Abstract

Over the last decade, researchers have been interested in personalised recommender systems as they have detected the many profitable applications they have. As a direct effect many algorithms and different techniques have been developed, some of them combining user and item categorisation techniques. The purpose of this work is to bring the capabilities of such systems to the context of museums and art galleries and combine them with advanced visualisation techniques. This will be done in order to help the users to explore a large dataset of European artworks from the 11th until the 19th century while taking advantage of a recommendation engine that will help them finding interesting artworks rapidly. To pursue the aforementioned goals, an existing application that already contains such advanced visualisation techniques and the dataset has been enhanced with recommendation capabilities. The recommendations are based on semantic user profiling techniques extracted unobtrusively from users.

Using semantic user profiling to provide recommendations enables the possibility to have a hybrid recommendation algorithm with content-based and collaborative features. The recommendation algorithm extracts the preferences from users through their interactions with the system and uses them to find related content that might be interesting for the user. An offline experiment has been conducted in order to assess the relative accuracy of the recommendations and the user categorisation based on semantic user profiling. The offline experiment reveals the advantages of combining such techniques and shows promise on a theoretical level on how to bring recommendations in such a restricted and relatively scarcely explored environment like the applications for museums and art galleries.

Acknowledgement

First and foremost I would like to thank everyone who helped making this accomplishment possible, but as most of my relatives are not familiar with the language in which this work has been written I will start with my supervisor Prof. Bruno Dumas and promoter Dr. Beat Signer, thanks for your patience and the time spent improving my work. Special gratitude to my beloved family, Muchas gracias por estar presentes, apoyarme y hacer hasta lo imposible para brindarme la oportunidad de lograr lo que he logrado.

Last but not least thank you, for lifting my heart during the hardest times.

Contents

Chapter 1

Introduction

Visitors of museums and art galleries find themselves quite often struggling to find content relevant to their liking, or for more advanced visitors their research or work. But bearing in mind that the average application user avoids giving appropriate feedback, finding a usable, unobtrusive way to deliver relevant content to these users can be a challenging task [41], especially since mobile devices or other external artifacts can dramatically change the visitor's experience and its way to appreciate the works of art in its surroundings.

These approaches have already been taken into account providing users with a whole set of augmented reality features for their visits in [6, 18], but in order to avoid distracting the visitors from their goals and still be able to provide relevant content, an in-depth study of the person behind the "system user" is required, and then model it according to an arbitrary set of categories [4]. These categories can be further processed to avoid major forms of interaction with external systems or devices to deliver tailored content to the final user [4].

The easiest path towards this study of the user starts with the modelling of the user and his environment, for which noticeable improvements have been made in the recent years [51, 8]. Semantic user profiling consists of studying the individual behind an application and allowing the system to learn his interests during his interactions with it over long periods of time [84]. The implementation of semantic user profiling accomplishes this by allowing an information system to gather and analyse information about its user to modify its own behaviour according to certain parameters.

To deal with this personalisation model some researchers have opted to define a small set of stereotypes which cover the mainstream users that interact with their systems. These stereotypes can be used to model the context of these users in order to offer them more relevant information and recommend tailored content [80].

Modelling and building user profiles in order to implement recommender systems based on statistical approaches seems to have a very promising future, especially while trying to improve the overall experience of the visitors in art galleries and museums, as these visitors are often overwhelmed with vast amounts of artworks in huge spaces, making it really difficult for them to find everything that they may consider interesting, missing a lot of content that under other circumstances would have probably enriched their visit.

Tsasou et al. [84] divided the goal of modelling the user profile in two main tasks during their research, user modelling based on user transactions and user modelling based on a user physical behaviour. These two categories have been united in this work as in the scope of applications for museums and art galleries part of the user behaviour can be explicitly mapped to transactions in the system.

Another challenging stage of this process is finding a way to analyse the gathered information to refine the search results so our users can constantly find relevant information to their current studies or interests. After defining the user context model, it will be easier to determine how this information will be stored and how it shall be further analysed.

Although the implementation itself can be rather simple, the design is not. The trick is to accurately define a set of constraints that will eventually allow us to query the analysed information. The core of this whole process is based on accepting the fact that building a system that implements semantic user profiling requires the construction of a whole subsystem or module in charge of collecting information about the user and processing it to improve the accuracy of the queries that the user unintentionally makes while he browses through the UI. The tasks of this module can be subdivided in two main categories: information extraction and annotation with partial analysis.

The information extraction phase consists of logging the user interactions with the system and mapping it to a previously defined ontology.

In the partial analysis phase, the system pre-processes the gathered information. When the user performs a query on the system it will consider the outcome of the latter process to give relevance to certain results displaying it in a more attractive way to the user.

This work will be built on top of the ArtVis system, a platform that allows its users to browse through the information of several art galleries and museums in Europe and their corresponding artworks. The system will be described in detail later.

1.1 Objectives

- Extract enough information about the users to build a profile that allows the system to assess its level of expertise in art-related subjects and recommend content using unobtrusive information extraction techniques.

- Give semantic value to the interactions users have with the system enabling it to analyse them to recommend tailored content.

- Explore the different kinds of recommender systems and techniques in order to determine which should be used in the context of art-related applications.

- Combine advanced visualisation techniques with user information management to improve usability in the existing ArtVis application.

- Encourage users to develop an interest in art by enhancing usability in museum applications through unobtrusive information extraction and advanced visualisation techniques.

1.2 Research Questions

Personalising the behaviour of applications depending on the user in interaction is an active research topic that has been explored since the emergence of the semantic web. The way to address this problem is through the creation of evolutionary semantic profiles that analyse the context of the user in different aspects to determine how to respond to future requests.

One of the most common usages of these profiles is personalisation on recommender systems, which will be addressed often since the goals of these kinds of systems is indeed close to the ones of this work: Deliver more relevant content to the users according to a predefined set of heuristics based on long term interactions. But the longer the period of interaction the harder it gets to accurately represent the users interests, since often these interests tend to be evolutionary, meaning that what today is the focus and center of attention for a user may be completely irrelevant tomorrow [46].

Besides the latter, challenges arise due to the multidisciplinary nature of the users which can be constantly interested in very different topics. This clouds the way to its real current interests while interacting with the system.

Furthermore, as users tend to avoid giving explicit feedback after using information systems [7], it is necessary to circumvent these options by providing the system with other tools to approach the user and learn about its interests, being as unobtrusive as possible for the purpose of gaining usability.

The following research questions will be addressed during the development of this work:

a. *Q1:* Which methods should be used to extract information about users preferences in an unobtrusive way?

b. *Q2:* How can the system analyse the information about a user in order to provide personalised search results that reflect its interests?

c. *Q3:* Is it possible to categorise and group users by their preferences in art?

d. *Q4:* What relevance weight should be assigned to each category of users?

To address these research questions the ArtVis system will be introduced.

Chapter 1 provided some insight in the domain of art related applications, Chapters 2 and 3 contain the state-of-the-art and an analysis on the aforementioned techniques. Chapter 4 consists of an overview of the application used as a starting point for this work. In Chapter 5 a description of a set of scenarios of applicability can be found. Chapter 6 explains the methods used to enhance the system with recommending capabilities and Chapter 7 the results of this work based mainly on an offline experiment and an expert user assessment. Lastly Chapter 8 contains the discussion and future work and the conclusions are presented in Chapter 9.

Chapter 2

Related Work

2.1 Overview

In order to find a way to effectively combine users' preferences management with advanced visualisation techniques in the context of museum applications, the main techniques involving those fields have been explored. The main goal is to exploit the semantic fingerprints users leave on the system with every interaction they make to improve the content presented to them. The closest research areas that address this issue are recommender systems and semantic user profiling, thus a special effort has been made to explore the past, present and future trends in such matters.

As a previous step, and with the purpose of increasing usability features even further, unobtrusive information extraction techniques have been explored. It is assumed that such techniques can be applied to gather the necessary data from the aforementioned semantic fingerprints. An important part of this research field is the security implications of applying remote unobtrusive information extraction techniques. Nonetheless, given the context that is being addressed in this work, those security implications were not taken into account. At last, other art related applications were looked into to provide a better insight on how the other preceding research areas have been applied in the same direction and explore the possibilities to go one step further from what has already been done.

The exploration into the world of recommender systems has been done categorising them in their main approaches: content-based recommendations, collaborative recommendations and hybrid techniques, then further subdivided in heuristic or memory-based systems and model-based systems.

Semantic user profiling is being used in the research on recommender systems as a tool to extract information about users and perform categorisation on items to improve the accuracy of recommendations. Therefore some of the most relevant trends in semantic user profiling and categorisation have also been addressed. In the end of this section Tables 1 and 2 illustrate the relationship between recommender systems and categorisation techniques and depict the main advantages and disadvantages of such techniques.

The art related applications section depicts the current research trends that combine user information management and recommendations. Two main approaches were studied. In the first one, a user categorisation is made depending on an arbitrary set of features such as museum size, visiting style, type of visitors, demographic factors and age. The second one presents an overview of the MNEMOSYNE system, which is probably the closest existing application to this work.

The MNEMOSYNE system combines unobtrusive information extraction techniques, semantic user profiling features and custom made recommendations with facial expressions and gestures analysis to determine users interests and filter categories of items that the user is not interested in. This system provides interactive tables in which users can have an overview of their visit and find recommendations based on the observations registered by the sensors. Nevertheless, the intention of this work is to provide a much broader view on the artworks, allowing the user to browse through a broader dataset and combine it with more accurate recommendations based on explicit interactions with the system.

2.2 Information Extraction

The first analysis concerning query results was based on a deterministic approach, in the sense that a query was supposed to return the same results for all users [76]. This perception quickly evolved to a more flexible one in which a long term based user profile could influence queries to filter or retrieve information based on user preferences [56, 9]. These user profile and preferences are built through a long phase of user information extraction, in which applications learn from the user through every single interaction they have [45], and in other more advanced approaches, external elements observe the user behaviour and physical reactions towards special items in their surroundings and perform the analysis with little or no human interaction [51].

Unobtrusiveness is an important aspect to take into account while extracting information from users, especially as it usually makes more difficult the process of learning from their behaviour. Monitoring users can be effectively used to register some examples of what users are looking for, and heuristics have been developed to infer negative examples as well as to help the filtering process, although it is a rather limited approach [65].

2.3 Recommender Systems

Recommender systems are software artifacts that provide suggestions to users in a specific domain [20, 72]. Since their first mentions in the literature, recommender systems have had an increasing importance as a research area, mainly because of the great amount of practical applications that it has. Recommendations can become useful to improve information systems as a way to browse efficiently through large datasets by highlighting interesting items in the domain [20, 72].

In the great majority of cases, personalised recommendation is done by ranking the available items on the system according to a fixed set of parameters, contrasted with the implicit or explicit preferences of a user, meaning that a user can provide explicit feedback of its preferences or they can also be deduced by its usage of the system [59]. Besides that, these recommender systems tend to gather information about the accuracy of these recommendations based on the reaction of the user to the recommendations made, and then tries to apply them to the next sets of users with similar characteristics to improve the rank of the selected items. This approach was first called collaborative filtering [59].

In order to preserve the information on users preferences to be used repetitively, recommender systems evolved to create long term user profiles, aiming to increase the accuracy of recommendations by providing the recommender engine with more data to analyse.

Enhancing queries with long term user profiles is somehow similar to providing recommender systems capabilities to search for content, in the sense that the system learns some information about the user throughout interactions in an arbitrary period of time, and uses this information to provide content that has a higher probability to be useful to the user. A formal definition of the base problem is given as follows:

Let C be the set of users and S the set of items that can be recommended, taking into account that S can be a very large set, containing thousands or even millions of items. Let u be a utility function that measures the usefulness of an item s to a user c,i.e., $u : CxS \rightarrow R$, where R is a totally ordered set.

For a particular user or subset of users, the selection of an arbitrary number of elements of R that maximise each user's utility such that $\forall c \in C, S'_c =$ arg max u (c, s) [1] is desired.

The most commonly accepted formulations of this problem were stated in [44, 70, 81].

Typically recommendations are made following one out of three methods:

- Content-Based Recommendations
- Collaborative Recommendations
- Hybrid Approaches

2.3.1 Content-Based Recommendations

The user gets recommended items similar to the ones with which he has had interactions in the past.

Taking the latter into account it would then be accurate to assume that content-based recommendation approaches find their basis on information retrieval and filtering [10]. Due to these similarities, many content-based systems recommend items using textual information, such as documents and URLs. As a result of this text-based angle, content is grouped by the use of keywords [1].

In information retrieval the measure the keywords weights is given by the term frequency/ inverse document frequency(TF-IDF) defined as [75]: Assuming N as the quantity of items available for recommendation, and that keyword k_j is present in n_i of them, $f_{i,j}$ is the number of times keyword k_i appears in item d_j. Then, $TF_{i,j}$, the term frequency (normalised) of keyword k_i in item d_j is defined as:

$$TF_{i,j} = \frac{f_{i,j}}{max_z f_{z,j}} \tag{2.1}$$

Being the maximum computed over the frequencies $f_{z,j}$ of all available keywords k_z in the item d_j. Nevertheless whenever keywords appear in several items, they cannot be used for differentiation between relevant and non relevant ones. In such cases the measure of inverse document frequency (IDF_i) is used in combination with the term frequency ($TF_{i,j}$), defining the IDF for keyword k_i as follows:

$$IDF_i = log \frac{N}{n_i} \tag{2.2}$$

Finally, the TF-IDF weight for keyword k_i in item d_j is:

$$w_{i,j} = TF_{i,j}XIDF_i \tag{2.3}$$

And the content of item d_j is:

$$Content(d_j) = (w_{1j}, ...w_{kj}). \tag{2.4}$$

To conclude, content-based recommendation systems make use of the available information of items liked by a user or a group of users in the past to recommend content, in several cases using keyword analysis techniques derived from information retrieval [74].

2.3.2 Collaborative Recommendations

The user gets recommended items that were liked by people with similar profiles in the past. This approach was first proposed as a solution to the new user problem of the content-based method. The new user problem or cold start problem states that in order for the system to recommend content, the user first has to rate a sufficient amount of items. Consequently, new users cannot be granted accurate recommendations [61].

In order to find users with similar preferences, the system establishes some parameters to find similarities in the available information on a set of users and classifies them as "peers". These "peers" comply with the stereotype paradigm, first seen in the Grundy system [73], which was used for defining models of users based on limited amounts of information about each individual user on the system.

The algorithms used by collaborative recommendation were grouped in two categories: memory based and model based [71, 17]. Memory or heuristic based algorithms consist of a set of heuristics that predict usability ratings based on all available information on items previously rated by the users, so, the usability index(rating) $r_{c,s}$ for the user c and item s gets computed by means of the ratings of a group of similar users (usually the N most similar) For the same item s:

$$r_{c,s} = aggr r_{c',s}, c' \in C. \tag{2.5}$$

Where C is the set of N users that were identified as the most similar to user c who have also rated item s. The aggregation function can be as simple a an average, but the most common approach is a weighted sum, in which ratings given by more similar users have a higher weight in the computation of the final predicted rating.

One problem not addressed in this approach is the fact that different users may use the rating scale in a different way, therefore more complex solutions have been proposed, such as the adjusted weighted sum, in which the absolute values of the ratings are not used, but the deviations from the average rating of the corresponding user are used in its place. Preference-based filtering has also been used to overcome the same limitations [34, 49].

To compute the similarity index between users, several approaches have also been taken into account. The most common one is based on their ratings of items [1]. More complex approaches take also into account information that may seem external to the problem but can help categorise users in a better way, for instance using their age, location, ethnicity or any other information that the system may have obtained while creating their profiles or registering them in the system. At a low level of implementation of these system, the most popular methods used to calculate these similarities are correlation and cosine, both optimised by modifications such as default voting (assuming a default rating for missing rating), inverse user frequency, case amplification and weighted majority prediction [1, 17, 28].

Model based methods, in opposition to heuristics, use the previously registered ratings to learn a model and to predict further ratings. In most cases the models are built based on statistical analysis or machine learning techniques. Bayesian networks, cluster models [17] and probabilistic procedures [11, 37] are the most noteworthy among them.

Having partially avoided the new user problem, new challenges arise. One of the most predictable ones is the new item problem, in which an item will not be recommended by the system until it gets enough ratings, but there is a chance that the item will remain unrated because it has never been recommended [1]. In order to cope with this new challenge, hybrid approaches were explored.

2.3.3 Hybrid Approaches

Hybrid approaches are combinations between the content-based and collaborative recommendation techniques.

Hybrid approaches can be further sub-categorised in four classes:

1. Combining predictions of an instance of both content-based and collaborative methods.

2. Adding content-based features to a collaborative implementation.

3. Adding collaborative features to a content-based implementation.

4. Constructing a model that unifies characteristics of both approaches.

Depending on the general purpose of the recommender system, one of the most challenging aspects to be considered is how to handle the influence capability that users may have among themselves. In web-based recommender systems, studies have shown that there are users whose goal is to influence ratings of particular items to influence others into viewing/purchasing those items. These kinds of biases are extremely hard to detect and counteract against [43].

Selecting the right data sets to study can also be a determining task at the time of testing a recommender system. Key aspects that need to be considered include the capability of performing offline evaluations, using simulated data when real data is not available, and establish the properties that such data should contain. According to [43], in cases where recommending new items is desired, it may be inappropriate to use only off-line evaluation. Taking into account that the system is generating recommendations that the user does not already know about, it is probable that the data set will not contain enough information to evaluate the quality of the items to recommend, while performing a live evaluation can help gain ratings on the spot for each item recommended [62].

As for the possibility to use synthetic data sets, it may be recommended in such cases where research has been carried out to roughly determine the structure of user preferences but no data sets have been generated. This way, promising algorithms can be identified for further study.

Herlocker et al. [43] also divide the properties of data sets for recommender systems in three categories, with the purpose of evaluating the usefulness of data sets to model the tasks for the recommender system that those data sets are addressing. Those categories are:

Domain features describe the nature of the content to be recommended, loosely coupled from the system. Usually includes information such as the topic under recommendation with its associated context, user tasks supported, novelty and quality desired in recommendations, the cost/benefit ratio of false/true positives/negatives and the granularity of user preferences.

