

(A) (B) (C) (D) (E)

FIGURE C.2: Screenshots: Rate of the most satisfying generic image.

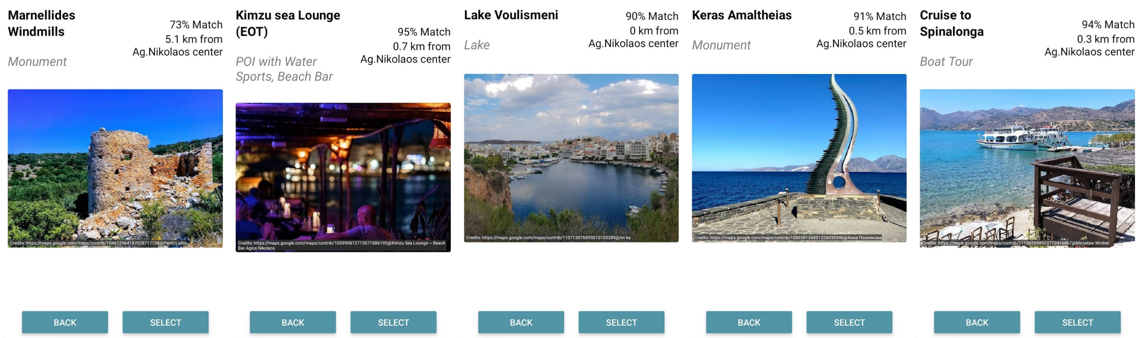

(A) (B) (C) (D) (E)

FIGURE C.3: Screenshots: Final recommendations to a specific user.

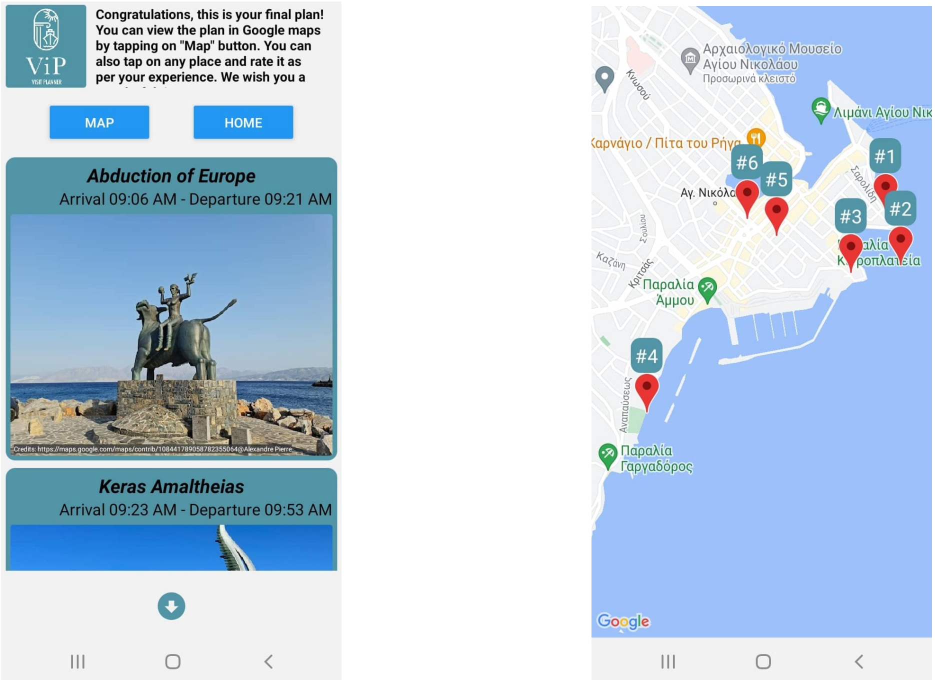

(A) (B)

FIGURE C.4: Screenshots: Personalized route and a map with the positions of each POI belonging to the route.

