

Chapter 1

Introduction

In this chapter we present a general overview of the thesis. In Sections 1.1 and 1.2 we provide the motivations and research goals of our work. In Section 1.3 we summarise the main contributions of the thesis, and in Section 1.4 we list the publications resulting from our research. Finally, in Section 1.5 we describe the structure of this document.

1.1 Motivation

Information Retrieval (IR) technologies have gained outstanding prevalence in the last two decades with the explosion of massive online information repositories, and much in particular the World Wide Web. IR systems are researched and designed in ways that seek to maximise the degree of satisfaction of certain objective conditions, typically – though not necessarily only – user satisfaction. IR research and development have revolved around the definition of models and algorithms that best achieve this goal, methodologies and metrics that let assess how well the goal is achieved by different systems, and sound theories providing a solid ground and orientation in the development of IR algorithms and their consistent evaluation. Among many new trends stemming from this main stream of research and developments, a new research goal started to be considered by the early 2000's: is it possible to predict how good a result returned by an IR system is going to be, before presenting it to the user, or even, before running the IR system at all (Cronen-Townsend et al., 2002)? This question has given rise to a fertile strand of research on so-called **performance prediction** in IR.

Performance prediction has many potential uses in IR. From the user's perspective it may provide valuable feedback that can be used to direct a search, from the system's perspective it may help to distinguish poorly performing queries, and from the system administrator's perspective it may let identify queries related to a specific subject that are difficult for the search engine. Performance prediction approaches are based on the analysis and characterisation of the evidence used by an IR system to assess the relevance (utility, value, etc.) of retrieval objects (documents, goods, etc.) at execution time (Cronen-Townsend et al., 2002). The most classic and basic retrieval scenario involves a user query and a collection of documents as the basic input to form a ranked list of search results, but other additional elements can be taken into account to select and rank results (Baeza-Yates and Ribeiro-Neto, 2011). Any information the retrieval system takes as input can be taken as input for the performance prediction as well, and often the prediction methods use additional information beyond that. The user context (current tasks, query logs, preferences, etc.), global properties of the document collection, comparisons with respect to other reference elements such as historic data, and the output from other systems, among others, are some examples of the different sources of information that a predictor may draw evidence from.

Predicting the performance of a subsystem, module, function, or input by contrasting the performance estimation for a query for each component, enables an array of dynamic optimisation strategies that select at runtime the option which is predicted to work best or, when larger systems or hybrid approaches are used, allows for adjusting on the fly the participation of each module. The IR field is pervaded with

cases where information relevance, retrieval systems, models, and criteria are based on a fusion or combination of sub-models. Personalised retrieval systems (including techniques such as personalised search, recommender systems, collaborative filtering, and retrieval in context) are clear examples where performance prediction can be applied since such systems combine several sources of evidence for relevance assessment, such as explicit queries, search history, explicit user ratings, social information, user feedback, and context models.

Performance prediction finds additional motivation in personalised recommendation, inasmuch these applications may decide to produce recommendations or hold them back, delivering only the sufficiently reliable ones. Furthermore, current Recommender Systems (RS) are characterised by an increasing diversification of the types and sources of data, content, evidence and methods, available to make decisions and build their output. In such context, predicting the performance of a specific recommendation approach or component becomes an appealing problem, as it lets properly combine the available alternatives, and make the most of them by dynamically adapting the recommendation strategy to the situation at hand. The question gains increasing relevance today, with the proliferation of hybrid recommendation techniques to improve the accuracy of the methods – the Netflix prize was a paradigmatic example of the use of this, where all the top ranked participants used combinations of large sets of recommendation methods. This calls for the research of hybrid approaches with a level of dynamic self-adjustment mechanisms, in order to optimise the resulting effectiveness of the recommendation systems, by opportunistically taking advantage of high-quality data when available, but avoiding sticking to fixed strategies when they can be predicted to yield poor results under certain conditions.

Performance prediction in IR is typically assessed in terms of the correlation between a predictor's scores and a system's performance values on a per-query basis. This requires reliable performance evaluation metrics and methodologies, which have been thoroughly analysed, and are currently well established in the IR field, mostly oriented to ad-hoc search. In contrast, evaluation in the RS field is more open, and the variability in evaluation approaches and experimental configurations is significant. How to measure the performance of a recommender system is a key issue in our research since the system quality measurements may be influenced by statistical properties of the measurement approach and/or the experimental design. Throughout this thesis we shall focus on the accuracy of the system, where we have to avoid that if a metric – i.e., precision – is biased towards some form of noise along with the recommender's quality, then a predictor capturing only that noise would appear as an (equivocal) effective performance predictor. Hence, statistical biases (noises) of the evaluation methodologies should be well understood in order to enable a meaningful assessment of performance predictors.