Inherent features reflect the characteristics of the recommender system for which the data set is going to be used along with its data collection techniques. Includes information about the obtrusiveness with which ratings are collected, the scale of the ratings, dimensions and presence or absence of timestamps on the ratings.

Sample features consider the distribution properties of the data, which can be manipulated selecting different subsets of the data set to be analysed. Inherent features include information about whether the recommendations made to a user were recorded or not, the availability of user demographic information and item content information.

Regardless of the collection technique, most data sets that involve user ratings present a certain level of bias, especially at the time to create the initial ratings before users are able to interact with the system [38, 69]. This is a challenging issue to overcome, especially in certain domains. Literature shows that one of the most significant challenges in the field is the lack of publicly available collaborative filtering data sets [43].

2.4 Semantic User Profiling

In [65] Middleton et al. presented their approach to apply semantic user profiling to recommender systems, which has become one of the main tools for developers to cope with the complexity of understanding the preferences of their users. Later on, Sieg et al. [82] used ontological user profiles to represent context in web search and in 2010, Hopfgartner and Joemon presented a method to apply semantic user profiling techniques to multimedia recommendation [46].

Before the semantic web and ubiquitous computing users had to cope with the complexity of formulating queries to find content on the web that was not necessarily well structured or tagged. Keywords were one of the first approaches to reduce the complexity of such tasks, but limiting the vocabulary available to fit the thoughts of inexperienced users into them is not that appealing, especially while browsing large data sets [65]. With the semantic web more intelligent queries are possible, and although it is still on a relatively early stage its potential is such that its most remarkable features such as data annotation can benefit the problem of effective searching of content.

The problem of creating user profiles to enhance search is usually either knowledge-based or behaviour-based. The first one develops static models of users and categorises them dynamically to the closest model, while behaviour-based use the users' behaviour as the model and often apply machine learning techniques to discover interesting patterns. Recommender systems for instance generally use behaviour-based techniques to categorise users, with binary models to establish what they find interesting and uninteresting. Machine learning techniques are then applied to rate the items in the domain that the user is searching in order to provide recommendations. In [77], Sebastiani and Fabrizio explain several examples of machine learning techniques used in automated text categorisation.

Ontologies are shared and agreed conceptualisations of ”the world”, referring to ”the world” as a specific domain in which generally information systems function and inter-operate [40]. These rich conceptualisations represent the main concepts and their relationships. Given the representation, they could be used to represent isolated pieces of information useful for instance, in this case to facilitate the creation of user and item categories that interact with each other to develop models.

The work described by Middleton et al. in [65] describe an ontological approach to content recommendation using a hybrid recommender system, containing both collaborative and content-based recommendation techniques and depicting user profiles in ontological terms. They performed their research with two experimental systems, Quickstep and Foxtrot. Quickstep is a recommender system to be used in a computer science laboratory by a set of researchers and Foxtrot is a recommender system and searchable database for a computer science department.

In order to address the unobtrusiveness problem, they used a proxy to monitor users' web browsing activities and added new items to the searchable database as users discovered them, making them available to all users. Recorded web browsing and relevance feedback is later used to engineer user profiles and synthesise users' preferences [64]. Collaborative filtering techniques were used to compile the recommendations on a daily basis. Having found the interesting papers to save, they are constrained to match the content-based profiles. Finally the papers left are used in recommendations.

The research paper topic ontology is mainly based on the classifications made by the dmoz open directory project [64]. In order to speed development time, an existing taxonomy is used thus it is also extensible with other external ontologies. The obtained profiles hold a set of topics and interesting tags for each day. User feedback is used to adjust interests in different topics and a time decay function is implemented to prioritise newer papers over old ones.

According to their results, it is possible to deduce that using ontological inference to create the profiles results in superior performance than using lists of unstructured topics. These profiles are also more accurate as the recommendations turned out to be more useful according to the feedback provided. The method also allows to discover interests that were not directly observed in the users' behaviour.

The daily learner [12] uses a k-NN clustering algorithm for short term profiling and a naive version of Bayesian networks for long term profiling, reporting a precision level of 33% and recall of 29%, for the top 4 recommendations monitoring the selected stories for users to read. There is also some research to expand profiles using contextual information in task modelling [19]. Next, some of the most popular approaches to the classification problem are presented.

2.4.1 Nearest Neighbours

The nearest neighbours approach is a classification method used to predict the class of unseen cases with the use of training records. A route learner for instance, memorises the training set and classifies items into classes only when the attributes of the new records match the ones in one of the training records exactly [19]. kNN classifiers are more elaborated examples of such kind. They are instance-based classifiers that given a point (item), find the closest points or nearest neighbours. They then assign the new item the class of such nearest neighbours [30]. The classification is done under the assumption that if an item falls in a particular neighbourhood, it is likely that this item belongs to that neighbourhood (class) [25].

Given a certain point p that needs to be classified, and having a training set $X = x_1, C_1, ..., x_n$ where C_j represents the class and x_j is the j-th element, the algorithm will find a subset $Y = y_1, C_1, ...y_k$ such that $Y \in X$ and $\sum_{i=1}^{k} d(p, y_k)$ is minimal. Y contains the k points in X closest to p. Then, $C = f(C_1, ..., C_k)$ is the class assigned to point p [59, 30]. Although the method in general is rather simple, finding the right k can be a challenging task, choosing a k value too small will sensitise the classifier to noise points, and choosing a value that is too high will cause the inclusion of too many points from other classes.

2.4.2 Rule-based Classifiers

Rule-based classifiers address the classification problem by the usage of a set of rules of the form "If A then B" where, A or the antecedent is an expression made by attribute conjunctions and B is a binary classification referred to as the consequent [3]. It is said then that an instance x is covered by a rule r if the attributes of the instance satisfy the condition. The coverage index of a rule is defined by the fraction of records satisfying the antecedent, and the accuracy by the fraction of records satisfying the antecedent and the consequent [59]. Other properties as mutual exclusion and exhaustive rules are also addressed as the inclusion of rules that are independent from each other and the capacity of accounting for every possible combination of attribute values respectively.

2.4.3 Clustering

Clustering techniques are broadly used to classify items. That is grouping elements into groups such that the elements in the same groups are more similar to each other than to elements in other groups [42]. The main goal is to discover natural groups existing in data, detect patterns and making meaningful categorisations. The similarity between items is determined using a distance measure, given by arbitrary features or dimensions of the data. Vectors of measurable and in some cases weighted properties are used for this purpose with the goal of maximising distances between clusters and minimising intra-cluster distances [53, 59].

Clustering techniques can be categorised in two groups: hierarchical and partitional. The first ones cluster elements inside other previously found clusters, creating a hierarchical structure similar to a tree. The second ones divide data in non-overlapping clusters. This way each item is found in exactly one cluster [83].

K-means clustering is a technique based on partitioning. A dataset of N elements is partitioned into k separate subsets S_j containing N_j elements being as close as each other as possible according to an arbitrary distance measure. Each cluster has a centroid λ_j which is the point that minimises the sum of all distances from all elements. The cluster then is defined by its centroid and its N_j items [16, 78]. The k-means technique is iterative and its goal is to minimise E that represents the sum of all distances and is defined as follows:

$$E = \sum_{i=1}^{k} \sum_{n|n \in S_j} d(x_n, \lambda_j) \tag{2.6}$$

Where d is the distance measure, x_n is a vector that represents the n-th element and λ_j is the centroid of the item in S_j. The algorithm stops when E cannot be decreased any further [59].

In general the centroids are selected randomly although modern approaches have been working in the optimisation of such selections [16, 50]. After selecting k centroids, all elements are assigned to the cluster with the closest centroid. Membership to the clusters need to be updated when adding new items.

2.4.4 Bayesian Classifiers and Bayesian Networks

Bayesian classifiers are another popular approach to the classification and categorisation problem. This method is purely based on probabilistic techniques [24]. The goal is to predict a class C_k that maximises the posterior probability of the class given the data $P(C_k|A_1, ..., A_n)$ being $(A_1, ...A_n)$ the vector of n attributes under evaluation. This is done by applying the Bayes' theorem:

$$P(C_k|A_1, ..., A_n) \propto P(A_1, ...A_n|C_k)P(C_k) \tag{2.7}$$

Meaning, posterior is proportional to prior times likehood [58]. Priors represent the expectations or the prior knowledge about the data and what the relationship might be. Posterior is the probability of the model given the data [36].

Naive Bayesian classifiers for instance, learn the conditional probability of every attribute A_i in data given the class label C [36]. The data is grouped then by calculating the probability of C given the particular instance of A_i and assigning them to the class with the highest posterior probability. It is assumed that all attributes are conditionally independent given the value of C. Then:

$$Pr(C) > 0 \Rightarrow Pr(A|B, C) = Pr(A|C) \forall A, B, C \tag{2.8}$$

That is, A is independent of B given C. Even though the previously mentioned assumption is utterly unrealistic, the performance of the classifier is indeed surprising [36].

Whenever the independence assumption does not hold, the method needs to be extended in order to encode such existing dependencies between the attributes of data. This is usually done by implementing Bayesian Belief networks [59]. Bayesian Belief Networks (BBN), or Bayesian Networks as they are also known, use acyclic graphs and a probability table that associates each node to its immediate parents for such purpose [36, 35]. BNNs use graphical models to capture the prior knowledge in the domain.

2.4.5 Artificial Neural Networks

Another alternative to the classification problem are the Artificial neural networks (ANN), whose architecture is inspired by the biological brain. ANNs are assemblies of nodes (neurons) and weighted links, that allow the network to learn the classification problem after training with sufficient data [88].

Artificial neural networks are composed by processing units, Figure 2.1 illustrates an example of such units as nonlinear summing nodes. Bias is taken into account and is represented by the grey square in the bottom. S_j is the incoming sum and a_j the activation value for unit j. w_{ji} represents the weight from unit i to unit j [27].

Layers in artificial neural networks are categorised into three classes:

- Input: Serve as façade to the data entering the network

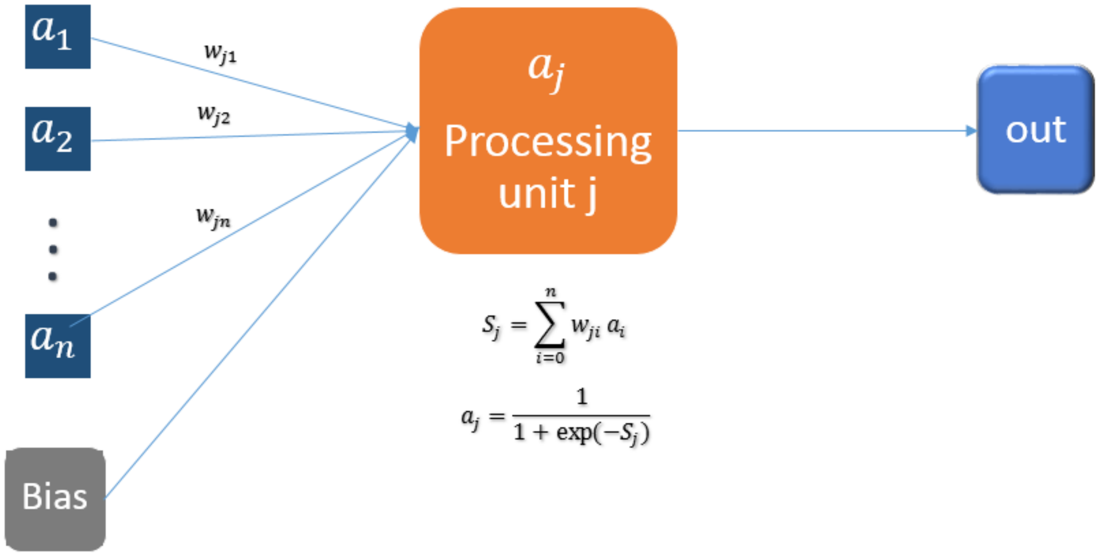

Figure 2.1: Processing unit of an ANN

- Hidden: Receive the weighted data from the input layers
- Output: Generate the output of the network after processing the weighted data from the hidden units.

An artificial neural network can have as many layers as needed. There are several possibilities to design an architecture for a network when using nodes as atomic functional units, nevertheless the most commonly followed approach due to its simplicity is the feed-forwards model in which data is always passed in one way: from input to output [88].

2.4.6 Decision Trees

Decision trees use a graph-like structure similar to a tree to make decisions and analyse their outcomes [87]. Classifiable items are seen as compositions of attributes and their target value. Their nodes can be either decision nodes or leaf nodes. Decision nodes are the ones to which attribute-values of items are tested to determine to which branch apply, whilst leaf nodes indicate the value of the target attribute [59]. Figure 2.2 illustrates a basic example of a decision tree. Square nodes represent decisions whilst circle nodes are circumstances that can have unexpected outcomes. They can contain values for attributes of data of which nothing or very little is known about but that can be used to detect patterns that facilitate the classification problem.

Some of the most noticeable examples of decision trees algorithms are: Hunts algorithm [67], CART [47], ID3 [85], C4.5 [68], SLIQ [63, 22] and SPRINT [79]. Recursive Hunt for instance, depends on a test condition, given to an attribute to discriminate observations by their target values. The partition is then induced by their target values and the process is repeated until the partition is empty or all observations have the same value [59]. The information gain is defined as follows:

$$\Delta_i = I(parent) - \sum_{j=1}^{k_i} \frac{N(v_j)I(v_j)}{N} \tag{2.9}$$

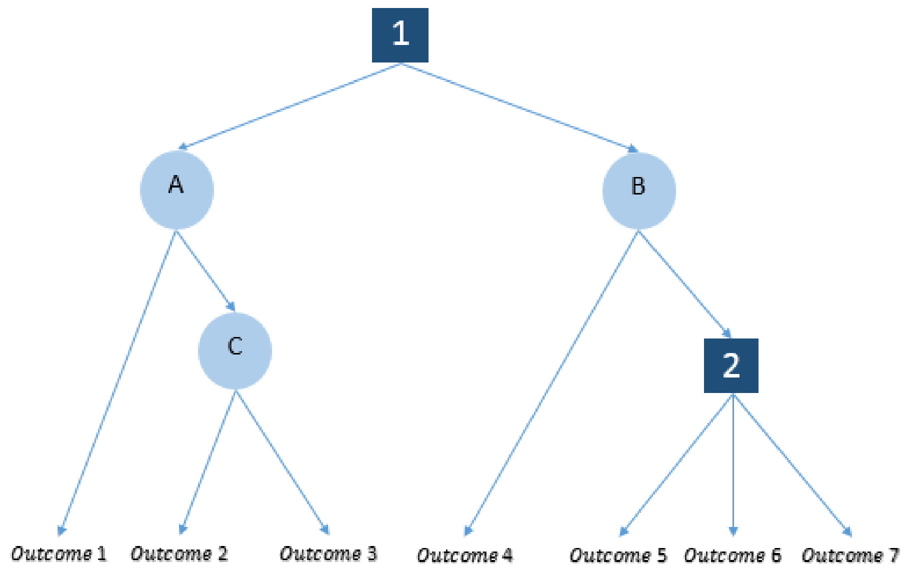

Figure 2.2: Topology of a basic decision tree.

Where I is a function designed to measure node impurity with methods such as entropy, misclassification or Gini index [59]. k_i are values of attribute i, N is the number of observations and v_j is the j-th partition of observations determined by k_j. If all observations belong to the same class, impurity of the leaf nodes is zero, although, for practical reasons, most algorithms prune using an impurity index that should be reached.

2.4.7 Customised Approaches

As previously discussed in the introduction, profiles can be also constructed around a set of highly correlated commonalities wrapped as ”stereotypes”, which can serve as a primitive clustering method. These clusters are created to enhance the user model with additional information, expanding its functionality. The main goal of this approach is to circumvent the obtrusiveness reduction problem by extracting a great part of the information about the users from their transactions. Extracting information from the transactions and applying an appropriate model for the specific domain can result into an accurate prediction of the system usage by users, an ideal scenario for recommender systems that use techniques such as text mining, user ratings, user categorisation through peers identification and collaborative filtering [70, 17, 32, 60].

Furthermore, given the fact that information systems are becoming more ubiquitous over time, user transactions can also be derived from their physical behaviour, storing the users' actions registered by sensors as transactions to include all sorts of non-verbal interactions, such as gestures, emotions reading and audio analysis [84].

In order to collect the information associated with user behaviour-based transactions a whole range of devices can be used. Among the most popular there are cameras of all sorts, movement detectors, recorders and even Microsoft's Kinect sensor which has an increasing popularity over the last years due to its relatively low price and capabilities such as gesture analysis.

2.5 Art-Related Applications

In 2010, Antoniou and Lepouras studied a method to model visitors of museums and art galleries based on different forms of categorisation [4]. They focussed mainly on the visiting styles and combined them with diverse factors related to the user's context. These factors were loosely coupled with the context of the artworks, which will be adapted to this work, since it is not possible to recommend artworks according to the users preferences without including sufficient information about their context.

Most of the methods used to create semantic user profiles from sensors in media related applications can be easily adapted to the planet art. There are several examples in the literature of applications using sensors to track users behaviour in museums and art galleries [51]. Nevertheless there are several considerations that need to be taken into account including the practical, ethical and financial aspects [14, 15]. Along with these considerations, when dealing with museum applications, other correlations can be assessed. In [4] these correlations are further explored along with other important properties that can facilitate the analysis previous to development of museum related applications in the following way:

- Visiting Style and Museum Size: How to determine the museum size and establish how it can affect the visiting style. The most determining factor to deduce a museum size in such a way that it is correlated with the visiting style was according to the number of exhibits, rather than physical size in square meters.

- Visiting Style and Type of Visitor: The only plausible conclusion from this part of the study is that most people visiting in groups would see exhibits important to everyone in the group rather than exhibits related to individual preferences or partial groups.

- Visiting Style and Cognitive Style: Using an inferential statistics-based approach a correlation between these factors was found grouping users in the following categories: Introverts-extroverts and judgers-perceivers.