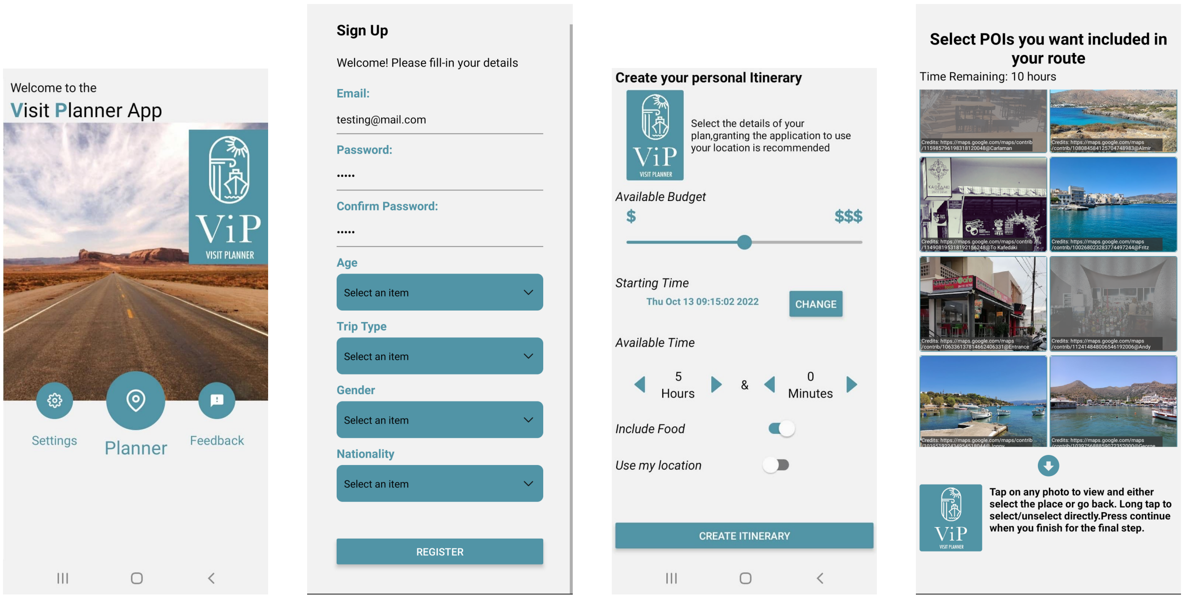

(A) (B) (C) (D)

FIGURE C.5: Screenshots of *(A)* welcome page of the application; *(B)* the registration of a new user; *(C)* the settings menu for the personalized route; and *(D)* the menu of the final recommendations for the personalized route.

Bibliography

[1] Gediminas Adomavicius and Alexander Tuzhilin. “Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions”. In: *IEEE transactions on knowledge and data engineering* 17.6 (2005), pp. 734–749.

[2] Mohammed Fadhel Aljunid and Manjaiah Dh. “An Efficient Deep Learning Approach for Collaborative Filtering Recommender System”. In: *Procedia Computer Science* 171 (2020). Third International Conference on Computing and Network Communications (CoCoNet’19), pp. 829–836. ISSN: 1877-0509. DOI: https://doi.org/10.1016/j.procs.2020.04.090. URL: https://www.sciencedirect.com/science/article/pii/S1877050920310619.

[3] H Aziz, M Brill, V Conitzer, E Elkind, R Freeman, and T Walsh. “Justified representation in approval-based committee voting”. In: *Social Choice and Welfare* 48.2 (2017), pp. 461–485.

[4] Haris Aziz, Felix Brandt, Edith Elkind, and Piotr Skowron. “Computational Social Choice: The First Ten Years and Beyond”. In: *Computing and Software Science: State of the Art and Perspectives*. Ed. by Bernhard Steffen and Gerhard Woeginger. Cham: Springer International Publishing, 2019, pp. 48–65. ISBN: 978-3-319-91908-9. DOI: 10.1007/978-3-319-91908-9_4. URL: https://doi.org/10.1007/978-3-319-91908-9_4.

[5] Haris Aziz, Serge Gaspers, Joachim Gudmundsson, Simon Mackenzie, Nicholas Mattei, and Toby Walsh. “Computational Aspects of Multi-Winner Approval Voting”. In: *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. AAMAS ’15. Istanbul, Turkey: International Foundation for Autonomous Agents and Multiagent Systems, 2015, pp. 107–115. ISBN: 9781450334136.

[6] Konstantinos Babas, Georgios Chalkiadakis, and Evangelos Tripolitakis. “You Are What You Consume: A Bayesian Method for Personalized Recommendations”. In: *Proc. of the 7th ACM Conf. on Recommender Systems*. RecSys ’13. Hong Kong, China: ACM, 2013, pp. 221–228. ISBN: 9781450324090.

[7] Linas Baltrunas, Tadas Makcinskas, and Francesco Ricci. “Group recommendations with rank aggregation and collaborative filtering”. In: *RecSys ’10*. 2010.