Drawing from the state of the art on performance prediction in IR as a starting point, the present work restates the problem in the field of Recommender Systems where it has barely been addressed so far. We research meaningful definitions of performance in the context of RS, and the elements to which it can sensibly apply, investigating the statistical biases that may arise when adapting the IR evaluation framework into RS. In doing so, we take as a driving direction the application of performance prediction to achieve improvements in two specific combination problems in the RS field, namely, the dynamic combination of recommendation methods in hybrid recommendation systems, and the dynamic aggregation of neighbours' signals in user-based collaborative filtering.

1.2 Research goals

The main objective of the research presented here is to find predictive methods for the performance of specific components in recommender systems, and to improve the performance of combined recommendation methods, based on the dynamic, automatic analysis and prediction of the expected performance of the constituents of the composite methods, whereupon the relative participation of each constituent is adjusted, in accordance to its predicted effectiveness. To address these problems, this work has the following specific research objectives:

RG1: Analysis and formalisation of how retrieval performance is defined and evaluated in recommender systems. We need to develop an in-depth study on how recommender systems can be reliably evaluated in terms of numeric metric values, since we aim to predict their performance. Moreover, we have to investigate whether there is any bias on the way the systems are evaluated – either by the evaluation methodologies or metrics, since any bias in the evaluation process would lead to inconclusive or misleading results about the predictive power of the performance prediction methods proposed. If these biases do exist, we aim to precisely understand them and develop methodologies to isolate them; then, we shall check the effectiveness of the predictors against well-known baselines and whether it changes when unbiased methodologies are used.

RG2: Adaptation and definition of performance prediction techniques for recommender systems. We aim to study the potential of performance prediction in specific problems and settings in the area of Recommender Systems. We shall investigate the definition of a formal framework where performance predictors can be integrated. As a starting point, we aim to explore the adaptation of specific effective predictors from Information Retrieval such as query clarity (Cronen-Townsend et al., 2002) to recommender systems. Complementarily to the adaptation of known techniques, we aim to research the definition of new predictors based on models from Information Theory and Social Graphs, besides other heuristic, domain-specific ap-

proaches. Once we have defined some recommendation performance predictors, we shall assess the effectiveness of such predictors in terms of their correlation to performance metrics to estimate the predictive power of the performance predictors.

RG3: Application of performance predictors to hybrid and compound recommender systems. We aim to identify and integrate the proposed predictors into combined recommendation methods, in order to achieve an actual improvement in the performance of the combined methods. With this goal in mind, we shall consider problems where an aggregation of recommendation methods is needed, and shall analyse how to apply the performance predictors mentioned above in such problems. Besides, a methodological study for the experimental approach, setup, and metrics should be performed in such a way that appropriate baseline methods and experimental designs are used. Finally, we shall assess the improvements and benefits of the combined methods when the performance predictors are applied.

1.3 Contributions

This thesis is devoted to the problem of estimating the performance of recommender systems for particular users and items. The main contributions of this thesis are related to the evaluation of the performance of a recommender system, and the prediction of such performance, where we have addressed several issues regarding both topics and we have proposed novel models and methods, which have been applied into two applications as we shall see next.

As a first step, this thesis analyses the Cranfield paradigm of Information Retrieval evaluation since recommender systems are usually considered as a particular problem of information filtering, and, thus, of information retrieval at large (Belkin and Croft, 1992). In Chapter 4 **we discuss the differences involved in the experimental design alternatives from the common assumptions made in the Cranfield paradigm**, which result in **substantial statistical biases arising in Recommender Systems**, and **we propose different methods to neutralise these biases**. Additionally, the following related contributions have been addressed:

- We propose a precise and systematic characterisation of design alternatives in the adaptation of the Cranfield paradigm to recommendation tasks. We identify assumptions and conditions underlying the Cranfield paradigm that are not granted in usual recommendation experiments.
- We detect and characterise resulting statistical biases, namely test sparsity and item popularity, which do not arise in common test collections from IR, but do interfere in recommendation experiments.