- Visiting Style and Demographic Factors: Although no significant correlation was found between the visiting style of people according to their age or gender, cultural background seemed to be a determining factor. More studies have to be performed to confirm this hypothesis.

Karaman et al. have been working for over three years on the MNEMOSYNE system, to enhance the museum experience with augmented reality features and personalised multimedia content [51]. One of the most interesting facts in their proposal is that one of the methods they use to retrieve information from the user does not require any special form of interaction, since the user is being observed and monitored with different sensors that passively recollect the users points of interest during its visit. The downside is that the user cannot also state which artworks he enjoyed and which ones he did not since people often spend some time detailing works of art that they find "ugly" or of low interest just to determine the factors that they dislike. Also social pressure can influence the user's behaviour during the visit as well as external factors like crowded places near points of interest. The system implementation also seems too complex for places with high concurrency since the level of adaptation is relatively low. The authors of the study propose to enhance the museum experience integrating user profiling based on passive observation and using the collected information in a recommender system set on interactive tables. At the end of their visit, users can have an overview of their experience on their mobile phones using a custom application.

The MNEMOSYNE system exploits the available tools that museums generally offer for security purposes such as surveillance cameras to monitor the users' behaviour in order to build their profile. An operator can easily calibrate the position of the artworks as registered by the cameras, that is to map the physical museum. To identify users the system uses a pedestrian detector, with which they obtain a set of N person bounding boxes. This technology was introduce by Bimbo et al. from the same research team [13].

The descriptor of a person bounding box is defined as follows:

$$d_i = \left\{ d_i^a, d_i^s, d_i^t, d_i^c \right\}, for\, i \in \{1, ..., N\} \tag{2.10}$$

where d_i^a is an appearance descriptor defined by HS colour histograms computed on overlapping stripes and the histogram of oriented gradients descriptor (HoG) and RGB [51, 26], d_i^s is the absolute position of the person detection on the ground plane, as described in [52]:

$$d_i^s = (d_i^x, d_i^y) \tag{2.11}$$

d_i^t represents the temporal dimension as an integer timestamp and d_i^c identifies the camera registering the input. All video streams are synchronised such that d_i^t and d_j^t are comparable. The key to the profile creation is the association of the detected inputs per user $D = \{d_i | i = 1...N\}$. The algorithm relies on the distance between a model cluster m_j and a detection description d_j computed as follows:

$$dist(m_j, d_j) = (1 - \alpha - \beta) \times ||m_j^a - d_i^a||_2 (appearance\, contribution) \tag{2.12}$$

$$+\alpha \times dist_w(m_j^s, d_i^s, w_s)(spatial\, contribution) \tag{2.13}$$

$$+\beta \times dist_w(m_j^t, d_i^t, w_t)(temporal\, contribution) \tag{2.14}$$

where $dist_w(x, y, w)$ is the windowed L2 distance:

$$dist_w(x, y, w) = min(\frac{||x - y||_2}{w}, 1) \tag{2.15}$$

Parameters w_s and w_t are the spatial and temporal window around observations respectively. α and β are the weights of the distribution of spatial and temporal distances, to the overall distance calculation defined in such a way that $\alpha, \beta \in [0, 1]\, and\, \alpha + \beta < 1$. There exists also a control variable δ such that a detection is associated with a model if its distance to the model is less than δ. In order for the system to associate a detection to a real model, it must first accumulate at least τ detections in a temporary model. The readings are also constrained such as multiple detections from the same camera to the same model at the same time are forbidden.

User profiles are built instantly when the visitors enter the interactive table area, and are updated with each detected association made to its model. Every time the user gets back at the interactive table after spending time viewing a set of artworks, its profile will be updated.

As expected, in order to provide recommendations, the MNEMOSYNE system also uses a hybrid recommendation engine, developed as web servlets exposing several web services through a REST interface. The semantic search engine exploits the capabilities of an RDF ontology with a great amount of contextual information on the artworks used to improve the recommendations.

As mentioned before, it is part of the purpose of this work to expand some of the capabilities provided by the MNEMOSYNE system, bringing art recommendations to broader datasets. Along with that, more accurate recommendations will be done if the users are allowed to explicitly express their interests without the necessity of being overly obtrusive. This will also help reduce the implementation complexity in a first stage by avoiding the usage of advanced gesture recognition techniques.

Chapter 3

Analysis

The research training surrounding this work is done by analysing the main features of the previously discussed techniques to build recommender systems and perform semantic profiling. The main advantages and disadvantages of the aforesaid techniques will be summarised and later presented in two tables to facilitate the comparison while providing a general overview.

3.1 Recommender Systems

A key factor for the success of information systems is user satisfaction [29, 43], a problem that is being tackled from several angles by the aforementioned research. The ability to use all the information available on the user of an information system to provide tailored content has helped improve user satisfaction in a big part of today's most used information systems worldwide. This feature becomes especially useful in systems that manage great amounts of information that are not easy to browse nor explore [72]. Some studies have shown the way in which machine learning techniques applied to rating/grading systems have exceeded human capability in text classification accuracy [57]. One of the most remarkable and profitable examples is Amazon.com, which personalises the on-line store for each registered costumer [48].

Besides the latter information systems can also provide non personalised recommendations, which can be very useful in scenarios where there is not much information available about a user in a particular point in time, yet it is desirable to provide a glimpse of the available content that this user might be searching. For this scenario a statistical approach is considered, such that the information available on the preferences of all the other users in the system is used to guess the items that the anonymous user is searching [48]. The accuracy of these recommendations is still under discussion, nevertheless it has been noted that individuals tend to rely on recommendations made by peers [60]. The most common scenarios are recommendations of books and films. This non personalised approach does nothing more than eliminate the man in the middle, providing in a synthesised manner the most popular items to the current user.

3.2 Semantic Profiling

Semantic user profiling is becoming a de facto standard to categorise users and items, this is due to the flexibility and expressiveness provided by the techniques that implement it.

Using ontological approaches to create user profiles provides several advantages. One of them is that the profiles can be easily visualised, as ontologies in general are built using a fairly understandable terminology. Binary profiles are normally represented as term vector spaces or in neural network patterns, which are difficult to understand. Also, with the ontological approach users can provide feedback on their own profiles to improve their accuracy.

The daily learner [12] uses a k-NN clustering algorithm for short term profiling and a naive version of Bayesian networks for long term profiling, reporting a precision level of 33% and recall of 29%, for the top 4 recommendations monitoring the selected stories for users to read. There is also some research to expand profiles using contextual information in task modelling [19]. Next, some of the most popular approaches to the classification problem are presented.

3.2.1 Nearest Neighbours

Due to their inherent simplicity, kNNs are considered lazy learners, and leave many decisions to the classification step making it computationally expensive to classify unknown records. Eager learners such as decision trees or rule-based systems explicitly build models to avoid such detriment [59]. kNNs are adaptable as they do not require to learn or maintain models, but such adaptability is expensive in the sense that the neighbourhoods and similarity matrices need to be recomputed in the classification process [25].

3.2.2 Rule-based Classifiers

The main advantages of rule-based classifiers are the expressiveness, ease to interpret and efficiency [59]. The expressiveness is gained through the lack of transformations of data combined with the fact that they are symbolic. The efficiency resides on the inherent simplicity of the categorisation method. For recommender systems explicitly, it is hard to achieve a good implementation with rule-based classifiers, since prior knowledge on the decision making process is assumed unless rules are derived from other models such as decision trees. Nonetheless rule-based systems can be used to improve efficiency in recommender systems by injecting such prior knowledge or a set of business rules [59, 2].

3.2.3 Clustering

Clustering techniques are used to reduced the complexity of calculating the distances of objects when scaling is needed for a classifier. This is done mainly because even reducing the dimensionality of features to take into account in the comparison of items, the amount of operations needed to calculate their distances might still be significantly high [59]. An example of this is trying to find the k-nearest neighbours of an item, or more closely related to the purpose of this research, the k most similar items to the one under analysis.

3.2.4 Bayesian Classifiers and Bayesian Networks

Due to the models used and the contextual information gathered, Bayesian classifiers and Bayesian networks can easily handle incomplete data and behave robustly in the presence of model overfitting [59]. They can also support extensibility to multidimensional classes [35].

Nevertheless when dependencies of data attributes need to be encoded the complexity increases significantly [59].

3.2.5 Artificial Neural Networks

Non-linear classification tasks can be performed by artificial neural networks. They can be efficient and resilient to loss of parts of the network. Nonetheless, their main disadvantage is the complexity of design, as it is hard to come up with an ideal topology, and once established, it will serve as a lower boundary for the classification error [59]. As they implement an approach similar to a classic black box, there is no support for semantic learning inside the network.

3.2.6 Decision Trees

Basic decision trees are generally easy to implement and extremely fast at classifying new items, they can also be given a certain set of rules to be interpreted and still maintain their accuracy. Nevertheless it is almost impossible to include all the variables involved in the decision making process [59].

3.2.7 Customised Approaches

The use of ontologies can be helpful to deal with the information overload generated by sensors readings, taking advantage of compact and efficient information representations. Semantic user profiling enables structures of user profiles with high levels of granularity and accurate data annotation. These approaches explore the semantics of user behaviour, in face of items directly presented to them, but two big challenges arise from the usage of such techniques: First, the presence of structured domain knowledge that usually has to be constructed manually and second, precise semantic mappings with data annotations are required.

Another important challenge that semantic user profiling based on user behaviour faces is the need to understand heterogeneous data from different kinds of sensors, to express it in a machine-readable format. A key factor to succeed in these tasks is the ontology definition, as it has to be suitable for the domain and the intended system capabilities. These ontologies need to provide support for the creation of potent knowledge bases to capture the intended user profiles. Several languages and tools are available to facilitate the creation and standardisation of such ontologies like RDFs, OWL, SKOS and FOAF among others, and there exist general-purpose ontologies oriented to provide cross-discipline knowledge bases like DBPedia [5].

Lastly, the purpose of identifying the most fitting modelling schema to use, is to determine the precise syntax to represent not only the users' preferences but also their relations while supporting contextual information to establish hierarchies of preferences and interactions [84, 54].

To summarise, a set of tasks and sub tasks desired in order to create semantic user profiles in an ubiquitous environment have been identified:

- *Identify and understand the data relative to a user and its preferences:*
  - Information extraction (preferably as unobtrusively as possible).
  - Information analysis and dissemination of semantic value.

  - Data refining and domain knowledge generation.

- *Expand user preferences knowledge:*
  - Analysis of semantically related information.
  - World knowledge integration based on the ontological knowledge base.
  - Publicly available information exploitation (when possible).

- *Formatting obtained knowledge in a machine-readable format:*
  - Align obtained information with the integrated ontology.
  - Produce compact and meaningful semantics.

- *Incorporation of learnt user behaviour:*
  - Rating estimation.
  - Discover preferences relations.

- *Knowledge enhancement:*
  - Detection of behavioural patterns.
  - User clustering.

- *Further analysis:*
  - Assessment and enhancement of user identities.
  - Understanding behavioural information to adapt profiles (re-engineering).

Table 1 illustrates a comparison between the most popular techniques discussed to address classification in the context of heuristic-based recommender systems. In table 2 the main advantages and disadvantages of classification techniques for model-based systems are displayed. Both are divided into the different categories of recommender systems: content-based, collaborative filtering and hybrid.

Table 1 Recommender systems overview - Heuristic based systems

Approach	Techniques	Advantages	Disadvantages
Content-based	TF-IDF	Facilitates keywords weighting [75]	New user problem [75]
	Clustering	Simplicity and efficiency [59]	Not easily scalable, assumptions need to be made on the data [59]
Collaborative	Nearest Neighbour	Simplifies rating estimation procedure, adaptability as it does not requite to learn nor maintain a model [25]	Computation overhead as it needs to recompute the neighbourhoods and similarity matrices [25]
	Clustering	Simplicity and efficiency [59]	Computationally expensive [69]
	Graph theory	Facilitates users categorisation	New item problem [86]
Hybrid	Linear combination of predicted ratings	Improved accuracy	Complexity of combining techniques
	Various voting schemes	Facilitates item categorisation	User intentional biasing of ratings
	Incorporing one component as a part of the heuristic for the other	Support for unobtrusive rating techniques	Complexity

Figure 3.1: Advantages and disadvantages of the different heuristic-based techniques for recommender systems.

Table 2 Recommender systems overview - Model based systems

Approach	Techniques	Advantages	Disadvantages
Content-based	Bayesian Classifiers	Supports extensibility to multidimensional classes [36]	The independence assumption may not hold as attributes may be correlated [59]
	Clustering	Uses models based on statistical analysis over the predefined heuristics [12]	Generally lack extensibility [69]
	Decision trees	Facilitates implementation of utility function	Limited attribute-based profiles
	Artificial neural networks	Efficiency and resiliency to network failures [59]	Design complexity and scalability [59, 27]. No support for semantics [59]
Collaborative	Bayesian networks	Supports extensibility to multidimensional classes and incomplete data [35]	Extra complexity of encoding the dependencies of data attributes [59]
	Clustering	Uses models based on statistical analysis over the predefined heuristics	Generally lack extensibility
	Artificial neural networks	Efficiency and resiliency to network failures [59]	Design complexity and scalability [59, 27]. No support for semantics [59]
	Linear regression	Improved accuracy	Complexity
	Probabilistic models	Improved trustworthyness	Complexity of dealing with data sparsity
Hybrid	Incorporating one component as a part of the model for the other	Simplifies users and items categorisation	Integration of predefined models
	Building one unifying model	Flexibility of implementation	Complexity of model creations

Figure 3.2: Advantages and disadvantages of the different model-based techniques for recommender systems.

3.3 Summary

After examining the state-of-the-art in artwork-related applications for the museum environment and the research trends in the areas of recommender systems and semantic profiling, it is possible to assume that the best way to address the obtrusiveness vs. accuracy problem in information extraction for profile building is through simple explicit interactions. The techniques explored for instance in the MNEMOSYNE system seem not only overly complex but also inaccurate. Performing large scale studies they realised that after observing users for several hours it is likely that several users that look alike are bound to the same profile. Also explicitly asking users about their preferences during their visits is overly obtrusive. Tagging artworks with RFID tags bound to the user's profile is a reasonable middle ground that allows to accurately extract the users' preferences without bothering the users during their visits.

When it comes to user categorisation for recommending features where so little information is known about the user, it is better to use a network approach, since the amount of categories is not as relevant as the ability to find similar profiles. The recommendation engine of the

ArtVis system will do two user categorisation tasks, the first one related to the user expertise assessment and the second one to find the k closest profiles in terms of preferences and users' similarities. For the first categorisation, a binary Bayesian classifier will be used, using the available information about the user and based on its preferences the system will try to associate its profile to the expert model or the amateur model. The second categorisation will be used to find the closest profiles to the current user in order to include their favourite artworks, artists, schools and artwork types in the recommendations.

The recommendations will be done following a hybrid technique, combining features from collaborative-filtering and content-based approaches. This with the main purpose of helping amateur users to find interesting artworks that may apply to their preferences and expert users to bookmark and keep track of artworks that may seem interesting to them.

Chapter 4

The ArtVis System

Exploring large data sets of artworks is a complex task. There are several compilations available but few of them are free and available to reproduce and even less provide an automatised way to access the data in order to be exported. Most of the available sources also contain several errors and imprecisions as it is very difficult to trace the origin of artworks in an accurate manner. Taking all of this into account it has been decided to use the Web Gallery of Art to feed the ArtVis system database [1]. After detecting some of the aforementioned inconsistencies on the data it has also been decided to apply a set of transformations to normalise and fix the schema making it better suited for the purpose of the ArtVis system [33].

In the next sections an overview of the original ArtVis system will be presented followed by the description of the modifications applied to the system in order to provide the recommending capabilities based on semantic profiling techniques.

4.1 Overview of the Original System

The ArtVis system is an application that combines advanced visualisation techniques with a tangible user interface to explore a large data set based on the Web Gallery of Art, which contains information over approximately 30 000 artworks including paintings, sculptures, architectural masterpieces, etc. The Web Gallery of Art is part of a greater project designed to develop new techniques to use the web in virtual education. The content of the Web Gallery of Art is publicly available and it can be downloaded in text format.

The system is originally built in Java language with the Prefuse visualisation toolkit [2] for the GUI. The data is stored in a MySQL database that contains materialised views to support the construction of prefuse tables to analyse data and create the visualisations.

Information visualisation theory is ”the use of computer-supported, interactive visual representations of data to amplify cognition” [21]. In ArtVis, applying information visualisation techniques allows users to browse information from different dimensions according to their interests, that is, filtering by time-spans, geographic location of the artworks, geographic location of the authors who created the artworks, etc.

Enabling users to explore information with such techniques facilitates users to generate,

[1]Krén, Emil and Marx, Daniel The Web Gallery of Art: http://www.wga.hu/index1.html Last checked: June 2014

[2]Prefuse: http:http://prefuse.org/ Last checked: June 2014

confirm and reject hypotheses by focussing only on pieces of information relevant to those hypotheses in larger sources. Besides the latter it is assumed that correctly applying visualisation techniques can inspire curiosity on users in unknown domains. Artvis contains information about European artworks produced between the 11^{th} and 19^{th} century, presented in such a way that users can easily explore artworks that might seem interesting to them.

The ArtVis interface follows a classification technique for visualisations proposed by Keim in 2002 [55], which considers three main factors:

- Data types can be unidimensional or multidimensional, enabling different forms of representations including textual, hierarchical or in graphs.
- Visualisation techniques are dependent on the later feature, varying greatly from 2D, 3D to n-dimensional data.
- The relation of the users with the visualisations are described by Interaction and Distortion, usually applied by altering the projection, filtering, zooming, linking and brushing.

These classifications are done in order to define the best way to apply the visualisation techniques, aiming to increase the value of the visualisation, or in terms of the previously given definition, that is amplify cognition.