[8] Susanne Becken and Jude Wilson. “Trip Planning and Decision Making of Self-Drive Tourists”. In: *Journal of Travel and Tourism Marketing* 20 (Jan. 2006), pp. 47–62. DOI: 10.1300/J073v20n03_04.

[9] J. Bobadilla, F. Ortega, A. Hernando, and A. Gutiérrez. "Recommender systems survey". In: *Knowledge-Based Systems* 46 (2013), pp. 109–132. ISSN: 0950-7051. DOI: https://doi.org/10.1016/j.knosys.2013.03.012. URL: https://www.sciencedirect.com/science/article/pii/S0950705113001044.

[10] Joan Borràs, Antonio Moreno, and Aida Valls. "Intelligent tourism recommender systems: A survey". In: *Expert Systems with Applications* 41.16 (2014), pp. 7370–7389. ISSN: 0957-4174.

[11] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. *Handbook of Computational Social Choice*. 1st. USA: Cambridge University Press, 2016. ISBN: 1107060435.

[12] Robin Burke. "Hybrid Recommender Systems: Survey and Experiments". In: *User Modeling and User-Adapted Interaction* 12 (Nov. 2002). DOI: 10.1023/A:1021240730564.

[13] David Carmel and Shaul Markovitch. "Exploration Strategies for Model-based Learning in Multiagent Systems". In: *Aut. Agents and Multi-agent Systems* 2.2 (1999), pp. 141–172.

[14] Kinjal Chaudhari and Ankit Thakkar. "A comprehensive survey on travel recommender systems". In: *Archives of Computational Methods in Engineering* 27.5 (2020), pp. 1545–1571.

[15] Li Chen and Weike Pan. "CoFiSet: Collaborative Filtering via Learning Pairwise Preferences over Item-sets". In: *SDM*. 2013.

[16] Yan-Ying Chen, An-Jung Cheng, and Winston H. Hsu. "Travel Recommendation by Mining People Attributes and Travel Group Types From Community-Contributed Photos". In: *IEEE Transactions on Multimedia* 15.6 (2013), pp. 1283–1295.

[17] David R Cox. "The regression analysis of binary sequences". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 20.2 (1958), pp. 215–232.

[18] Amine Dadoun, Raphael Troncy, Michael Defoin-Platel, and Gerardo Ayala Solano. "Predicting your next trip: A knowledge graph-based multi-task learning approach for travel destination recommendation". In: *2021 Workshop on Recommenders in Tourism, RecTour 2021* 2974 (2021).

[19] Sriharsha Dara, C. Ravindranath Chowdary, and Ch. Naveen Kumar. "A survey on group recommender systems". In: *Journal of Intelligent Information Systems* 54 (2019), pp. 271–295.

[20] Berardina De Carolis, Nicole Novielli, Vito Leonardo Plantamura, and Enrica Gentile. "Generating Comparative Descriptions of Places of Interest in the Tourism Domain". In: *Proceedings of the Third ACM Conference on Recommender Systems*. RecSys '09. New York, New York, USA: Association for Computing Machinery, 2009, pp. 277–280. ISBN: 9781605584355. DOI: 10.1145/1639714.1639767. URL: https://doi.org/10.1145/1639714.1639767.

[21] Amra Delic and Julia Neidhardt. "A Comprehensive Approach to Group Recommendations in the Travel and Tourism Domain". In: *Proc. of the 25th Conf. on User Modeling, Adaptation and Personalization*. UMAP '17. Bratislava, Slovakia: ACM, 2017, pp. 11–16. ISBN: 9781450350679.

[22] Michael D. Ekstrand, John T. Riedl, and Joseph A. Konstan. "Collaborative Filtering Recommender Systems". In: *Found. Trends Hum.-Comput. Interact.* 4.2 (Feb. 2011), pp. 81–173. ISSN: 1551-3955. DOI: 10.1561/1100000009. URL: https://doi.org/10.1561/1100000009.

[23] Mehdi Elahi, Matthias Braunhofer, Tural Gurbanov, and Francesco Ricci. "User Preference Elicitation, Rating Sparsity and Cold Start". In: *Collaborative Recommendations*. Chap. Chapter 8, pp. 253–294. DOI: 10.1142/9789813275355_0008. eprint: https://www.worldscientific.com/doi/pdf/10.1142/9789813275355_0008. URL: https://www.worldscientific.com/doi/abs/10.1142/9789813275355_0008.