- We propose two novel experimental designs in order to neutralise these biases. We observe that a percentile-based evaluation considerably reduces the margin for the popularity bias, whereas a uniform-test approach removes any statistical advantage provided by having more positive test ratings. Furthermore, we find that both approaches discriminate well between pure popularity-based recommendation and an efficient personalised recommendation algorithm.

Additionally, in this thesis **we show how query performance prediction techniques developed in Information Retrieval can be adapted to Recommender Systems, and result in effective predictors in this domain.** We present these performance predictors in Chapter 6, where we propose different adaptations of the query clarity predictor based on different interpretations of the underlying language models along with models from Information Theory and Social Graphs. Furthermore, in the same chapter **we assess the effectiveness of such predictors by measuring the correlation with respect to performance metrics**, where we also test the methods proposed in Chapter 4 to neutralise biases on evaluation. Specific contributions regarding performance prediction for recommendation are summarised as follows:

- We define and elaborate several predictive models in the Recommender Systems domain according to different formulations and assumptions, and based on three types of preference data: rating-based, log-based, and social-based.
- Formulations for rating preferences are based on adaptations of query clarity from IR and concepts from Information Theory such as entropy. In this adaptation we propose different probability estimations, where Bayesian derivations and non-parametric estimations are developed.
- We also exploit temporal features when defining log-based predictors. Specifically, we use a time-aware version of the Kullback-Leibler divergence, along with other time series concepts such as a user's autocorrelation.
- We use graph-based metrics from Graph Theory to define predictors leveraging social network structures, and correlations between topological properties of users and the success of recommendations delivered to them.
- We find strong correlations between the outputs of the predictors and the performance metrics, thus finding empirical evidence of the predictive power of the proposed approaches. Furthermore, when unbiased evaluation methodologies are used, the predictors still obtain good correlation values, evidencing that our proposed predictors are not just capturing and benefitting from the analysed biases, especially when we compare them against other trivial predictors.

Finally, Chapters 7 and 8 present two applications of performance predictors on Recommender Systems. In Chapter 7 **we propose several linearly weighted hy-**

brids where the weights are dynamically adjusted based on the predictors' output. We observe that the correlations obtained in Chapter 6 help decide which are the best combinations to experiment with. More importantly, **the correlation between the predictor and the recommender tends to anticipate well when a hybrid will outperform its baseline.** Besides, Chapter 8 **presents a unified framework where the performance predictors are used to select and weight nearest neighbours in a standard user-based collaborative filtering algorithm.** The standard methodology from performance prediction is adapted and translated into this problem, where novel neighbour performance metrics are defined and the predictive power of the predictors is assessed.

The contributions related to the application part of the thesis are, in summary:

- We propose a dynamic hybrid framework to automatically decide when and how dynamic hybridisation should be done, depending on different conditions, namely the correlations between the recommenders and the predictors, and the relative performance level of the combined recommenders.
- In several experiments with the aforementioned performance predictors, our results indicate that a strong correlation with performance tends to correspond with enhancements in dynamic hybrid recommendation when the predictors are used for the adjustment of the combination weights.
- We propose a theoretical framework for neighbour selection and weighting in user-based recommender systems. This framework is based on performance prediction by casting the neighbourhood-based rating prediction task as a case of dynamic output aggregation.
- We compare several state-of-the-art rating-based trust metrics and other proposed neighbour scoring techniques, interpreted as neighbour performance predictors. We also propose several neighbour performance metrics that capture different notions of neighbour quality.

1.4 Publications related to the thesis

In the following international journal and conference papers we presented descriptions, results and conclusions related to this thesis:

Performance prediction and evaluation

1. Bellogín, A., Cantador, I., Díez, F., Castells, P., and Chavarriaga, E. (2012). An empirical comparison of social, collaborative filtering, and hybrid recommenders. *ACM Transactions on Intelligent Systems and Technology*, to appear.