4.2 The Graphical User Interface

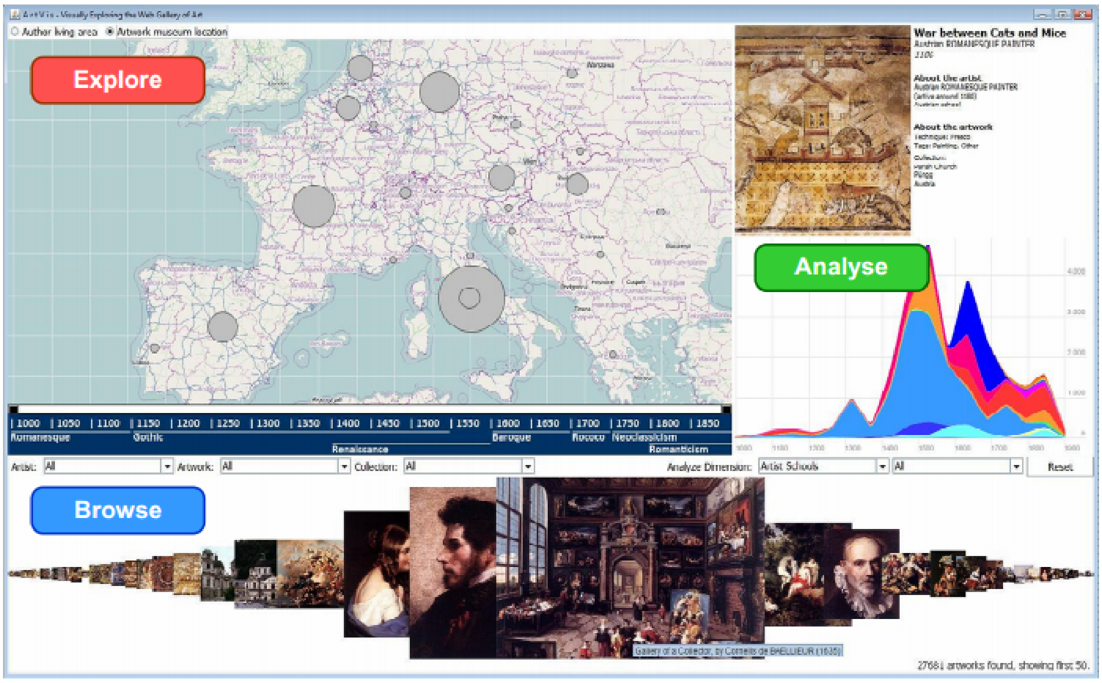

Figure 4.1: ArtVis Interface

As seen in Figure 4.1, the system interface is divided in four panels: The Explore Panel, the Analyse Panel, the Browse Panel and the Detail Panel, each serving a different but complementary purpose in the task of providing users the tools to navigate through the aforementioned dataset as follows:

- Explore Panel: Includes a zoomable 2D map representation and a time range explorer bar to filter in the spatial and time dimensions. These filters are applied in order to find artworks that the users may be interested in. The artworks are initially grouped by location and displayed in the map representation as bubbles, selecting a bubble will load a sample of the filtered data set in the browse panel. Figure 4.2 shows an example of the explore panel as loaded by default without filters.

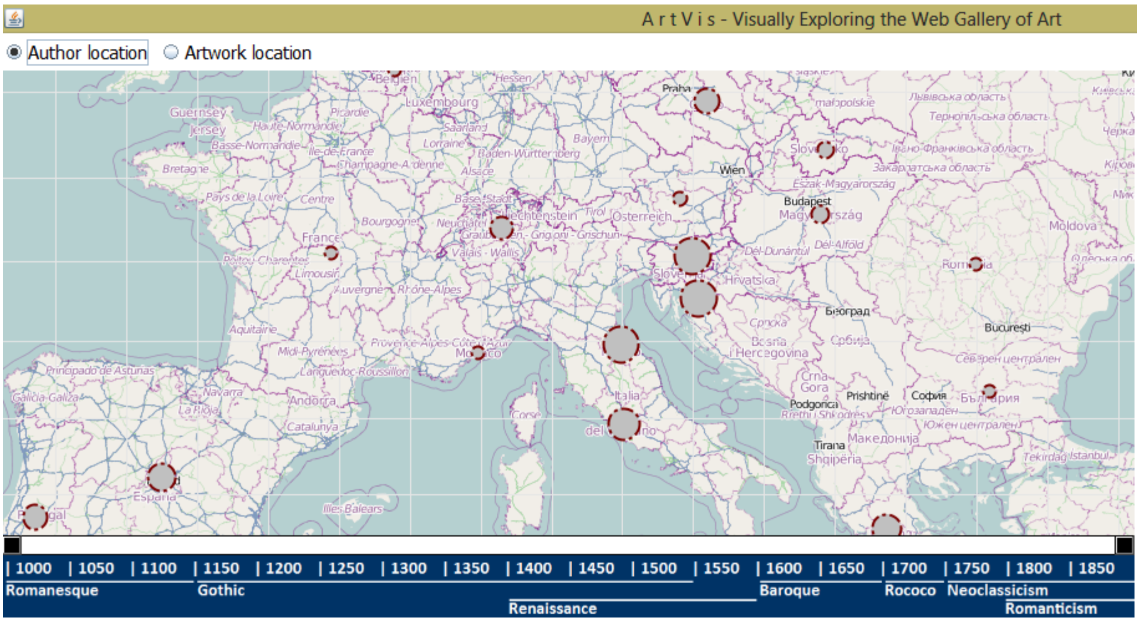

Figure 4.2: The Explore Panel

The top sub panel shows two radio buttons to manipulate the visualisation. The first one consist of emphasising the location of the artworks shown according to the artwork current location and the second one according to the author's birth area. This second option uses the label "Author living area" which will be later modified to avoid confusion.

The time range selector allows users to select a timespan between the 11^{th} and the 19^{th} century. The artworks shown will then be filtered accordingly and the bubbles resized to the amount of artworks scaled to the selected timespan. As an additional visual help, the time range selector contains a set of labels at the bottom which highlight the main artwork styles of the related period.

- Browse Panel: Displayed at the bottom part of the interface, the browse panel provides an insight of the artworks contained in the timespan and geographical locations selected in the explore panel. As seen in Figure 4.3, it uses a fish-eye technique to display up to 50 artworks and browse through them to have an overview of the related artworks. The browse panel also contains three combo boxes from which users can further filter artworks by artist, artwork name or collection.

- Analyse Panel: This panel displays the changes in the amount of artworks created over the time through different dimensions. It contains a stacked area chart and two combo boxes to filter according to art schools, inspiration, and artwork form.

  Figure 4.4 shows the analyse panel in its default state. Selecting a value from the Dimension combo box will filter the information accordingly, giving the corresponding values

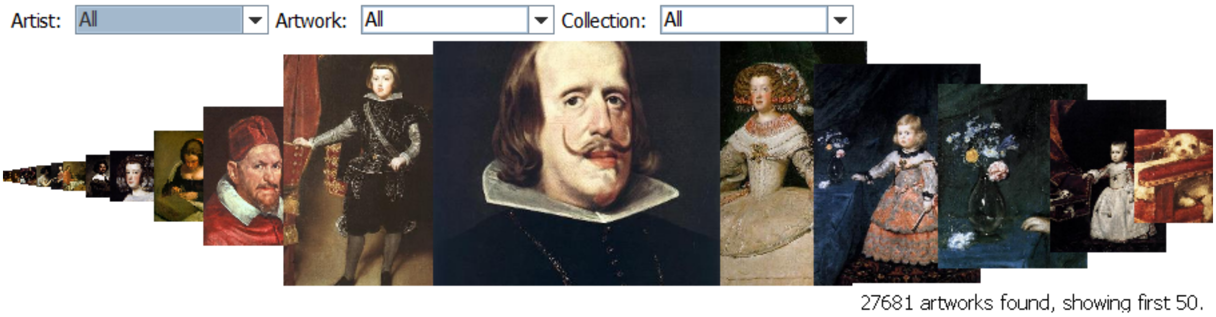

Figure 4.3: The Browse Panel

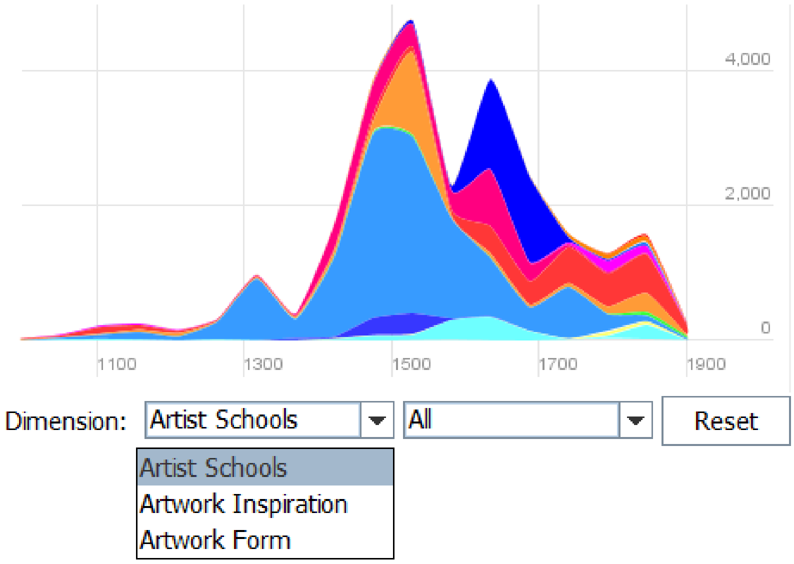

Figure 4.4: The Analyse Panel

to perform detailed exploration in the second combo box. For instance, selecting artist schools in the first combo box will load all the available European schools, enabling users to choose the evolution of a certain school over time, the same as selecting inspiration and further filtering by religious will show how the evolution of religious artworks has been from the 11^{th} until the 19^{th} century.

- Detail Panel: Provides detailed information on selected artworks from the browse panel, including an image of the artwork, year, author, school, collection, technique, size and tags.

  In Figure 4.5 the Detail Panel is displayed showing the artwork ”Portrait of Philip IV” by Diego Rodríguez de Silva y Velásquez from the Spanish school, painted in 1652.

Figure 4.5: The Detail Panel

4.3 The Tangible User Interface

The ArtVis tangible user interface is done through a setup of hardware components named Phidgets, which are low cost USB controlled artifacts that can be used as sensors to register interactions in an ubiquitous environment [39]. The Phidgets manufacturer provides a publicly available API to control the devices in wrappers for various programming languages such as Java, which is used in the implementation of the ArtVis system.

Figure 4.6: The tangible user interface used to interact with the system.

As displayed on 4.6, the original ArtVis system prototype contains a picture frame on a painter's easel, a painter's box and palette, paint tubes, postcards and some other objects enhanced with RFID tags. These objects are used either to control the system interface or to emulate tagged objects in a museum or art gallery that can interact with the system in various ways. The output screen has been embedded in the picture frame to emulate the creative environment of artists, and encourage people to use the interface.

The painter's palette contains controls for the main features provided in the GUI. The panning of the map visualisation is controlled with a small joystick and the zoom level by a rotation sensor. The time range slider can also be controlled using a slider sensor while another one is used to control the fish-eye view of the artworks in the browse panel.

Finally the palette (Figure 4.7) contains an RFID reader on the bottom right corner, which

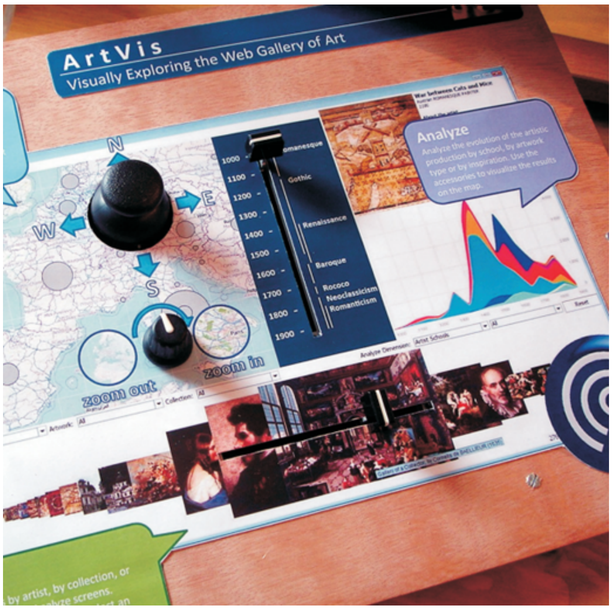

Figure 4.7: The painter's palette contains a small set of phidgets used to control the GUI of the ArtVis system.

allows users to extract information from several objects. A school can be selected using a tagged paint tube, artwork forms can also be detected with some of the aforementioned objects like a small sculpture or a painting, an inspiration can be selected approaching a postcard from a predefined set and the modality of artwork location (collection or author living area) can also be switched using one of the tagged small paintings set up for this purpose.

Other specific RFID tags can be used in other objects of the museum environment to select collections for visualisation and filtering, making the tangible interface serve the purpose of suggesting queries through tangible objects. Each of this objects will have a semantic value or meaning when input signals are detected from it.

4.4 The Database

The information in the Web Gallery of Art is not necessarily stored in a homogeneous, normalised way, which eventually led to a series of transformations conducting to the design depicted in figure 4.8.

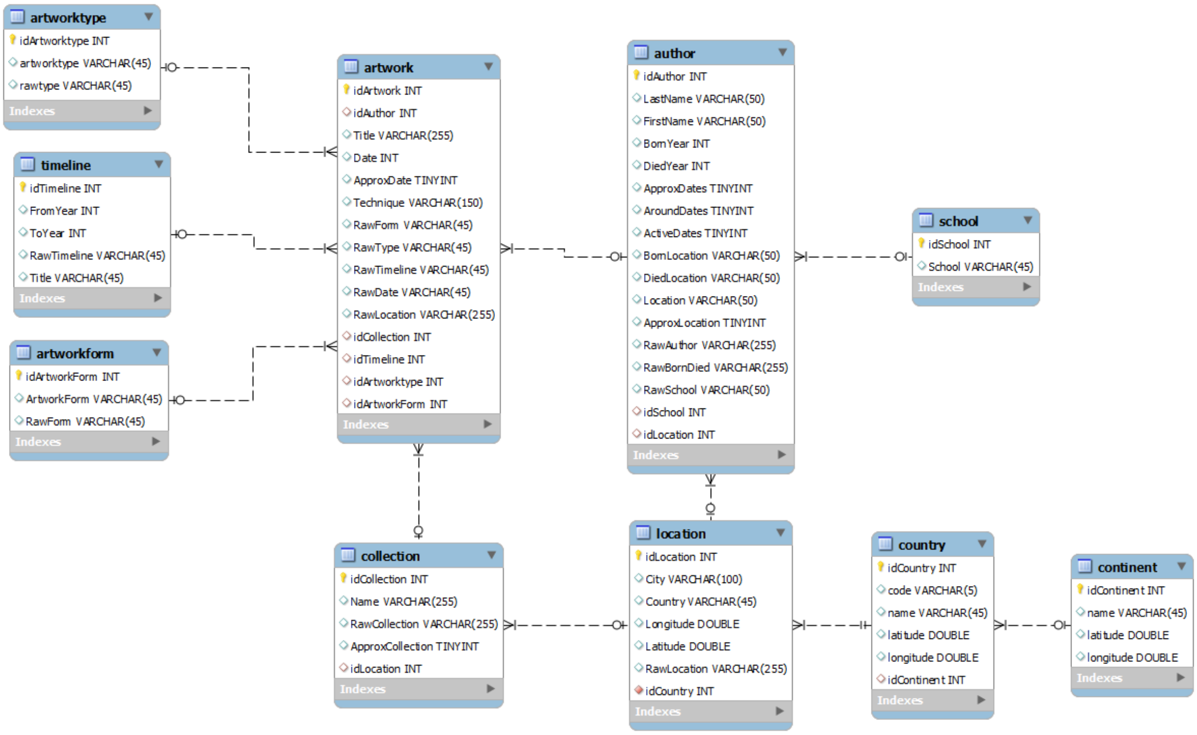

Figure 4.8: Database schema.

4.5 System Enhancement

This section focusses in the features added to ArtVis system during this thesis. The modifications are presented categorised in four groups:

- User information and preferences extraction related
- Recommendation engine implementation
- Interface adaptation
- Information storage and representation

4.5.1 User Information and Preferences Extraction

As depicted in the previous chapter, an interesting way to address the problem of choosing between usability and accuracy in terms of collecting user information and preferences was to think of a way to have explicit interactions with the system without forcing the user to make statements or fill in forms. As explained in the original system overview, ArtVis already had support for phidgets and more specifically for RFID tags, so an authentication mechanism was implemented using such hardware components. At the time of registration, users will be assigned an RFID tag. The id number of the tag will be used from that moment on to authenticate the user into the system. Registration will not be mandatory and as unauthenticated users approach the interactive tables the system will load preferences based on collective wisdom. The same RFID tags used to authenticate the user into the system will also be used to tag artworks expressing interest and linking those interests to the user's profile.

Whenever users interact with the browse panel and select artworks new interactions will be created to keep track of other artworks that the user might have found interesting. A compilation of these interactions is used to display artworks to the user in a certain way that interesting and related artworks will be discovered with more ease. In the physical museum, the artworks will contain RFID tag readers and as the users approach their tags to those readers they will have expressed their interest on the related artworks creating more interactions to be used by the recommendation engine to provide similar content. In order to process this interaction, a set of stored procedures were created in the database and are triggered when a user authenticates into the system with an RFID tag. The procedures implement a hybrid recommender algorithm to rank the artworks and sort them so the user gets a better overview of the artworks that it might consider interesting. This guarantees that with each interaction the system will learn about the user preferences improving the recommendations as the user tags artworks on the spot or selects them on the GUI for better visualisation.

To store the interactions in the database the system implements a DAO pattern (Data Access Objects) [66], providing a modular solution to the data access problem. The rest of the objects are loaded using Prefuse's built in data access capabilities as no writing in the database is necessary in any of the other possible interactions. Due to the visualisation tool's limitations, non-normalised views were required to access the data units necessary for the visualisations mapped through table objects.