[24] E Elkind, P Faliszewski, P Skowron, and A Slinko. "Properties of multiwinner voting rules". In: *Social Choice and Welfare* 48.3 (2017), pp. 599–632.

[25] Edith Elkind, Piotr Faliszewski, Jean-Francois Laslier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. "What Do Multiwinner Voting Rules Do? An Experiment over the Two-Dimensional Euclidean Domain". In: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*. AAAI'17. San Francisco, California, USA: AAAI Press, 2017, pp. 494–501.

[26] Ulle Endriss. *Trends in Computational Social Choice*. Lulu.com, 2017. ISBN: 1326912097.

[27] Piotr Faliszewski, Arkadii M. Slinko, and Nimrod Talmon. "Multiwinner Voting: A New Challenge for Social Choice Theory". In: 2017.

[28] Damianos Gavalas, Vlasios Kasapakis, Charalampos Konstantopoulos, Grammati Pantziou, Nikolaos Vathis, and Christos Zaroliagis. "The eCOMPASS multimodal tourist tour planner". In: *Expert systems with Applications* 42.21 (2015), pp. 7303–7316.

[29] Damianos Gavalas, Charalampos Konstantopoulos, Konstantinos Mastakas, and Grammati Pantziou. "Mobile Recommender Systems in Tourism". In: *Journal of Network and Computer Applications* 39 (Mar. 2014), pp. 319–333.

[30] Grzegorz Gawron and Piotr Faliszewski. "Using Multiwinner Voting to Search for Movies". In: *EUMAS*. 2022.

[31] Jagadeesh Gorla, Neal Lathia, Stephen Robertson, and Jun Wang. "Probabilistic Group Recommendation via Information Matching". In: *Proc. of the 22nd International Conference on World Wide Web*. WWW '13. Rio de Janeiro, Brazil: Association for Computing Machinery, 2013, pp. 495–504. ISBN: 9781450320351.

[32] Shengbo Guo. "Bayesian recommender systems : models and algorithms". In: 2011.

[33] Ruining He and Julian McAuley. "VBPR: Visual Bayesian Personalized Ranking from Implicit Feedback". In: *Proceedings of the AAAI Conference on Artificial Intelligence* 30 (Oct. 2015). DOI: 10.1609/aaai.v30i1.9973.

[34] F.O. Isinkaye, Y.O. Folajimi, and B.A. Ojokoh. "Recommendation systems: Principles, methods and evaluation". In: *Egyptian Informatics Journal* 16.3 (2015), pp. 261–273. ISSN: 1110-8665. DOI: https://doi.org/10.1016/j.eij.2015.06.005. URL: https://www.sciencedirect.com/science/article/pii/S1110866515000341.

[35] Anthony Jameson and Barry Smyth. "Recommendation to Groups". In: *The Adaptive Web: Methods and Strategies of Web Personalization*. Ed. by Peter Brusilovsky, Alfred Kobsa, and Wolfgang Nejdl. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 596–627. ISBN: 978-3-540-72079-9.

[36] Simon J. Julier and Jeffrey K. Uhlmann. "New extension of the Kalman filter to nonlinear systems". In: *Defense, Security, and Sensing*. 1997.

[37] Rudolph Emil Kalman. "A New Approach to Linear Filtering and Prediction Problems". In: *Journal of Basic Engineering* 82.1 (1960), p. 35. DOI: 10.1115/1.3662552. URL: http://dx.doi.org/10.1115/1.3662552.

[38] Ondrej Kaššák, Michal Kompan, and Mária Bieliková. "Personalized hybrid recommendation for group of users: Top-N multimedia recommender". In: *Information Processing & Management* 52.3 (2016), pp. 459–477.

[39] Mesut Kaya, Derek Bridge, and Nava Tintarev. "Ensuring Fairness in Group Recommendations by Rank-Sensitive Balancing of Relevance". In: *Fourteenth ACM Conference on Recommender Systems*. NY, USA: ACM, 2020, pp. 101–110. ISBN: 9781450375832.

[40] Mohamed Elyes Ben Haj Kbaier, Hela Masri, and Saoussen Krichen. "A Personalized Hybrid Tourism Recommender System". In: *2017 IEEE/ACS 14th International Conference on Computer Systems and Applications (AICCSA)*. 2017, pp. 244–250.