2. Bellogín, A., Castells, P., and Cantador, I. (2011). Predicting the Performance of Recommender Systems: An Information Theoretic Approach. In Amati, G. and Crestani, F., editors, *ICTIR*, volume 6931 of *Lecture Notes in Computer Science*, pages 27–39, Berlin, Heidelberg. Springer Berlin / Heidelberg.
3. Bellogín, A., Castells, P., and Cantador, I. (2011). Self-adjusting hybrid recommenders based on social network analysis. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information*, SIGIR '11, pages 1147–1148, New York, NY, USA. ACM.
4. Bellogín, A., Castells, P., and Cantador, I. (2011). Precision-oriented evaluation of recommender systems: an algorithmic comparison. In *Proceedings of the fifth ACM conference on Recommender systems*, RecSys '11, pages 333–336, New York, NY, USA. ACM.
5. Bellogín, A. and Castells, P. (2010). A Performance Prediction Approach to Enhance Collaborative Filtering Performance. In Gurrin, C., He, Y., Kazai, G., Kruschwitz, U., Little, S., Roelleke, T., Rüger, S., and Rijsbergen, editors, *Advances in Information Retrieval*, volume 5993 of *Lecture Notes in Computer Science*, pages 382–393, Berlin, Heidelberg. Springer Berlin / Heidelberg.
6. Bellogín, A. and Castells, P. (2009). Predicting neighbor goodness in collaborative filtering. In And, T. A., Yager and, R. R., And, H. B., And, H. C., and Larsen, H. L., editors, *FQAS*, volume 5822 of *Lecture Notes in Computer Science*, pages 605–616, Berlin, Heidelberg. Springer Berlin / Heidelberg.

Content-based recommendation

7. Cantador, I., Bellogín, A., and Vallet, D. (2010). Content-based recommendation in social tagging systems. In *Proceedings of the fourth ACM conference on Recommender systems*, RecSys '10, pages 237–240, New York, NY, USA. ACM.
8. Cantador, I., Bellogín, A., and Castells, P. (2008). News@hand: A Semantic Web Approach to Recommending News. In Nejdl, W., Kay, J., Pu, P., and Herder, E., editors, *Adaptive Hypermedia and Adaptive Web-Based Systems*, volume 5149 of *Lecture Notes in Computer Science*, chapter 34, pages 279–283. Springer Berlin / Heidelberg, Berlin, Heidelberg.
9. Cantador, I., Bellogín, A., and Castells, P. (2009). Ontology-Based Personalised and Context-Aware Recommendations of News Items. In *Web Intelligence and Intelligent Agent Technology, 2008. WI-IAT '08. IEEE/WIC/ACM International Conference on*, volume 1, pages 562–565.
10. Cantador, I., Bellogín, A., Fernández-Tobías, I., and López-Hernández, S. (2011a). Semantic Contextualisation of Social Tag-Based Profiles and Item Recommendations. In Huemer, C., Setzer, T., Aalst, W., Mylopoulos, J.,

- Sadeh, N. M., Shaw, M. J., Szyperski, C., Aalst, W., Mylopoulos, J., Sadeh, N. M., Shaw, M. J., and Szyperski, C., editors, *Electronic Commerce and Web Technologies*, volume 85 of *Lecture Notes in Business Information Processing*, chapter 9, pages 101–113. Springer Berlin Heidelberg, Berlin, Heidelberg.
11. Fernández-Tobías, I., Cantador, I., and Bellogín, A. (2011). cTag: Semantic contextualisation of social tags. In *Proceedings of the Workshop on Semantic Adaptive Social Web (SASWeb 2011)*. *CEUR Workshop Proceedings*, vol. 730, pages 45–54. RWTH, Aachen (2011).

Collaborative filtering recommendation

12. Bellogín, A., Wang, J., and Castells, P. Bridging Memory-Based Collaborative Filtering and Text Retrieval. *Information Retrieval Journal*, to appear.
13. Bellogín, A., Cantador, I., and Castells, P. A Comparative Study of Heterogeneous Item Recommendations in Social Systems. *Information Sciences*, to appear.
14. Bellogín, A. and Parapar, J. (2012). Using Graph Partitioning Techniques for Neighbour Selection in User-Based Collaborative Filtering. In *Proceedings of the sixth ACM conference on Recommender systems*, RecSys '12, pages 213–216, New York, NY, USA. ACM. (best short paper award)
15. Bellogín, A., Wang, J., and Castells, P. (2011). Structured collaborative filtering. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, CIKM '11, pages 2257–2260, New York, NY, USA. ACM.
16. Bellogín, A., Wang, J., and Castells, P. (2011). Text Retrieval Methods for Item Ranking in Collaborative Filtering. In Clough, P., Foley, C., Gurrin, C., Jones, G., Kraaij, W., Lee, H., and Mudoch, V., editors, *Advances in Information Retrieval*, volume 6611 of *Lecture Notes in Computer Science*, chapter 30, pages 301–306. Springer Berlin / Heidelberg, Berlin, Heidelberg.
17. Bellogín, A., Cantador, I., and Castells, P. (2010). A study of heterogeneity in recommendations for a social music service. In *Proceedings of the 1st International Workshop on Information Heterogeneity and Fusion in Recommender Systems*, HetRec '10, pages 1–8, New York, NY, USA. ACM.