4.5.2 Recommendation Engine Implementation

In order to provide users a better, personalised insight on the artworks of the data set, the system was enhanced with recommendation capabilities based on semantic user profiling techniques. To accomplish such a task, a hybrid recommendation algorithm was implemented. The algorithm starts whenever a user either resets the interface or gets authenticated in the system using its RFID tag. If the user resets the interface a default profile will be assigned bringing recommendations based on collaborative filtering and a collective wisdom approach. If the user authenticates, the system will assess its expertise and based on that it will assign an expert model or an amateur model. Whenever the expert model is selected a content-based approached will be used to recommend content based on the users' preferences. If the amateur model is selected, a hybrid approach that combines collaborative filtering and content-based techniques will rank the artworks before filtering according to the settings the user sets using the interface. More details on the algorithm implementation will be presented in Chapter 6.

4.5.3 Interface Adaptation

- Explore panel: The top sub panel shows two different options to drastically manipulate the visualisation. The first one consists of emphasising the location of the artworks shown according to either their current location or their author's living area. The second one manipulates the size of the bubbles to emphasise the amount of artworks in the filter or their relevance according to the user's profile.

  The artworks shown will then be filtered accordingly and when the recommendation is enabled the bubbles will be resized to the average relevance scaled to the selected timespan. As an additional visual help, the time range selector contains a set of labels at the bottom which highlights the main artwork styles of the related period.

- Browse Panel: The artworks are displayed in order of relevance, which is calculated using the previously mentioned hybrid recommendation algorithm. In case of authenticated users the order will be determined by the user profile and its previous interactions with the system and in case of a non authenticated user the relevance will be determined through collective wisdom, giving more importance to the most popular artworks in the system.

  As is the goal of recommender systems, the browse panel will show similar artworks as they share some common features extracted from the preferences of users.

4.5.4 Information Storage and Representation

Part of the success in integrating new capabilities to the system resides in its inherent simplicity and modularity. For instance, to allow the recommendations features to work the only changes made to the data schema were the users table and the interactions table, which link the preferences expressed by a user to a specific work of art.

Figure 4.9 displays the new schema highlighting the changes made to include the user information and the preferences that support the recommendations capabilities. As mentioned earlier in this document, the algorithm implementation was also done at a database level using stored procedures to improve the performance avoiding the use of complex layers and other heavy components to interact with the data.

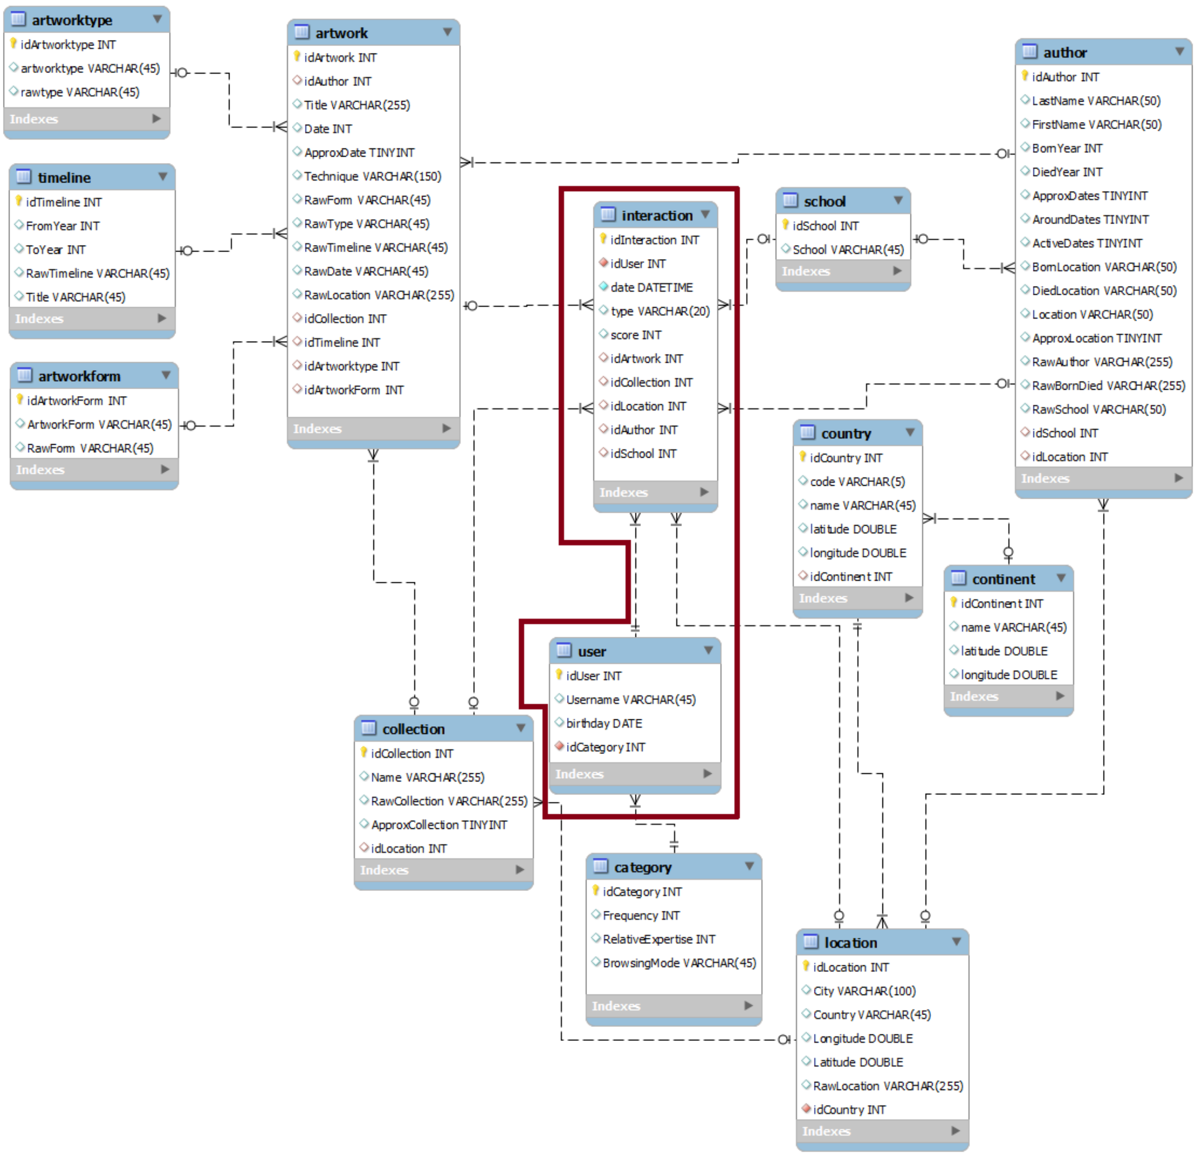

Figure 4.9: Database schema.

Chapter 5

Scenarios

During the time of exponential growth for commercial applications, a lot of research was done to improve capabilities of recommender systems. This is due to the profitable potential that good recommendations of items can have. For example in platforms such as amazon.com, providing good recommendations will be an immediate source of income. Some of the most interesting results of such research were shown in the related work section where the state of the art techniques in recommender systems were addressed.

Application developers who want to implement recommender system capabilities to their applications have now a large variety of techniques and algorithms to choose from. Given those techniques it is up to the system designer and its domain to select the right approach to implement. But, selecting the most suitable way to implement recommendation capabilities given the scope and context of the application can be a very complex task. Such decisions are generally taken based on experiments, comparing the performance of an arbitrary amount of options [59].

The most suitable approach can be selected depending on the structural constraints there might be, such as a predefined type, reliability, or availability of the data. One of the strongest tendencies is evaluation based on predictability, that is, the ability to predict a user's preferences. Nevertheless, predictability although utterly important is not always the only determining factor to evaluate a recommendation system, especially since there are people who use information systems with the purpose of finding new and interesting content, content that they do not know about instead of an anticipation of their preferences. Users may also be interested in exploring diverse items rapidly [59].

There are three kinds of performance assessment methods intended for this system, the first one consists of an offline study, which uses the data set and protocols that model the user behaviour. These protocols are based on the scenarios that will be defined in the next section and the actual results of the recommendations made. The accuracy of the recommendations will be assessed according to the predefined arbitrary preferences for each scenario. The second method is a user study, which would consist of a small set of users performing tasks in the system and then asking them for feedback. This is typically done by asking questions about their experience with the system, ideal to test usability features around functional requirements. This study was already made for the original version of the ArtVis system and the results can be seen on [33]. In order to complement such study, during this work one additional user study was made. The goal of this additional study was to get an insight from experts in the field of art history to direct the design of the recommendation algorithm and the user expertise assessment.

Lastly, a large scale experiment can be performed on a deployed version of the system. This method is called an online experiment, such experiments are meant to evaluate the accuracy and performance of the recommender engine on a large scale.

Nevertheless, as this is considered to be an initial stage of the system, the last kind of assessment method will not be applied at this first stage. This part will be set aside as part of the future work to tune up the recommendation engine. For the offline experiment some basic scenarios were created emulating an estimated use of the system by regular users. These scenarios will also help to build models around them to test the user categorisation feature.

5.1 Scenarios Overview

In order to assess the performance of the system and its recommendation capabilities, a set of scenarios were defined, from which additional challenges arose. Besides the research questions the performance assessment and implementation of the scenarios will be based on the following questions:

- What information can be extracted from a user based only on his interactions with the system?
- How can the latter be combined with the available user information?
- How can the users be categorised in a relatively unexplored environment by recommender systems?
- What level of accuracy can be reached using artificial scenarios to test a recommender system?

In order to register on the system, some basic information of the user is asked to build its profile, including name, birthday and home town. This information could also be extracted automatically from an identity card (eId) with the appropriate card reader. Next some common use scenarios of the ArtVis system with a semantic user profiling mechanism implemented are described:

5.1.1 The Student

Consider a student of any discipline that regularly attends exhibitions in museums and art galleries around the world. In the beginning, the system has no information whatsoever on this student, so his profile will be built from scratch. After the student starts attending new events the system begins to collect information about his preferences, which will be computed in order to manipulate the semantic queries done through the interactions with the Phidgets to prioritise artworks that he has already seen and liked, or that are closely related to those. The system will also take into account artworks that are popular to other users that the algorithm has considered similar to him.

Taking into account that this user travels and attends exhibitions in different places, the system could also assist him in finding interesting artworks in his surroundings during his travels, including artworks from his favourite artists that are being displayed in his next travel location or from related artists that may have a similar style and that he would like to see.

Table 1 Student profile - Interactions example

Date	Type	Artwork	Author	Collection	School	Timeline	Type	Form	Location
7/12/13	Reader	The Birth of Venus	Alexandre CABANEL	M. D'Orsay	French	1900	Mythological	Painting	Paris
7/12/13	Reader	The Lightning	Alexandre ANTIGNA	M. D'Orsay	French	1850	Genre	Painting	Paris
7/12/13	Reader	Strength	Antoine-Louis BARYE	M. D'Orsay	French	1850	Mythological	Sculpture	Paris
7/12/13	Reader	Woman Bitten by a Snake	Auguste CLESINGER	M. D'Orsay	French	1900	Genre	Sculpture	Paris
7/12/13	Reader	Pan with Bear Cubs	Emmanuel FREMIET	M. D'Orsay	French	1900	Mythological	Sculpture	Paris
7/12/13	Reader	Victory	Francisque-Joseph DURET	M. D'Orsay	French	1850	Mythological	Sculpture	Paris
7/12/13	Reader	The Source	Jean-Auguste-Dominique INGRES	M. D'Orsay	French	1850	Mythological	Painting	Paris
7/12/13	Reader	Cenotaph of the Gracchi	Jean-Baptiste-Claude-Eugene GUILLAUME	M. D'Orsay	French	1850	Historical	Sculpture	Paris
7/12/13	Reader	Despair	Jean-Joseph PERRAUD	M. D'Orsay	French	1900	Mythological	Sculpture	Paris
7/12/13	Reader	Strength	Jean-Louis-Nicolas JALEY	M. D'Orsay	French	1850	Mythological	Sculpture	Paris
15/12/13	Reader	Epitaph for the Nun Janne Colijns	Flemish UNKNOWN MASTER	Staatliche Museen	Flemish	1500	Religious	Painting	Berlin
15/12/13	Reader	Venus at Vulcan's Forge	Frans FLORIS	Staatliche Museen	Flemish	1600	Mythological	Painting	Berlin
15/12/13	Reader	Portrait of a Woman	Frans HALS	Staatliche Museen	Dutch	1650	Portrait	Painting	Berlin
15/12/13	Reader	Allegory of the Arts	Hans I ROTTENHAMMER	Staatliche Museen	German	1600	Mythological	Painting	Berlin
15/12/13	Reader	Rinaldo and Armida	Giovanni Battista TIEPOLO	Staatliche Museen	Italian	1750	Mythological	Painting	Berlin
15/12/13	Reader	Apollo, Pan, and Marsyas	Johann Karl LOTH	Staatliche Museen	German	1700	Mythological	Painting	Berlin

Figure 5.1: Sample of the interactions made by the student profile.

Table 1 depicts a sample of the possible set of interactions that a person with this profile may have. The interactions shown in table 1 are based on a visit this user made to two museums, Musée D'Orsay in Paris the 7th of December 2013 and Staatliche Museen in Berlin the 15th of December 2013. All the interactions seen were registered using the reader, which means all of them were tags on artworks using the RFID readers.

After analysing the interactions it is possible to conclude that the user has a certain preference for mythological artworks, and a particular affinity for the French and Italian schools. In the ideal scenario, if this user approaches an interactive table with the ArtVis system and authenticates using his tag, the system should be able to recommend content similar to the one the user has shown interest in, that is mostly mythological artworks from the French and Italian school among others that may turn out from other interaction the user has with the system. The more interactions the user has the more accurate these recommendations will be, since patterns will be detected with more ease. This user is most likely not an expert in art, therefore, approaching the interactive tables will probably give him a better insight of what he might be interested in as in some cases inexperienced users are more open to discover things about their preferences that they might not know before.

5.1.2 The Art Professor

Now, taking into account a more challenging profile, like an art professor, from whom nothing is known, but, due to the nature of his work it can be assumed that he has visited several exhibitions before. It is very natural for such kind of users to try and find in the system the artworks that he may consider the most interesting, as well as the ones that relate the most with his lectures and current research or work. After finding such artworks once, it is very possible that this user will want to bookmark them to find them easily again and browse them, reasons for which he would tag them knowing in advance how the ArtVis system works.

This way, without asking anything about the user's preferences explicitly and obtrusively and just with a few hours of interaction, and using the years of experience and vast knowledge of the user, the system already has a good idea of his preferences and is capable of recommending or displaying artworks that the user may find interesting by proximity (in terms of school, style, location, artist, etc...).

Table 2

Art professor profile - Interactions example

Date	Type	Artwork	Author	Collection	School	Timeline	Type	Form	Location
6/7/2013	Reader	Shield	Benvenuto CELLINI	M. Du Louvre	Italian	1550	Other	Metal-work	Paris
6/7/2013	Reader	Study for a seal	Benvenuto CELLINI	M. Du Louvre	Italian	1550	Other	Graphics	Paris
6/7/2013	Reader	Nymph of Fontainebleau	Benvenuto CELLINI	M. Du Louvre	Italian	1550	Mythologi-cal	Sculpture	Paris
6/7/2013	Reader	Juno	Benvenuto CELLINI	M. Du Louvre	Italian	1550	Mythologi-cal	Graphics	Paris
6/7/2013	Reader	Helmet	Benvenuto CELLINI	M. Du Louvre	Italian	1550	Other	Metal-work	Paris
6/7/2013	Reader	Holy Family	Bernaert van ORLEY	M. Du Louvre	Flemish	1550	Religious	Painting	Paris
6/7/2013	Reader	The Virgin Holding the Sleeping Child, with St John and Two Angels	Bernardino LUINI	M. Du Louvre	Italian	1550	Religious	Painting	Paris
6/7/2013	Reader	Nude Man Turned to the Right	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Study	Graphics	Paris
6/7/2013	Reader	Piet?	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Religious	Graphics	Paris
6/7/2013	Reader	Nude Man from the Front	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Study	Graphics	Paris
6/7/2013	Reader	Nude Woman on her Knees	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Study	Graphics	Paris
6/7/2013	Reader	Study of a Nude Man	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Study	Graphics	Paris
6/7/2013	Reader	Back View of a Woman	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Study	Graphics	Paris
6/7/2013	Reader	Madonna and Child with St John	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Study	Graphics	Paris
6/7/2013	Reader	Male Figures	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Study	Graphics	Paris
6/7/2013	Reader	Satyr's Head	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Study	Graphics	Paris
6/7/2013	Reader	Male Figure	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Study	Graphics	Paris
6/7/2013	Reader	Christ Crucified between the Virgin and Nicodemus	Michelangelo BUONARROTI	M. Du Louvre	Italian	1550	Religious	Graphics	Paris
6/7/2013	Reader	Adoration of the Shepherds with a Doonor	Palma VECCHIO	M. Du Louvre	Italian	1550	Religious	Painting	Paris
6/7/2013	Reader	Micheletto da Cotignola Engages in Battle	Paolo UCCELLO	M. Du Louvre	Italian	1450	Historical	Painting	Paris
6/7/2013	Reader	Portrait of Thomas Stachel	Paris BORDONE	M. Du Louvre	Italian	1550	Portrait	Painting	Paris

Besides the latter, using time-stamps on the interactions the system can also prioritise and rearrange the artworks in case the user's interests start to change due to his work or personal taste.

In table 2 a subset of the interactions expected of such a profile can be seen. This interactions have been created with the purpose of illustrating the way in which a more experienced and focussed user can interact with the system. As this user is meant to be an art professor, the system will most likely assess him as an expert user, due to his probable age, set of interactions, amount of interactions and places in which he has interacted with the system. Based on the subset of interactions visible in table 2, it is possible to conclude that this user has a particular interest on Italian renaissance, and more specifically on Michelangelo Buonarroti. The subset of interactions on table 2 belongs to a visit made the 6th of July 2013 to the Louvre museum in Paris.