[41] Blerina Lika, Kostas Kolomvatsos, and Stathes Hadjiefthymiades. "Facing the cold start problem in recommender systems". In: *Expert Systems with Applications* 41.4, Part 2 (2014), pp. 2065–2073. ISSN: 0957-4174. DOI: https://doi.org/10.1016/j.eswa.2013.09.005. URL: https://www.sciencedirect.com/science/article/pii/S0957417413007240.

[42] Kwan Hui Lim, Jeffrey Chan, Christopher Leckie, and Shanika Karunasekera. "Personalized Tour Recommendation Based on User Interests and Points of Interest Visit Durations". In: *Proceedings of the 24th International Conference on Artificial Intelligence*. IJCAI'15. Buenos Aires, Argentina: AAAI Press, 2015, pp. 1778–1784. ISBN: 9781577357384.

[43] Siping Liu, Xiaohan Tu, and Renfa Li. "Unifying explicit and implicit feedback for top-N recommendation". In: *2017 IEEE 2nd International Conference on Big Data Analysis (ICBDA)*. 2017, pp. 35–39. DOI: 10.1109/ICBDA.2017.8078860.

[44] Pasquale Lops, Dietmar Jannach, Cataldo Musto, Toine Bogers, and Marijn Koolen. "Trends in Content-Based Recommendation". In: *User Modeling and User-Adapted Interaction* 29.2 (Apr. 2019), pp. 239–249. ISSN: 0924-1868. DOI: 10.1007/s11257-019-09231-w. URL: https://doi.org/10.1007/s11257-019-09231-w.

[45] David Massimo and Francesco Ricci. "Building effective recommender systems for tourists". In: *AI Magazine* 43.2 (2022), pp. 209–224. DOI: https://doi.org/10.1002/aaai.12057. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/aaai.12057. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/aaai.12057.

[46] David Massimo and Francesco Ricci. "Harnessing a Generalised User Behaviour Model for Next-POI Recommendation". In: *Proceedings of the 12th ACM Conference on Recommender Systems*. RecSys '18. Vancouver, British Columbia, Canada: Association for Computing Machinery, 2018, pp. 402–406. ISBN: 9781450359016. DOI: 10.1145/3240323.3240392. URL: https://doi.org/10.1145/3240323.3240392.

[47] David Massimo and Francesco Ricci. "Next-POI Recommendations Matching User's Visit Behaviour". In: *Information and Communication Technologies in Tourism 2021*. Ed. by Wolfgang Wörndl, Chulmo Koo, and Jason L. Stienmetz. Cham: Springer International Publishing, 2021, pp. 45–57. ISBN: 978-3-030-65785-7.

[48] Judith Masthoff. "Group Recommender Systems: Combining Individual Models". In: Jan. 2011, pp. 677–702.

[49] Joseph F. McCarthy. "Pocket RestaurantFinder: A situated recommender system for groups". In: *Proc. of the Workshop on Mobile Ad-Hoc Communication at the 2002 ACM Conference on Human Factors in Computer Systems*. Minneapolis: ACM, 2002.

[50] Joseph F. McCarthy and Theodore D. Anagnost. *MUSICFX: An Arbiter of Group Preferences for Computer Supported Collaborative Workouts*. 1998.

[51] Kevin McCarthy, Maria Salamó, Lorcan Coyle, Lorraine McGinty, Barry Smyth, and Paddy Nixon. "Group Recommender Systems: A Critiquing Based Approach". In: *Proc. of the 11th Int. Conf. on Intelligent User Interfaces*. IUI '06. Sydney, Australia: ACM, 2006, pp. 267–269. ISBN: 1595932879.

[52] Douglas R. Miller. "Markov processes". In: *Encyclopedia of Operations Research and Management Science*. Ed. by Saul I. Gass and Carl M. Harris. New York, NY: Springer US, 2001, pp. 486–490. ISBN: 978-1-4020-0611-1. DOI: 10.1007/1-4020-0611-X_582. URL: https://doi.org/10.1007/1-4020-0611-X_582.

[53] Ram Krishn Mishra, Siddhaling Urolagin, and J. Angel Arul Jothi. "Sentiment Analysis for POI Recommender Systems". In: *2020 Seventh International Conference on Information Technology Trends (ITT)*. 2020, pp. 174–179. DOI: 10.1109/ITT51279.2020.9320885.