Social filtering recommendation

18. Díez, F., Chavarriaga, J. E., Campos, P. G., and Bellogín, A. (2010). Movie recommendations based in explicit and implicit features extracted from the Filmtipset dataset. In *Proceedings of the Workshop on Context-Aware Movie Recommendation*, CAMRa '10, pages 45–52, New York, NY, USA. ACM.

Time-aware recommendation

19. Campos, P. G., Bellogín, A., Díez, F., and Cantador, I. (2012). Time Feature Selection for Identifying Active Household Members. In *Proceedings of the 21st ACM international conference on Information and knowledge management, CIKM '12*, New York, NY, USA. ACM (to appear).
20. Campos, P. G., Díez, F., and Bellogín, A. (2011). Temporal rating habits: a valuable tool for rating discrimination. In *Proceedings of the 2nd Challenge on Context-Aware Movie Recommendation, CAMRa '11*, pages 29–35, New York, NY, USA. ACM.
21. Campos, P. G., Bellogín, A., Díez, F., and Chavarriaga, J. E. (2010). Simple time-biased KNN-based recommendations. In *Proceedings of the Workshop on Context-Aware Movie Recommendation, CAMRa '10*, pages 20–23, New York, NY, USA. ACM.

Hybrid recommender systems

22. Cantador, I., Castells, P., and Bellogín, A. (2011). An enhanced semantic layer for hybrid recommender systems. *International Journal on Semantic Web and Information Systems*, 7(1):44–78.
23. Cantador, I., Bellogín, A., and Castells, P. (2008). A multilayer ontology-based hybrid recommendation model. *AI Commun.*, 21(2-3):203–210.
24. Cantador, I., Castells, P., and Bellogín, A. (2007). Modelling Ontology-based Multilayered Communities of Interest for Hybrid Recommendations. In *Workshop on Adaptation and Personalisation in Social Systems: Groups, Teams, Communities, at the 11th International Conference on User Modeling*.

Recommender evaluation

25. Bellogín, A., Cantador, I., Castells, P., and Ortigosa, A. (2011). Discerning Relevant Model Features in a Content-based Collaborative Recommender System. In Fürnkranz, J. and Hüllermeier, E., editors, *Preference Learning*, chapter 20, pages 429–455. Springer Berlin Heidelberg, Berlin, Heidelberg.
26. Bellogín, A., Cantador, I., Castells, P., and Ortigosa, A. (2008). Discovering Relevant Preferences in a Personalised Recommender System using Machine Learning Techniques. In *Preference Learning Workshop (PL 2008), at the 8th European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD 2008)*, pages 82–96.

These publications are related to the contents of this thesis as follows. In [4] we analyse different evaluation methodologies available in the recommendation literature (Chapters 3 and 4). In [2], [5], and [6] we define the formulations for the concept of user clarity based on ratings (Chapter 6), whereas in [1] and [3] we define the social-

based predictors (again, Chapter 6). Besides, in [1] and [3] we also investigate the use of performance predictors for dynamic hybrid recommendation (Chapter 7). Moreover, in [5] and [6] we address the problem of neighbour weighting based on neighbour performance predictors (Chapter 8).

Additionally, during the course of the thesis, the research presented here has motivated a number of publications that address broader topics in the field, such as content-based recommendation [7-11], collaborative filtering [12-17], social filtering techniques [18], time-aware recommendation [19-21], hybrid recommender systems [22-24], and recommendation evaluation [25, 26]. These publications have resulted in the use and construction of datasets, the development of algorithms and the research and use of some evaluation methodologies and metrics that appear in this thesis.