Table 3
Art professor profile - Interactions example from interactive table

Date	Type	Artwork	Author	Collection	School	Timeline	Type	Form	Location
6/7/2013	Table	Virgin and Child with St John and Angels	Michelangelo BUONARROTI	National Gallery	Italian	1550	Religious	Painting	London
6/7/2013	Table	Entombment	Michelangelo BUONARROTI	National Gallery	Italian	1550	Religious	Painting	London
6/7/2013	Table	Leda and the Swan	Michelangelo BUONARROTI	National Gallery	Italian	1550	Mythologi-cal	Painting	London
6/7/2013	Table	Battle Scene	Michelangelo BUONARROTI	Ashmolean Museum	Italian	1550	Study	Graphics	Oxford
6/7/2013	Table	Study for a Deposition	Michelangelo BUONARROTI	Ashmolean Museum	Italian	1550	Study	Graphics	Oxford
6/7/2013	Table	St Anne with the Virgin and the Christ Child	Michelangelo BUONARROTI	Ashmolean Museum	Italian	1550	Religious	Graphics	Oxford
6/7/2013	Table	Staircase	Michelangelo BUONARROTI	Biblioteca Medicea Laurenziana	Italian	1550	Interior	Architec-ture	Florence
6/7/2013	Table	Allegorical figure	Michelangelo BUONARROTI	Casa Buonarroti	Italian	1550	Study	Graphics	Florence
6/7/2013	Table	The eighth bay of the ceiling	Michelangelo BUONARROTI	Cappella Sistina	Italian	1550	Religious	Painting	Vatican City
6/7/2013	Table	Resurrection of Christ	Michelangelo BUONARROTI	Casa Buonarroti	Italian	1550	Religious	Graphics	Florence
6/7/2013	Table	The ninth bay of the ceiling	Michelangelo BUONARROTI	Cappella Sistina	Italian	1550	Religious	Painting	Vatican City

The interactions also show that after tagging some artworks in the system, the user approached an interactive table and started searching for related content. From the recommended items the ArtVis system showed to the user, he selected some of the artworks, creating therefore the second set of interactions. Theoretically this user was looking for artworks by Michelangelo Buonarroti, so approaching the interactive tables and using the explore panel and the recommendations made by the system according to his profile and previously shown preferences table 3 would be a perfectly feasible set of complementary interactions that in the future will enhance this user's recommendations.

5.1.3 The Art Enthusiast

What about an art enthusiast that regularly attends exhibitions in his favourite local museum and that already has a membership card? Let us assume that due to the static nature of his profile, some basic information about him is already known, especially regarding his location and interests since those are more likely to be rather static. Browsing through the ArtVis system this user may find new information about the artworks that he already knows and connect them with similar others through the information that has been gathered. For instance, if in the permanent exhibition of his local museum he has shown interest for artworks belonging to a particular school, he can easily track down artworks belonging to the same institutions and interact with them.

Following the same workflow, Table 4 shows a subset of feasible interactions for the art enthusiast profile. From the interactions it is possible to infer that this enthusiast resides in the city of Amsterdam or nearby and he regularly visits the Rijksmuseum. Most of the artworks exhibited at the Rijksmuseum belong to the Dutch school which was very influential from the 16th until the 18th century. These features shared by the artworks will be used in further recommendations made to this user's profile and will facilitate the user find Dutch artworks when he visits other museums around the world.

Table 4
Art enthusiast profile - Interactions example

Date	Type	Artwork	Author	Collection	School	Timeline	Type	Form	Location
5/3/2013	Reader	Still-Life	Abraham van BEYEREN	Rijksmuseum	Dutch	1700	Still-life	Painting	Amsterdam
5/3/2013	Reader	Cornelis Tromp in Roman Costume	Abraham WILLAERTS	Rijksmuseum	Dutch	1650	Portrait	Painting	Amsterdam
5/3/2013	Reader	Silver candlestick	Adam LOOFS	Rijksmuseum	Dutch	1700	Other	Metal-work	Amsterdam
5/3/2013	Reader	Boatmen Moored on the Shore of an Italian Lake	Adam PYNACKER	Rijksmuseum	Dutch	1700	Landscape	Painting	Amsterdam
20/4/2013	Reader	Covered ewer	Adam van VIANEN	Rijksmuseum	Dutch	1650	Other	Metal-work	Amsterdam
20/4/2013	Reader	Winter Landscape	Jacob Isaackszon van RUISDAEL	Rijksmuseum	Dutch	1700	Landscape	Painting	Amsterdam
20/4/2013	Reader	Landscape	Jacob van GEEL	Rijksmuseum	Dutch	1650	Landscape	Painting	Amsterdam
20/4/2013	Reader	Portrait of the Artist and His Family	Jacob Willemsz I DELFF	Rijksmuseum	Dutch	1600	Portrait	Painting	Amsterdam
12/7/2013	Reader	Portrait of a Young Boy	Jacob Willemsz I DELFF	Rijksmuseum	Dutch	1600	Portrait	Painting	Amsterdam
12/7/2013	Reader	Landscape with a Tall Tree on the Right	Jacob Woutersz VOSMAER	Rijksmuseum	Dutch	1650	Landscape	Graphics	Amsterdam
12/7/2013	Reader	Italian Landscape with the Ruins of a Roman Bridge and Aqueduct	Jan ASSELYN	Rijksmuseum	Dutch	1650	Landscape	Painting	Amsterdam

5.1.4 The One-Time User

The last scenario will show how one-time users can also profit from the semantic user profiling to improve their user experience with the system. As the focus of this work is to explore how detailed can long term profiles be built without being overly intrusive, login is an optional feature. One-time users can browse freely without giving up any information about themselves, but letting the system know about their preferences while they do it will not only influence queries for all users but it will also help building an anonymous profile for all the casual visitors that can guide them through popular content.

As explained in the algorithm section, all the interactions registered by users, including one-timers will be stored and used to determine the popularity of artworks. The most popular artworks will be recommended to users with the generic profile (unauthenticated) and when filters are applied using the explore panel, the recommendations will be modified to display the most popular artworks meeting the criteria the user selects in the filters.

This feature is known as social tagging, and even though it is usually prone to polysemy, it is easily scalable as a large number of taggers usually help improve the quality of the tags [59]. This scalability and quality enhancement is realised due to what is known as wisdom of the crowds. This phenomena is not strange to the planet art, since even though the perception of art is subjective, the marked trends in art help the artwork categorisation. In the case of the ArtVis system this property is materialised when the user gets recommended the most popular artworks in the filters it sets using the interactive tables.

To set an example Table 5 shows a controlled set of interactions made by two users. At the point these interactions were captured, these two users were the only ones registered in the system. This will point out how the system behaves in presence of the cold start problem (new user/item - not enough interactions).

Table 5
All users - Interactions sample

Username	Artwork	Author	School	Timeline	Type	Form
test profile 2	Storm at Sea	Pieter the Elder BRUEGEL	Netherlandish	1550	Landscape	Painting
test profile 2	Summer	Pieter the Elder BRUEGEL	Netherlandish	1550	Landscape	Graphics
test profile 2	The Conversion of Saul	Pieter the Elder BRUEGEL	Netherlandish	1550	Landscape	Painting
test profile	The Little Street	Johannes VERMEER	Dutch	1700	Landscape	Painting
test profile 2	The Suicide of Saul	Pieter the Elder BRUEGEL	Netherlandish	1550	Landscape	Painting
test profile	View of Delft	Johannes VERMEER	Dutch	1700	Landscape	Painting
test profile	View of Zaragoza	Diego Rodriguez de Silva y VELAZQUEZ	Spanish	1650	Landscape	Painting
test profile	Villa Medici, Grotto-Loggia Facade	Diego Rodriguez de Silva y VELAZQUEZ	Spanish	1650	Landscape	Painting
test profile	Aesop	Diego Rodriguez de Silva y VELAZQUEZ	Spanish	1650	Mythological	Painting
test profile	Allegory of Music or Erato	Filippino LIPPI	Italian	1500	Mythological	Painting
test profile	Arachne	Diego Rodriguez de Silva y VELAZQUEZ	Spanish	1650	Mythological	Painting
test profile	Diana and her Companions	Johannes VERMEER	Dutch	1700	Mythological	Painting
test profile	View of the Strozzi Chapel	Filippino LIPPI	Italian	1500	Interior	Painting
test profile	View of the Strozzi Chapel	Filippino LIPPI	Italian	1500	Interior	Painting
test profile 2	Landscape with the Fall of Icarus	Pieter the Elder BRUEGEL	Netherlandish	1550	Landscape	Painting
test profile	View of the frescoes in the Strozzi Chapel	Filippino LIPPI	Italian	1500	Interior	Painting
test profile	Carafa Chapel	Filippino LIPPI	Italian	1500	Interior	Painting
test profile	View of the Carafa Chapel	Filippino LIPPI	Italian	1500	Interior	Painting
test profile	Allegory of Music or Erato	Filippino LIPPI	Italian	1500	Mythological	Painting
test profile	View of the Vaulting in the Strozzi Chapel	Filippino LIPPI	Italian	1500	Interior	Painting
test profile	Allegory of Music or Erato	Filippino LIPPI	Italian	1500	Mythological	Painting
test profile	The Delphic Sibyl	Filippino LIPPI	Italian	1500	Mythological	Painting
test profile	Mercury and Argus	Pieter the Elder BRUEGEL	Netherlandish	1650	Mythological	Painting
test profile	Sibyl	Pieter the Elder BRUEGEL	Netherlandish	1650	Mythological	Painting

From the random subset of interactions by these two users shown in Table 5, it is a safe guess to assume that if a one-time user approaches an interactive ArtVis table, he will get recommended artworks mostly from the Netherlandish, Spanish and Italian schools, and more specifically artworks made by the authors Diego Rodriguez de Silva y VELAZQUEZ, Filippino LIPPI and Pieter the Elder BRUEGEL.

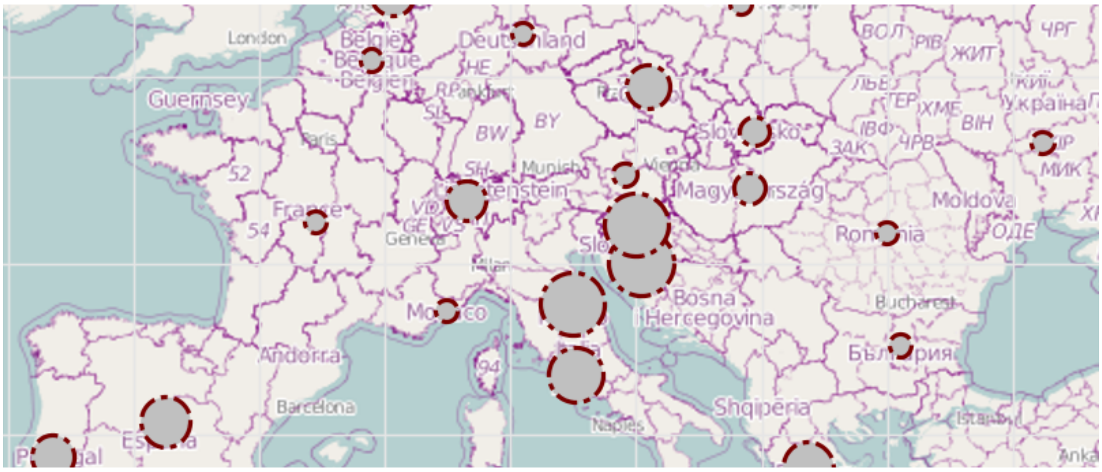

Figure 5.2: Map visualisation of the locations where content is being recommended to a non-authenticated user according to the collective wisdom principle

The map visualisation in Figure 5.2 displays bigger bubbles near the aforementioned areas, which means that several artworks similar to those tagged by the two current registered users in the system are stored in collections in those areas. The initial fish-eye visualisation (Figure 5.3) depicts the most recommended artworks which in this controlled test environment would be artworks from the Spanish school between the years 1650 and 1700. The most popular artist is Diego Rodriguez de Silva y VELAZQUEZ.

Figure 5.3: Fish-eye visualisation of the most popular artworks in the system recommended to unauthenticated users

If a one-time user is more interested in other periods, authors or schools, he can use the explore panel and the filters to visualise the most popular artworks according to his current interest.

Figure 5.4 displays the recommended content after the user filters the recommended content to display only artworks from the years 1550 to 1600.

As shown in Table 5, the most popular author in this time for the users of the system is Pieter the Elder BRUEGEL and most of his artworks are placed in The Netherlands and Germany. The Bubbles on these areas are significantly bigger than in any other part of Europe and the fish-eye visualisation displays the top 50 artworks in the selected area.

The filters in the bottom part of the explore panel can also be helpful to find artworks stored in collections close to the location of the user, that way, when users move to different locations they can easily find interesting artworks in their surroundings. The analyse panel can also help them understand the way in which their favourite art schools evolved over time (Figure 5.5) and where they can find the greatest collections related to those schools.

Figure 5.4: Recommendations made to unauthenticated users based on the current interactions and filters

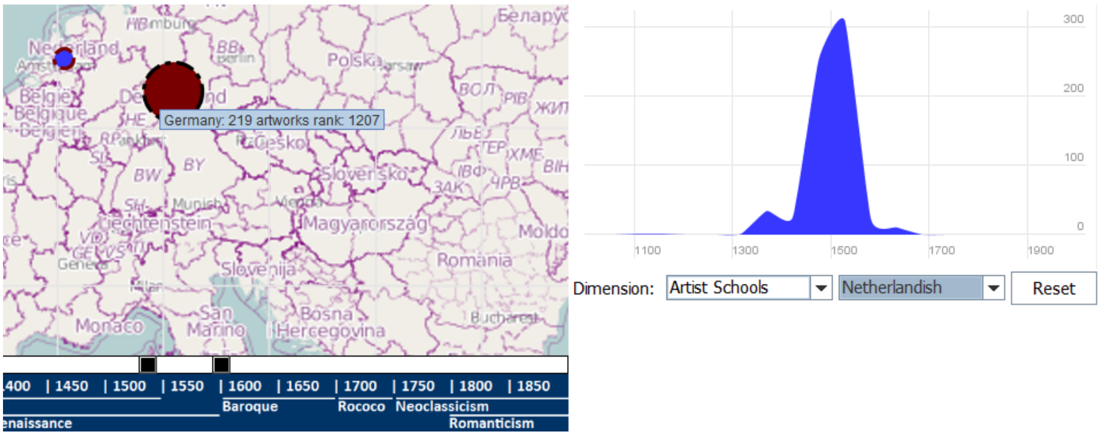

Figure 5.5: Evolution of the Netherlandish school and location of the most important collections related to that school.

Another side effect that highlights the importance of this feature resides in the art looting phenomenon, i.e. the way in which conflicts, wars and invasions between different countries has a direct effect on the artwork locations. In Figure 5.5 can be seen the way in which most of the artworks belonging to the Netherlandish school are being displayed in Germany instead of The Netherlands. Some of the other most known examples are the artworks stolen by the British Empire, Napoleon or The Third Reich. These historical facts explain the reasons why it is not only possible but common to find great collections of Greek artwork in the United Kingdom, Egyptian masterpieces in Paris and Italian renaissance paintings all over Germany. The bright side of it is indeed that users can find art from most of their favourite movements, time-spans and schools in different museums and art galleries all over the world and can analyse them against these historical facts.

Chapter 6

Design Overview

6.1 Workflow Overview

The general idea of the recommendation engine is to receive as an initial input a user from whom nothing or very little is known about, and get to know this user through the interactions this user has with the system. In order to do so, the semantic fingerprints the user leaves in the system will be analysed. Those semantic fingerprints are defined as the set of interactions from which useful information can be extracted.

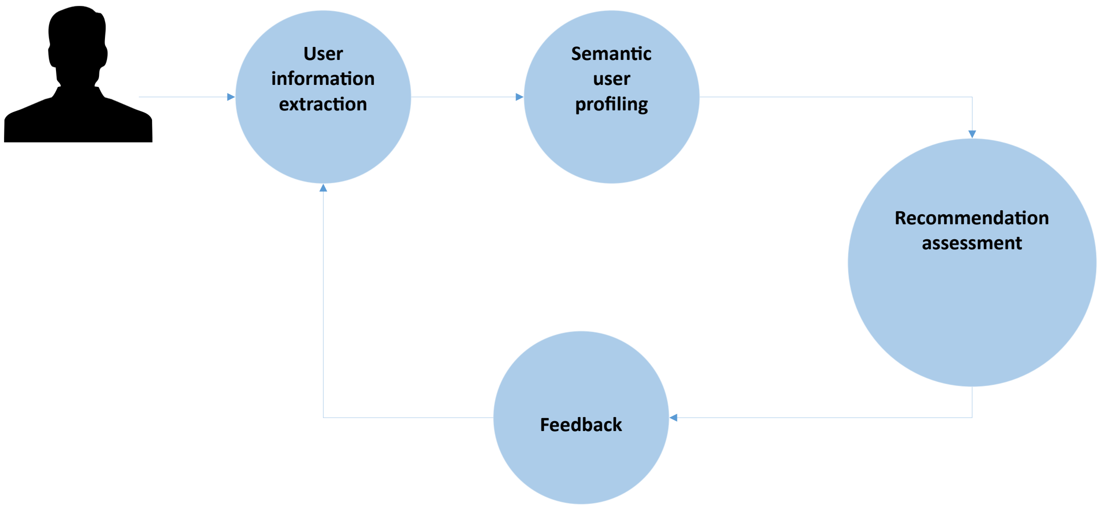

Figure 6.1: ArtVis recommendation engine workflow

In the context of the ArtVis system, and more specifically, the recommendation engine, useful information is such that it allows the system to know the user's preferences in a format that the system can interpret and analyse. Figure 6.1 illustrates the workflow of the recommendation engine. Whenever a user interacts with the system, a user information extraction stage begins, capturing all the information the user leaves about its preferences in terms of artworks. Then, a classification and semantic profiling stage begins to associate users with their peers based on such preferences. That information is used as an input by the recommendation assessment step in which rankings will be assigned to artworks. The artworks with the higher rankings that meet the criteria selected by the user in the filters of the explore panel will be the ones presented to him as recommendations in the browse panel. The map representation will also highlight the geographical locations where such artworks can be found or where their authors lived, according to the user settings. Finally, as the user selects artworks from the browse panel to show in the display panel, new interactions will be created as a feedback that will later enhance the information extracted from the user.