[54] T.K. Moon. "The expectation-maximization algorithm". In: *IEEE Signal Processing Magazine* 13.6 (1996), pp. 47–60. DOI: 10.1109/79.543975.

[55] Julia Neidhardt, Rainer Schuster, Leonhard Seyfang, and Hannes Werthner. "Eliciting the Users' Unknown Preferences". In: *Proc. of the 8th ACM Conf. on Recommender Systems*. RecSys '14. Foster City, Silicon Valley, California, USA: ACM, 2014, pp. 309–312. ISBN: 9781450326681.

[56] Julia Neidhardt, Leonhard Seyfang, Rainer Schuster, and Hannes Werthner. "A picture-based approach to recommender systems". In: *Information Technology & Tourism* 15 (2015), pp. 49–69.

[57] Frank Nielsen and Richard Nock. "Clustering Multivariate Normal Distributions". In: *Emerging Trends in Visual Computing: LIX Fall Colloquium, ETVC 2008, Palaiseau, France, November 18-20, 2008. Revised Invited Papers*. Ed. by Frank Nielsen. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 164–174. ISBN: 978-3-642-00826-9.

[58] Weike Pan and Li Chen. "Group Bayesian Personalized Ranking with Rich Interactions for One-Class Collaborative Filtering". In: *Neurocomputing* 207 (May 2016). DOI: 10.1016/j.neucom.2016.05.019.

[59] Costas Panagiotakis and Harris Papadakis. "Session-Based Recommendation by Combining Probabilistic Models and LSTM". In: *Proceedings of the Recommender Systems Challenge 2022*. RecSysChallenge '22. Seattle, WA, USA: Association for Computing Machinery, 2022, pp. 39–44. ISBN: 9781450398565. DOI: 10.1145/3556702.3556846. URL: https://doi.org/10.1145/3556702.3556846.

[60] Moon-Hee Park, Han-Saem Park, and Sung-Bae Cho. "Restaurant Recommendation for Group of People in Mobile Environments Using Probabilistic Multicriteria Decision Making". In: July 2008, pp. 114–122. ISBN: 978-3-540-70584-0.

[61] Dip Paul and Subhradeep Kundu. "A Survey of Music Recommendation Systems with a Proposed Music Recommendation System". In: *Emerging Technology in Modelling and Graphics*. Ed. by Jyotsna Kumar Mandal and Debika Bhattacharya. Singapore: Springer Singapore, 2020, pp. 279–285. ISBN: 978-981-13-7403-6.

[62] Michael J. Pazzani and Daniel Billsus. "Content-Based Recommendation Systems". In: *The Adaptive Web: Methods and Strategies of Web Personalization*. Ed. by Peter Brusilovsky, Alfred Kobsa, and Wolfgang Nejdl. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 325–341. ISBN: 978-3-540-72079-9. DOI: 10.1007/978-3-540-72079-9_10. URL: https://doi.org/10.1007/978-3-540-72079-9_10.

[63] Huihuai Qiu, Yun Liu, Guibing Guo, Zhu Sun, Jie Zhang, and Hai Thanh Nguyen. "BPRH: Bayesian personalized ranking for heterogeneous implicit feedback". In: *Information Sciences* 453 (2018), pp. 80–98. ISSN: 0020-0255. DOI: https://doi.org/10.1016/j.ins.2018.04.027. URL: https://www.sciencedirect.com/science/article/pii/S0020025516315742.

[64] Lawrence R Rabiner. "A tutorial on hidden Markov models and selected applications in speech recognition". In: *Proceedings of the IEEE* 77.2 (1989), pp. 257–286.

[65] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. "BPR: Bayesian Personalized Ranking from Implicit Feedback". In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. UAI '09. Montreal, Quebec, Canada: AUAI Press, 2009, pp. 452–461. ISBN: 9780974903958.

[66] Francesco Ricci, Lior Rokach, and Bracha Shapira. "Recommender Systems Handbook". In: vol. 1-35. Oct. 2010, pp. 1–35. DOI: 10.1007/978-0-387-85820-3_1.

[67] Maria Salamó, Kevin McCarthy, and Barry Smyth. "Generating recommendations for consensus negotiation in group personalization services". In: *Personal and Ubiquitous Computing* 16 (2011), pp. 597–610.