Additional publications

Preliminary work towards the approaches presented in this thesis was published in my Master's Thesis entitled "Performance prediction in recommender Systems: Application to the dynamic optimisation of aggregative methods" (Bellogín, 2009); specifically, the concept of performance prediction for recommendation is proposed in such work. Apart from that, the motivation, potential impact, and initial main results of our research were published as contributions in two international doctoral symposiums:

- Bellogín, A. (2011). Predicting performance in recommender systems. Doctoral Symposium. In *Proceedings of the fifth ACM conference on Recommender systems*, RecSys '11, pages 371–374, New York, NY, USA. ACM.
- Bellogín, A. (2011). Performance Prediction in Recommender Systems. Doctoral Symposium. In Konstan, J., Conejo, R., Marzo, J., and Oliver, N., editors, *User Modeling, Adaption and Personalization*, volume 6787 of *Lecture Notes in Computer Science*, pages 401–404, Berlin, Heidelberg. Springer Berlin / Heidelberg.

Furthermore, the following submissions are under revision, some of them closely related to the topics of the thesis:

- Bellogín, A., Castells, P., and Cantador, I. Statistical Biases in IR Metrics for Recommender Systems: A Methodological Framework for the Adaptation of the Cranfield Paradigm. Under review.
- Bellogín, A., Castells, P., and Cantador, I. Neighbour Selection and Weighting in User-Based Recommender Systems: A Performance Prediction Approach. Under review.
- Parapar, J., Bellogín, A., Castells, P., and Barreiro, Á. Relevance-Based Language Modelling for Recommender Systems. Under review.

1.5 Structure of the thesis

The thesis is divided into six parts. The first part introduces and motivates the problem addressed, along with a survey of the Recommender Systems field, where this thesis is framed. The second part describes the different evaluation techniques used in the recommender systems literature and provides an analysis of the design alternatives and statistical biases that may arise. The third part gives background knowledge and a literature survey on performance prediction, proposes translations of this concept into the recommender system space, and evaluates the predictive power of these approaches. The fourth part provides two applications of the proposed recommender performance predictors. The fifth part concludes and summarises the main contributions of this thesis. Additional information and details are provided in the last part.

In more detail, the contents of this thesis are distributed as follows:

Part I. Introduction

- **Chapter 1** presents the motivation, research goals, contributions and publications related to the thesis.
- **Chapter 2** provides an overview of the state of the art in recommender systems, considering a classification of the main types of recommendation approaches. We also describe the weaknesses of the different recommendation techniques and present a broader class of hybrid recommenders that aim to overcome these limitations.

Part II. Evaluating Performance in Recommender Systems

- **Chapter 3** describes the main evaluation metrics and methodologies used in the recommender systems field. The public datasets commonly used in the field are also described.
- **Chapter 4** provides an analysis and formalisation of the different evaluation methodologies reported in the literature. First, we present a systematic characterisation of the experimental design alternatives. Next, we identify and analyse specific statistical biases arising when some methodologies are applied to recommendation, and propose two alternative experimental designs that effectively neutralise such biases to a large extent.

Part III. Predicting Performance in Recommender Systems

- **Chapter 5** presents the problem of performance prediction in Information Retrieval, surveys the main research works in that area, both in the definition of (query) performance predictors and also in the predictor evaluation in order to infer their predictive power.

- **Chapter 6** states the problem of performance prediction in recommender systems. We define several performance predictors based on three recommendation input spaces where we qualitative analyse the predictive power of the predictors.

Part IV. Applications

- **Chapter 7** proposes a framework where recommender performance predictors are used to build dynamic hybrid recommender systems. We evaluate these recommenders in the three input spaces previously considered for the definition of performance predictors and using different experimental design alternatives where some statistical biases are neutralised.
- **Chapter 8** restates the user-based recommendation problem, providing a generalisation as a performance prediction problem. We investigate how to adopt this generalisation to define a unified framework where we conduct an objective analysis of the effectiveness (predictive power) of neighbour scoring functions.

Part V. Conclusions

- **Chapter 9** concludes with a summary of the main contributions of this thesis, and a discussion about future research lines.

Part VI. Appendices

- **Appendix A** provides details about the methods proposed in this thesis: configuration of the recommendation algorithms and parameters of the experimental designs used in the evaluation. Detailed statistics about the datasets used in the experiments are provided, complementary to those given in previous chapters.
- **Appendix B** contains the translation into Spanish of Chapter 1.
- **Appendix C** contains the translation into Spanish of Chapter 9.