6.2 Information Extraction

As stated in the objectives section, one of the goals of the ArtVis system is to encourage people to develop an interest in art by enhancing the usability of museum applications. Unobtrusive information extraction techniques assist this purpose by helping the user to avoid giving explicit feedback during its visits to museums and art galleries and gathering information about its preferences by monitoring its interactions with the artworks. This is possible by giving them the possibility to express interest in particular artworks with a simple action such as tagging with RFID tags that at a later stage can be added to museum cards or small portable devices like the ones used to guide the users through their visits.

The only time the user would have to explicitly provide information to the system is when it registers to create its profile, process in which assistance of operators can be expected to facilitate and accelerate registration. The system will learn about the users' preferences and calculate the rankings of artworks at login time. The login will also be done with RFID tags, so the user does not have to memorise any usernames or passwords.

Unobtrusiveness is an important part of the usability features provided by the system since in the context of art related applications it is crucial that the users will not get distracted by external elements that can prevent them to enjoy their visits.

As stated in the related work section, important efforts have been made on this area. As a possibility to extend this work, the techniques proposed by Karaman et al. in the MNEMOSYNE system [51] can be implemented in ArtVis in the future. This way the interactions can be reduced to the natural interactions visitors have with artworks in the museum environment instead of making the users tag their favourite artworks.

6.3 Data Annotation

The Museum conceptual schema is built around the most important concepts for the user profiling, such as user, art center, art school, art movement, artist and user transaction. This way the creation of a correct model of the relationships between these entities is guaranteed to accurately represent the user's interests.

The user profile is defined by the user's personal information and its vector of interests constructed throughout its interactions with the system on a long term basis. In order to transform these interactions into the vector of interests it is necessary to annotate and categorise them, so the algorithm can accurately weigh the concept base of the interaction to manipulate the search results and display content more relevant to the user.

A sample of one of these interactions will contain annotated tags such as its source (tag on the artwork or interaction on a terminal) and its type (artwork, artist, art school, museum, etc...). The main goal of the annotations is then to provide a context for each transaction. These annotations will be used by the algorithm to categorise and weigh the interaction.

The artworks in the system contain tags, which are used to provide additional contextual information and group them for instance according to inspiration or art movement. This metadata is processed by the algorithm to provide higher ranks to artworks with similar tags than the ones the user has expressed its interest for.

6.4 Semantic User Profiling

So far, user profiling has been addressed in either following and analysing user transactions or capturing and studying the user's physical behaviour [84, 23]. Given the context of this application and its inherent ubiquitousness, for this work it has been decided to mix those two approaches as most of the user's transactions are done through physical gestures. In ArtVis an important part of the gathered information about the users' interests is gathered on the spot, at the museum or art gallery when the user explicitly tags an artwork and rates it in a very simple way, while his impression of it is still fresh. The other part is done while the user browses through the application on an interactive table. Every time the user interacts with the system in either of those ways and interacts with a work of art he is informing the system that he is somehow interested in that work of art. All this information is gathered and merged as semantic fingerprints and after every interaction it is used to enhance the queries the user performs while searching for relevant content.

This query enhancement is achieved by asynchronously updating a vector of interests with a constant length after every interaction using both the information of the query the user performed (when present) and the available information about the work of art with which the user interacted [46, 82].

Every time a user logs in the ArtVis System its preferences will be loaded. To keep up with the interactions made by users that constantly visit museums and art galleries, the relative expertise of the user will be calculated on the spot. The parameters taken into account to calculate the expertise of the user include its age, the popularity of the artworks it has interacted with, the amount of preferences stated and the amount of places in which it has interacted with the system. The categorisation contains two levels, expert and amateur, and according to the category in which the system locates the user, the recommender algorithm will set up a path of action for its recommendations. There are three predefined exploration profiles for the recommendations, browse, discover and default defined as follows:

- Browse: The browse profile is set up for users the system has assessed as experts and it will rank higher the artworks it has already interacted with and the ones closely related. This profile is meant for users that already have had a considerable contact with art and want to have easy access to their favourite works of art.

- Discover: This profile is meant for users the system has assessed as amateurs and are more interested in discovering new artworks with similar features than the ones they have already liked before. For instance artworks from the same authors, authors from the same schools and time spans will rank higher so the users can find them more easily and develop a certain interest in them.

- Default: The default profile is set up for unauthenticated users and it is configured so it takes into account the most popular artworks in the system to display and give a higher ranking. The popularity of the artworks is assessed using the amount of interactions users have had with them in the past, so naturally the most popular artworks will rank higher, the algorithm will also consider the ones from the same authors, authors from the same schools and time spans to help the one-time and occasional unregistered users discover the masterpieces of the planet art.

6.5 The Algorithm

With the purpose of achieving the most accurate recommendations with the information available, a hybrid recommendation algorithm has been developed, taking elements from content-based and collaborative filtering techniques depending on the profile of the user and what it is known about him. The default profile, made for anonymous and one-time users follows a pure content-based approach, in which the interactions created by all users in the system are combined to rank artworks in such a way that the most popular ones over the past six months that are covered by the filters the user selects in the explore panel will be the ones that will rank higher. As previously mentioned in this document, the system will consider as the most popular artworks the ones that have the most interactions by users.

The range of six months has been set up to assist users in finding trending artworks present in temporal exhibits. Due to the static nature of the museum environment and the scope that limits the artworks to a predefined, past time span (11^{th} to 19^{th} Century), it is assumed that no major changes will be detected in the popularity of artworks belonging to permanent exhibitions. This way the six month filter will only take effect when external events change the trends on artwork viewings, most likely due to temporary displays in determined places. The default profile will also help circumvent the new user problem since it will also be set up for users that have not had any interactions with the system yet. No additional techniques have been applied to circumvent the new item problem since in the context of museums and art galleries where the interactions are created, artworks are displayed regardless of their popularity index on the ArtVis system.

The first part of the algorithm assesses the relative expertise of the current user. In order to do so, the system takes into account parameters such as its age, amount of artworks rated, amount of places in which interactions have been made and the popularity the artworks present in its interactions. When the user has no interactions or an unauthenticated user is interacting with the system, the default profile will be used to provide recommendations. The expertise assessment only contains two levels, expert and amateur, one leading to a browse profile and the other to a discover profile.

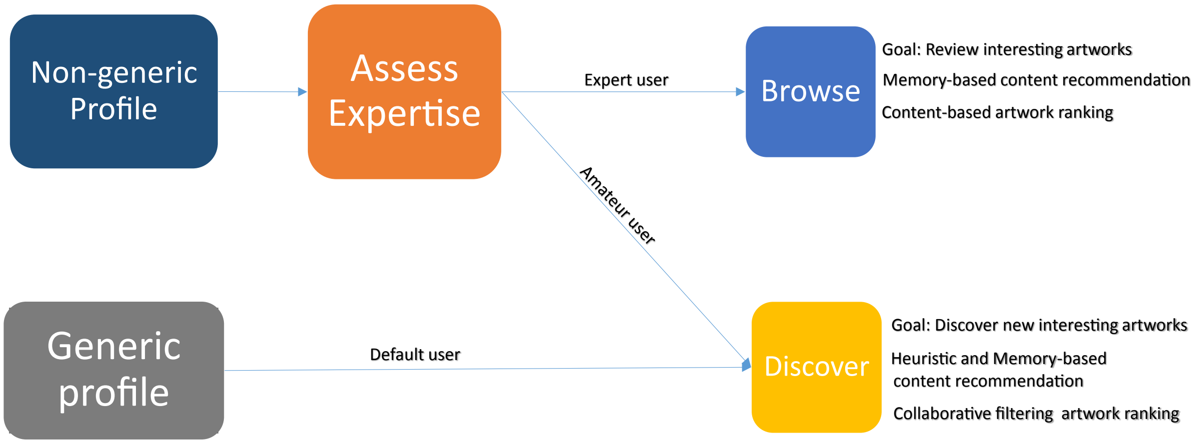

Figure 6.2: Recommender algorithm workflow overview.

6.5.1 Generic and Non-generic Profiles

A generic profile is a temporary profile created for one-time and non-authenticated users. This profile is set-up so these users can also benefit from the recommendation engine as amateur users looking for interesting content. The interesting content recommended for users with the generic profile is based on the collective knowledge of preferences by all the users registered in the system. Content is recommended to these users according to the most popular artworks in the last six months using a heuristic-based, collaborative filtering recommendation technique.

On the other hand, as users approach interactive tables and authenticate themselves with RFID tags, the system looks for that tag in the database. If the tag is found, the user is registered. Registered users have what has been defined as a non-generic profile, which means the system already contains some information about them and their preferences. When a non-generic profile is detected the algorithm proceeds to assess the expertise of the user based on the information extracted from its interests.

6.5.2 Expertise Assessment

Assessing the relative expertise of a user in the ArtVis system is probably the most fuzzy part of the algorithm, especially since art appreciation is fully subjective. In contemplation of recommending the most accurate content to each user two levels of expertise have been created: one for fully amateur users which, according to their taste and interactions, the system will assume want to discover new artworks and a second one for users that have more experience in the museum setup. For the second case the system will assume these users want to have a closer look into artworks they have already interacted with and other ones closely related.

To determine which users belong to which group, the algorithm implements a ranking system, awarding points according to the user's age, the amount of artworks tagged, the average popularity of artworks tagged and the amount of places in which the user has interacted with the system.

To become an expert, the user has to rank at least one hundred points which will be awarded as follows:

- Age: The user will be granted 10 points for every 15 years of age, since users increase their knowledge of art with age.

- Amount of tagged artworks: If users have rated more than 25 artworks, 5 points will be granted for every 25 interactions
- Average popularity of artworks: The popularity of artworks has been scaled from one to five according to the amount of interactions registered for the artwork with the most interactions, so, the popularity of an artwork is calculated as:

$$P_{(a)} = \frac{\#i_{(a)} * 5}{\#i_{(m)}} \tag{6.1}$$

  Where $P_{(a)}$ is the popularity of artwork a, $\#i_{(a)}$ is the amount of interactions the artwork a has registered in the system by different users and $\#i(m)$ is the amount of interactions of the most popular artwork in the system.

  After calculating the popularity index of the artworks the user has interacted with, the average is computed. Points will be granted to the users according to this average. With the purpose of assessing the most probable average for expert users, the higher and lower ends will have the least granted points. This way users that mostly like very popular or very unpopular artworks will get less points than the users who generally like artworks with a medium popularity. Users whose artworks' average popularity is 3 will get 50 points, then users with average 2 will get 40 points, users with average 4 will get 30 points, 20 points for average 1 and 10 for the rest. This was decided since the average inexperienced user will most likely attend exhibitions and tag the "main attractions" such as the Gioconda by Leonardo da Vinci or David by Michaelangelo and other users of the same level will most likely tag relatively unknown artworks of their local museum without really grasping the concept of the exhibition. The users that are considered experts by the system will most likely have a purpose in keeping track of their favourite artworks as part of their work or research, such as art students and professors. Therefore, there is a higher chance that these users are interested not only in the masterpieces that are usually the center of attraction in an exhibition but in the related artwork, their context and a set of influential artworks belonging to a particular art movement or style.
- Amount of places the user has interacted with the system: Since ArtVis and the Web Gallery of Art were conceived as learning platforms, it is also assumed that the more a user interacts with the system, the more it learns about art. Accordingly, in the expertise assessment users that have interacted with the system in most places will rank higher. On account of the latter if a user has tagged artworks in three different places it will get ten additional points for its ranking, if it has done so in five different places then it will get twenty points and thirty for eight or more.

This categorisation method is part of a heuristic-based approach to semantic user profiling. Nevertheless it is only the first step into the creation of the real model that will determine the recommendations set up for each user. The categorisation is mostly based on the goals assumed for the user based only on its extracted preference and personal information. In the next section, the steps that follow will be discussed in detail to gain a better understanding of the reasons that led to the selected categorisation. As shown in figure 6.2, the expertise assessment will dictate the recommendation techniques for each user belonging to a non-generic profile.

6.5.3 Browse Profile

The browse profile is based on a memory-based content recommendation technique. This means that based on a set of heuristics artworks get ranked. Such heuristics are based on the properties shared by the artworks the user has interacted with in the past. As it is designed to help expert users to keep track of the artworks they have previously shown interest for and other ones closely related, it uses a content-based approach to go through the user's interactions and gives a higher ranking to the artworks tagged by the user that share most things in common and the ones that have not been tagged yet by this user but that are strongly related.

As depicted in figure 6.2, the goal of this profile is to review interesting artworks, therefore it is assumed that the artworks that are strongly correlated are the ones in which the user is most interested. The previous assumption is made since tagging strongly correlated artworks implies that the user has attended exhibitions of particular art movements or authors or has chosen to explicitly find them in bigger exhibitions.

The most simple scenario to illustrate this case would be an art enthusiast who is studying a particular movement and attends one or multiple exhibitions to find artworks that belong to it. After this user tags a certain amount of artworks belonging to this movement and most particularly artworks by a particular subset of authors in the aforementioned movement, the system will start recommending similar content. This and the other scenarios were explored in detail in the corresponding section.

As mentioned, the browse profile uses a content-based ranking technique to recommend artworks. Similar to most recommendation engines that use rankings, the artworks will receive ranking points in concordance to the fulfilling of certain characteristics. These characteristics point to how similar they are to the ones previously tagged by the user in terms of context and the tangible information available about them. This means that two artworks that humans may consider totally different but that share sufficient contextual information are going to be considered similar by the system. This is to be expected due to the semantic gap between the contextual information of artworks and the visual properties they have. This is not considered a major issue since artists usually deliver content in their artworks according to their context, so the artworks made by an author or a group of artworks from a specific school at a determined point in time are most likely similar. In case some of those artworks are utterly different, that twist can also be a point of interest valuable to an experienced art observer.

Contrary to the user expertise assessment, the rankings set up for artworks to be recommended do not have a lower limit or a minimum amount of points to get so they can be recommended to a user, they also do not have an upper limit of points they can get. This feature gives the algorithm more flexibility in the recommendations, allowing to recommend content even when the system is not able to extract its preferences. In such cases popular content that resembles the artworks tagged by the user will be considered for recommendation.

To accomplish an accurate recommendation the browse profile takes into account the following parameters:

- Timespan: The algorithm grants 500 points to artworks belonging to the same timespan as the one the user has shown most interest in. This is done by grouping the artworks the user has expressed interest in by timespan and selecting the two most popular. Artworks belonging to the second most popular timespan in the user's history of preferences will be granted 400 points.

- Author: The system will select the most popular authors included in the user's preferences and grant from 600 to 50 points to their artworks according to the scaled amount of artworks tagged.
- Art school: Similar to the authors, the system will grant rank points for the artworks belonging to authors in the same schools as the ones the user has shown previous interest. As the art schools are too generic the awarded points will be less, that is in a scale from 300 to 25 points.
- Artwork type: The same technique is applied to grant points for artwork type. Since this is in general a great form of extracting the inspiration that authors had at the moment of the creation of the artwork, the points awarded for artworks with the same type as the ones the user has shown interest in will be the same as the ones awarded by author. Artwork types reflect the kind of artwork in a way that they give hints to make a better categorisation. Some examples of artwork types are: religious, portrait, landscape, etc.
- Artwork form: Artwork forms are less descriptive as they reflect only the tangible nature of the artwork, therefore the scale used to rank artworks by this parameter will be from 100 to 10 points.
- Collection: Including the artwork collection may help users find interesting artworks in their local museum and also enhances the correlation between artworks that are presented in the same location. The algorithm assigns a range from 300 to 25 points to artworks in the same collections as the ones tagged by the user according to the same scale as art schools.

After granting the rank points to the artworks, the system will return the values of the rankings to the interface where the visualisation tool will scale the size of the bubbles in the map representation according to the average rank of the artworks left in the filters selected by the user in the explore panel. The fish eye visualisation in the browse panel will show the top ranked artworks. If the user selects one of these artworks a new interaction will be created, the system will know the user expressed interest in that artwork. This action is considered as a feedback action to the recommender algorithm which will take into account these selected artworks to improve the recommendations.

6.5.4 Discover Profile

In contrast to the browse profile, the discover profile is designed to help inexperienced users to discover new and interesting artworks similar to the ones they have shown interest for. The discover profile also helps the user keeping track of their tagged artworks but will rank higher new artworks keeping a balance in the recommended content for comparison.

The main difference then, is the usage of a collaborative filtering technique, this means that the algorithm at this point will use the information about the preferences of other profiles to recommend content instead of using only what is stated in the user's preferences. According to the scenario the algorithm will follow one out of two paths: the first case, for the default user that can be either a one-time or an unauthenticated user, the system will take into account all the interactions created in the last six months by all the registered users in the system. As explained in the algorithm overview, the six month rule applies to attract interest of users in art trends and the recent or current temporary exhibitions. As the time range of the artworks in the system is fixed it is assumed that the permanent exhibitions will not change the preferences of users in a significant way. The second path is for the authenticated amateur users. In this case, the algorithm uses a k-nearest neighbours technique to select a fixed number of profiles similar in preferences to the current user to influence the recommended content.

The k-nn classifier uses a simple probabilistic technique that assigns a weigh to the input parameters and then extracts the profiles with the most similar interests according to those features, without taking into account the artworks themselves to avoid extreme repetition. For performances enhancement, the algorithm preselects the users that have tagged artworks with similar features and calculates the similarity level on them.

The probability of a user classifying to the same category of the one in evaluation, and therefore be eligible for recommendation is calculated as follows:

$$P(u) = \sum_{i=1}^{n} P(f_i) \tag{6.2}$$

Where n is the amount of parameters or features taken into account, and $P(f_i)$ is the weighted probability of feature i according to the interactions the user has had with the system compared with the ones of the user in current interaction. $P(f_i)$ is computed as:

$$P(f_i) = P(f_{i|u}) \times W(f_i) \tag{6.3}$$

$P(f_{i|u})$ is the probability of the artworks tagged by the current user being similar to the ones of the user interacting with the system according only to feature i. Variable u represents the artworks tagged by the user in interaction and $W(f_i)$ the weight assigned to feature i, since not all features have the same relevance weight. The relevance weight of the features is given by the relative expressiveness of the features in question given the context of the system.