[68] Pablo Sanchez and Alejandro Bellogin. "Point-of-Interest Recommender Systems Based on Location-Based Social Networks: A Survey from an Experimental Perspective". In: *ACM Comput. Surv.* (2022). Just Accepted. ISSN: 0360-0300. DOI: 10.1145/3510409. URL: https://doi.org/10.1145/3510409.

[69] Lara Quijano Sánchez, Juan A. Recio-García, Belén Díaz-Agudo, and Guillermo Jiménez-Díaz. "Happy Movie: A Group Recommender Application in Facebook". In: *FLAIRS Conference*. 2011.

[70] Joy Sarkar, Abhishek Majumder, Chhabi Panigrahi, Sudipta Roy, and Bibudhendu Pati. "Tourism recommendation system: a survey and future research directions". In: *Multimedia Tools and Applications* (Apr. 2022), pp. 1–45. DOI: 10.1007/s11042-022-12167-w.

[71] J. Ben Schafer, Dan Frankowski, Jon Herlocker, and Shilad Sen. "Collaborative Filtering Recommender Systems". In: *The Adaptive Web: Methods and Strategies of Web Personalization*. Ed. by Peter Brusilovsky, Alfred Kobsa, and Wolfgang Nejdl. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 291–324. ISBN: 978-3-540-72079-9. DOI: 10.1007/978-3-540-72079-9_9. URL: https://doi.org/10.1007/978-3-540-72079-9_9.

[72] Josef Schmee and T. Anderson. "An Introduction to Multivariate Statistical Analysis". In: *Technometrics* 28 (1986), p. 180.

[73] Dimitris Serbos, Shuyao Qi, Nikos Mamoulis, Evaggelia Pitoura, and Panayiotis Tsaparas. "Fairness in Package-to-Group Recommendations". In: WWW '17. Perth, Australia, 2017, pp. 371–379. ISBN: 9781450349130.

[74] Mete Sertkan, Julia Neidhardt, and Hannes Werthner. "PicTouRe - A Picture-Based Tourism Recommender". In: *Fourteenth ACM Conference on Recommender Systems*. RecSys '20. Virtual Event, Brazil: Association for Computing Machinery, 2020, pp. 597–599. ISBN: 9781450375832. DOI: 10.1145/3383313.3411526. URL: https://doi.org/10.1145/3383313.3411526.

[75] Piotr Skowron, Piotr Faliszewski, and Jérôme Lang. "Finding a collective set of items: From proportional multirepresentation to group recommendation". In: *Artificial Intelligence* 241 (2016), pp. 191–216. ISSN: 0004-3702.

[76] M Sridevi, N ManikyaArun, M Sheshikala, and E Sudarshan. "Personalized fashion recommender system with image based neural networks". In: *IOP Conference Series: Materials Science and Engineering* 981.2 (Dec. 2020), p. 022073. DOI: 10.1088/1757-899X/981/2/022073. URL: https://dx.doi.org/10.1088/1757-899X/981/2/022073.

[77] Errikos Streviniotis and Georgios Chalkiadakis. "Multiwinner Election Mechanisms for Diverse Personalized Bayesian Recommendations for the Tourism Domain". In: *Proceedings of the Workshop on Recommenders in Tourism (RecTour 2022) co-located with the 16th ACM Conference on Recommender Systems (RecSys 2022), Seattle, WA, USA and Online, September 22, 2022*. Ed. by Julia Neidhardt, Wolfgang Wörndl, Tsvi Kuflik, Dmitri Goldenberg, and Markus Zanker. Vol. 3219. CEUR Workshop Proceedings. CEUR-WS.org, 2022, pp. 65–82. URL: http://ceur-ws.org/Vol-3219/paper5.pdf.

[78] Errikos Streviniotis and Georgios Chalkiadakis. "Preference Aggregation Mechanisms for a Tourism-Oriented Bayesian Recommender". In: *PRIMA 2022: Principles and Practice of Multi-Agent Systems - 24th International Conference, Valencia, Spain, November 16-18, 2022, Proceedings*. Ed. by Reyhan Aydogan, Natalia Criado, Jérôme Lang, Víctor Sánchez-Anguix, and Marc Serramia. Vol. 13753. Lecture Notes in Computer Science. Springer, 2022, pp. 331–346. DOI: 10.1007/978-3-031-21203-1_20. URL: https://doi.org/10.1007/978-3-031-21203-1%5C_20.

[79] Jianing Sun, Wei Guo, Dengcheng Zhang, Yingxue Zhang, Florence Regol, Yaochen Hu, Huifeng Guo, Ruiming Tang, Han Yuan, Xiuqiang He, and Mark Coates. "A Framework for Recommending Accurate and Diverse Items Using Bayesian Graph Convolutional Neural Networks". In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD '20. Virtual Event, CA, USA: Association for Computing Machinery, 2020, pp. 2030–2039. ISBN: 9781450379984. DOI: 10.1145/3394486.3403254. URL: https://doi.org/10.1145/3394486.3403254.

[80] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 1998. ISBN: 0-262-19398-1. URL: http://www.cs.ualberta.ca/%7Esutton/book/ebook/the-book.html.

[81] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005. ISBN: 0262201623.

[82] Chuhan Wu, Fangzhao Wu, Yongfeng Huang, and Xing Xie. "Personalized News Recommendation: Methods and Challenges". In: *ACM Transactions on Information Systems (TOIS)* (2022).

[83] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. "Graph Neural Networks in Recommender Systems: A Survey". In: *ACM Comput. Surv.* (Mar.

2022). Just Accepted. ISSN: 0360-0300. DOI: 10.1145/3535101. URL: https://doi.org/10.1145/3535101.

[84] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. "Graph Neural Networks in Recommender Systems: A Survey". In: *ACM Computing Surveys* (May 2022). DOI: 10.1145/3535101.

[85] Byung-Jun Yoon. "Hidden Markov Models and their Applications in Biological Sequence Analysis". In: *Current Genomics* 10 (2009), pp. 402–415.

[86] Quan Yuan, Gao Cong, and Chin-Yew Lin. "COM: A Generative Model for Group Recommendation". In: *Proc. of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '14. New York, New York, USA: Association for Computing Machinery, 2014, pp. 163–172. ISBN: 9781450329569.

[87] Jun Zeng, Feng Li, Haiyang Liu, Junhao Wen, and Sachio Hirokawa. "A restaurant recommender system based on user preference and location in mobile environment". In: *2016 5th IIAI International Congress on Advanced Applied Informatics*. IEEE. 2016, pp. 55–60.

[88] Jiang Zhang, Yufeng Wang, Zhiyuan Yuan, and Qun Jin. "Personalized real-time movie recommendation system: Practical prototype and evaluation". In: *Tsinghua Science and Technology* 25.2 (2020), pp. 180–191. DOI: 10.26599/TST.2018.9010118.

[89] Qian Zhang, Jie Lu, and Yaochu Jin. "Artificial intelligence in recommender systems". In: *Complex and Intelligent Systems* 7 (Nov. 2020). DOI: 10.1007/s40747-020-00212-w.

[90] Yi Zhang and Jonathan Koren. "Efficient Bayesian Hierarchical User Modeling for Recommendation System". In: *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. SIGIR '07. Amsterdam, The Netherlands: Association for Computing Machinery, 2007, pp. 47–54. ISBN: 9781595935977. DOI: 10.1145/1277741.1277752. URL: https://doi.org/10.1145/1277741.1277752.

[91] Yu Zheng, Chen Gao, Liang Chen, Depeng Jin, and Yong Li. "DGCN: Diversified Recommendation with Graph Convolutional Networks". In: *Proceedings of the Web Conference 2021*. WWW '21. Ljubljana, Slovenia: Association for Computing Machinery, 2021, pp. 401–412. ISBN: 9781450383127. DOI: 10.1145/3442381.3449835. URL: https://doi.org/10.1145/3442381.3449835.

[92] Tao Zhou, Zoltan Kuscsik, Jian-Guo Liu, Matus Medo, Joseph Wakeling, and Yi-Cheng Zhang. "Solving the apparent diversity-accuracy dilemma of recommender systems". In: *Proceedings of the National Academy of Sciences of the United States of America* 107 (Feb. 2010), pp. 4511–5. DOI: 10.1073/pnas.1000488107.

[93] Ioannis Panagiotis Ziogas, Errikos Streviniotis, Harris Papadakis, and Georgios Chalkiadakis. "Content-Based Recommendations Using Similarity Distance Measures with Application in the Tourism Domain". In: *Proceedings of*

the 12th Hellenic Conference on Artificial Intelligence. SETN '22. Corfu, Greece: Association for Computing Machinery, 2022. ISBN: 9781450395977. DOI: 10.1145/3549737.3549772. URL: https://doi.org/10.1145/3549737.3549772.