The features taken into account are the same used in the browse profile and are weighted according to the following values:

- Timespan: 0.1
- Author: 0.4
- Art school: 0.1
- Artwork type: 0.3
- Artwork form: 0.1

These weights were assigned in accordance to the expressiveness of the features, given by the amount of different possible values for each feature in the database. The calculations of the rankings are done following the same principles and values as the browse profile. The difference resides in the fact that instead of taking into account the artworks tagged by the user in question, the algorithm gathers the preferences of the 20 most similar users according to the aforementioned semantic profiling technique to enhance the recommendations in discovery mode. This way, the chances of a user to find more interesting artworks increase by recommending artworks that people with similar interests have also liked in the past.

This combination makes of this algorithm a hybrid recommendation algorithm, given the fact that it contains collaborative and content-based filtering techniques to provide recommended content. Such techniques were described in section 2.2.3.

Chapter 7

Results

The main purpose of this work was to provide art enthusiasts and people from the museum environment with an interactive tool that facilitates browsing and discovering artworks in a large dataset. The results will be presented according to the specific objectives for each research area involved in this study.

7.1 Information Extraction

The goal of this part of the work was to find a suitable unobtrusive way to allow users to state their preferences in order to recommend content related to such preferences. As it was stated in the related work section, users are reluctant to provide feedback by using questionnaires or interviews so it was imperative to find a method to extract their preferences without using questions about them. Both of the advanced users that participated in the user study agreed that external portable devices would distract users and alter their perception of the displayed. It is important that users are attentive to the message the artist or the people who organise the exhibits are trying to send. A mobile phone or other device would in most cases prevent users to grasp such message.

The RFID-tags are the middle ground between traditional information extraction techniques and gesture recognition/analysis. They not only allow users to express their preferences in an unobtrusive way but also provide the advantage of reliability as the user is explicitly expressing its preferences to the system. In order to find a way to obtain feedback on the recommendation engine that would eventually serve as another input for preferences, a user transaction monitoring technique was implemented. This user transaction monitoring is as unobtrusive as the tagging itself as it is based on the selections the user makes on the browse panel to view their details. Selecting such artworks states an interests that is registered to create new interactions that will help building its profile.

7.2 Semantic User Profiling

The offline experiment was useful to determine the accuracy of the user profiles at a small scale and especially the user categorisations made for the discover profile. To evaluate the user profiling a set of twenty artificial users was created and categorised within four groups. An additional set of five users was created that did not belong to any category. Those groups were based on arbitrary models to which a set of preferences was assigned. Interactions were created randomly selecting artworks corresponding to those preferences to ensure diversity and objectivity. The tags were used to authenticate the users in the system and a log table in the database would store the profiles that the system considered similar to the one in interaction while the recommendations were built. The experiment was repeated 3 times changing the number of interactions for each users from 10, then 20 and finally 30 interactions per user.

The purpose of this artificial experiment is to ensure that the distribution of interests is easily traced and that the system is able to categorise users and extract their preferences based only on their interactions with the system. As some of those interests overlapped with each other, the first experiment registered three anomalies, that is, three users that were placed in the wrong categories. The additional users were not placed in any category, meaning, no similar profiles were found for them, so their recommendations were based purely on their own interactions.

Increasing the number of interactions registered a positive effect as with twenty and thirty interactions per user no anomalies were detected during the experiment. This shows a tendency to increase the accuracy of the recommendations when the number of interactions is also larger.

Table 1
Semantic user profiling - Classification example

Profile	Points for artists proximity *0.4	Points for art schools proximity *0.1	Points for artwork type proximity*0.3	Points for artwork form proximity*0.1	Total Points	Position
test user 1	200	200	200	200	180	1
test user 4	158.28	158.28	200	100	149	2
test user 9	117.69	187.69	200	100	135.846	3
test user 8	117.69	187.69	200	100	135.846	4
test user 6	117.69	192.7	200	95	135.7	5
test user 5	115.3	187.5	190	105	134.56	6
test user 7	112.69	187.69	200	100	133.846	7
test user 24	0	112.23	200	100	81.2308	8
test user 20	0	112.3	200	100	81.2308	9
test user 21	0	95	200	117.23	81.225	10
test user 23	0	95	200	117.23	81.2308	11
test user 22	0	95	200	117.23	81.2308	12
test user 2	0	60.113	200	100	76.113	13
test user 28	0	22.3	190	104	72.2308	14
test user 29	0	22.3	190	104	72.2308	15
test user 13	0	30	120	145	53.5	16
test user 14	0	30	120	145	53.5	17
test user 15	0	30	120	145	53.5	18
test user 12	0	30	120	145	53.5	19
test user 19	0	30	120	145	53.5	20

Table 1 Displays the profiles selected as the most similar for one of the test users in the offline experiment. To simplify the tracking of the users the profiles have been identified with consecutive numbers according to the order in which they were created. The user under consideration is number 1. As it is to be expected the algorithm chooses the same number as the most similar one to itself. This constitutes the content-based approach of the recommender algorithm as a user's preferences will be taken into account along with the ones of other similar profiles. According to the way in which the experiment was designed the users 4 to 9 belong to the same group as number 1. Users 10 to 14 have preferences closer to the French school, users 15 to 19 to the Italian school, users 20 to 24 to the German school and users 25 to 30 have tagged purely random artworks in the system.

The points column in Table 1 represent the actual ranking each user got in the semantic user profiling step and these 20 users got the higher ranking. As it was to be expected besides user

1, users 4 to 9 were the higher ranked of all. The floor of the ranking algorithm is 100 which means only the profiles that were ranked higher than that will be used to recommend content to the current user. Profiles 1 and 4 to 9 were effectively the only ones with a higher ranking than 100 points (marked in green in Table 1), constituting then the ”neighbourhood” of user 1 according to the nearest neighbours approach. The other profiles will be discarded from the classification.

7.3 Recommendation Engine

According to the offline experiment, the recommender engine was able to find related content for all the scenarios described. One-time users are presented with the most popular content that meets the filters they give to the system to explore. The recommendations to expert users tend to favour artworks that they have already interacted with, serving them as highlights for interesting items and helping them to rapidly browse through their highlighted artworks. As for amateur users, content is recommended according to their preferences which are extracted from their transactions and the ones made by users the system has considered its ”peers” in the user categorisation step.

User 1 in the offline experiment has 65 interactions with the system. From those 65 interactions, 57 correspond to artworks by authors from the Netherlandish school, especially the author Bruegel. The artwork type distribution for the artworks tagged by user 1 are as follows:

- Genre: 11.05%
- Landscape: 104%
- Religious: 8.45%
- Still-life: 5.2%
- Other: 2.6%
- Historical: 1.3%
- Mythological: 1.3%
- Portrait: 1.3%
- Study: 0.65%

After running the recommender algorithm without any filters for user 1 and analysing the first 100 results without any filters the results were the following:

- 52% of the artworks recommended are artworks by Bruegel
- 82% of the artworks recommended belong to the Netherlandish school
- 38% Have religious artwork type
- 26% Have genre artwork type

- 13% Have landscape artwork type
- 7% Have portrait artwork type
- 4% Have Still-life artwork type
- 3% Have mythological artwork type
- 2% Have historical artwork type
- 1% Have Study artwork type
- 6% Have other artwork type

The mayor influence of the authors and schools is notorious while the artwork types, artwork forms and timelines features are considerably less expressive. The similarities between the recommended artworks and the artworks tagged by the user suggest that users have great chances of discovering interesting artworks through the recommendation engine.

7.4 Expert User Assessment

User-based evaluations are generally considered the most reliable and valid way to estimate usability in applications [31]. Therefore, in order to assess satisfaction and given the context of the application, it was decided to perform a user-based study. A complete evaluation of the original ArtVis system was carried out in [33]. To complement such evaluation an expert user assessment was done to conduct the implementation of the recommendations. During the user studies, the system received very positive feedback. The art expert users found fascinating the ease with which they could follow historical changes in the art perception and production while using the analyse panel and the filters. This shows how the used visualisation techniques help the analysis of large datasets.

The complementary expert assessment was conducted as an interview with two art experts. The art experts both hold a Bachelor of Arts degree in Art history and are 25 and 30 years old respectively. They both apply concepts of art history in their current studies which makes them suitable for such assessment.

The art experts were asked to freely play with the system for ten minutes to see how fast they could understand the system workflow and the interface without any further training. Then they were asked to perform some basic tasks while explaining the panels on the interface and lastly, they were asked to provide feedback on the functional requirements desired for users at their level and the graphical user interface.

The Artvis system is considered a user-centered system, therefore usability is one of its main concerns. An iterative approach was used in which adjustments were made or planned according to the feedback provided by the users that performed the evaluation. Next some of the results from the user studies are displayed.

7.4.1 Intuitive Interfaces

The art expert users found the graphical user interface very intuitive and easy to use. They could understand the workflow and the main features offered by the application with just a couple of minutes of interaction without any further training or explanations. The users gave themselves tasks to accomplish and analysed their favourite artists, art movements, schools and artworks. Nevertheless, they struggled finding a way to integrate the filters in the explore panel with the filters in the analyse panel. Some rearrangements are possible in the components to cover that gap.

Another point that users found somehow confusing was the labels in the expert panel for the components that indicate whether the artworks should be placed in the map according to the collection location or the user location. The original labels were ”Author living area” and ”Artwork museum location”. First it was noticed that ”Author living area” was actually referring to the author’s place of birth, making it all even more confusing to an expert user who knows this kind of information. One of the experts made the remark that a great part of the artists in the dataset moved constantly from city to city and country to country and that it would be very interesting to be able to trace those movements according to the place in which the artworks were created. For now that information is not available in the dataset. As for the labels, it was decided to change the ”Author living area” to ”Author birthplace” and ”Artwork museum location” to ”Collection location”.

Lastly, it was also stated that in the analyse panel when the user hovered over a particular item, such item should stand out before selecting it in the features. One of the suggestions given was to add a transparency factor to the other items. This suggestion will also be considered for a future version of the system exploring the capabilities of the visualisation tool.

7.4.2 Tangible User Interfaces

The use of Phidgets was perceived as a positive thing as the expert users considered that such tangible interfaces trigger the curiosity of people and encourage them to approach the interactive tables to explore the system. The RFID tags used to authenticate a user, compile its preferences and tag artworks were also considered an innovative solution to the obtrusiveness problem against the trustworthiness of the extracted preferences.

7.4.3 Look and Feel

The users were also very pleased with the general look and feel of the application. They found the interface very clean and transparent which made it easier to find items and explore their properties. The use of colours was generally considered appropriate and sober. Nonetheless it was stated that traditionally darker colours were used as background for art related applications and exhibitions. The expert user explained that as the white is the combination of all colours it can be exhausting to look for prolonged periods of time, while darker colours would help prevent this effect and make the artworks and their visual properties stand out more easily.

7.4.4 Performance

The overall performance of the application is good as the interface adapts and displays the results of the filters rapidly. The map visualisation quickly reacts to the scrolling and the bubbles are updated in real time according to the time spans selected by the users. On the other hand, during the evaluation it was noted that the recommendation algorithm usually takes between 8 and 16 seconds to load the users' preferences. The main reason for such a delay is the inefficiency of using cursors in MySQL stored procedures to rank the artworks. This time is regarded as considerably high and some measures will be taken to improve the efficiency in the recommendation engine.

In order to improve the performance of the recommendation engine, the most logical bottlenecks are going to be monitored which will determine the way to proceed. In the future work section it is described how to perform some performance tuning over the classifications, expertise assessment and the overall recommendation engine.

Chapter 8

Discussion and Future Work

As artworks are constantly relocated, the data in the Web Gallery of Art is also updated. Nevertheless, the ArtVis system uses a local database that was fed once with the information it contained in 2011. An automated synchronisation with the original source would be desirable, but, due to limitations regarding the lack of normalisation of the data this is currently impossible. If the data in the WGA gets normalised and services are exposed to access the updated information at any time any external system would be able to access real-time information providing the users with better ways to discover and explore artworks. Having information of artwork relocations for temporal exhibitions will indeed help users find interesting content in their surroundings at the time when it is most relevant. This feature could serve as invitation to visit such temporal exhibitions in the user's local museums or their future destinations.

Thanks to the expert user assessments some inconsistencies were detected in the information provided by the Web Gallery of Art project. Some of those inconsistencies were corrected through automated techniques, nonetheless the information is still not totally reliable. Although it is said that at certain points in history it is almost impossible to trace the precedence of artworks, the goal is to get a dataset as complete and precise as possible.

Besides the latter, to complete the dataset it would be an ideal step to obtain also the place of origin of artworks and an approximate date of completion. The expert users expressed a special interest in such a feature as it would be highly desirable to be able to trace down the artists in the space and temporal dimensions. For art history students and professionals obtaining a visualisation of the locations in which specific artists worked during certain periods of time would help them understand better the evolution of art schools and movements in terms of highly influential historical events. The ability to compare the geographical locations of several artists in parallel during a determined timespan would also be a desired feature to facilitate the analysis of the influence that can exist over the evolution of one or multiple artists.

The idea of providing the users with RFID tags to tag their favourite artworks was conceived as a usability advantage and a middle ground in which users did not have to provide feedback by answering questions or polls, but with the advantage of accuracy enhancement through explicit interactions. The possibility of using external devices such as mobile phones was also studied but discarded since such devices often distract users from their tasks and change the appreciation of the artworks displayed. Besides the latter, most museums already provide their own devices to complement the user experience in art exhibitions, coordination and integration with such devices would be an extremely challenging task due to the lack of standardisation. Nonetheless, the positive impact of the monitoring features of the MNEMOSYNE system were acknowledged, and at a future stage their monitoring techniques could be combined with the RFID-tagging mechanism used by the ArtVis system to have a more complete information extraction scheme. The ability of analysing the users' reactions and gestures while appreciating art and also provide them with an unobtrusive but explicit way to state their preferences would be the ideal scenario to explore.

Some further adjustments can still be done to improve the graphical user interface, such as full linking and brushing. Currently, linking and brushing is partially supported, so items selected on the analyse panel are reflected in the explore panel but not the other way around. It also would be ideal to provide the possibility to select multiple items in the explore panel to be further displayed in the analyse panel for efficient comparison and general overview.

Lastly, during the offline experiment some performance issues were detected. To improve the system performance, indexes were created on the database and the algorithm was implemented using stored procedures instead of services in the application server. Some improvements were registered, nonetheless, loading a user's preferences still takes between eight and 15 seconds. Alternatives to cache the user profiling are being explored so the expertise assessment only needs to be updated when needed. Also the use of sets and insert from select statements could be explored to replace some of the cursors used to perform the semantic user profiling step. For design reasons the queries performed by the algorithm in the database are done through the physical tables. Some tuning opportunities arise as views were created exclusively for the usage of the visualisation tool Prefuse. Such views could also be used by the algorithm since they contain the necessary information on the artworks avoiding unnecessary table joins at recommendation time.

8.1 About the Research Questions

a. *Q1:* Which methods should be used to extract information about users preferences in an unobtrusive way?

   A: In order to obtain information about users' preferences it is best to guarantee a certain level of accuracy by providing a way to explicitly state preferences. But, these method should not distract the user from the visit nor take unnecessary amounts of time to be carried out. Tagging is probably the middle ground between usability and accuracy in terms of stating preferences in artworks as it consists of an explicit, simple physical gesture that allows the system to gather such information.

b. *Q2:* How can the system analyse the information about a user in order to provide personalised search results that reflect its interests?

   A: The system is able to provide personalised recommendations through the application of probabilistic techniques that allow comparisons of items based on their features. Those features can be weighted in terms of an arbitrary index of expressiveness to get a better insight of what the user might be interested in.

c. *Q3:* Is it possible to categorise and group users by their preferences in art?

   A: Using the previously mentioned features in items (in this specific case artworks) and having the information about the preferences of each user, it is also possible to compute a similarity level between users. With such similarity index certain artificial intelligence techniques can be used to create categories of users.

d. *Q4:* What relevance weight should be assigned to each category of users?

   A: The relevance weight should be assigned according to the expressiveness level of each feature used to categorise items. The expressiveness level is a measure of the amount of ”information” that can be extracted from it and it can be calculated for instance as the relation between the possible different values available for that feature in the dataset.

Chapter 9

Conclusion

An improvement over the ArtVis system has been presented, a highly interactive application that combines advanced visualisation techniques, tangible interfaces and user preferences exploitation inspired by recommendations of artworks in a museum environment. Offline experiments suggested the capabilities of a hybrid recommendation engine based on semantic user profiling techniques. One of the main focusses of this work was the improvement of quality in the content presented to users in a currently existing application which was achieved through unobtrusive information extraction techniques, content personalisation and context awareness.

Addressing the research questions, it is possible to conclude that in order to extract information about user preferences unobtrusively it is recommended to select a method that implies simple explicit interactions. This way the information is precise, accurate and can be used to recommend similar content to other users. Obtaining this information enables the system to apply classification techniques to provide personalised search results for users.

This work has shown how to categorise and group users based only on their stated preferences in art. This was done through semantic user profiling techniques that weighted the similarity of users according to their preferences and information. The weight assigned to each feature extracted from the known information about the user and its preferences was calculated according to its expressive power.

The application illustrates the potential of using advanced visualisation techniques to improve cognition. Such a feature is achieved by providing users with the ability to easily browse large datasets, highlighting interesting information without filtering any information from the scope of the user. The tangible user interface exploits the users' curiosity, attracting them and encouraging them to explore the dataset browsing through different dimensions. ArtVis provides a good example of a user-friendly, intuitive interface that facilitates the detection of relevant content.

It has also been demonstrated how to bring content recommendations on the go in the context of museums and art galleries beyond the content of the museum itself. One of the strongest points of the recommendation algorithm is the ability to easily change the recommendation techniques according to the profile created for a user. Although the recommended content is quite similar for both profiles, the difference resides in the possibility given by the discover profile to prioritise content that is most likely unknown by the user but that fits its taste with great precision and accuracy.

The main contribution of this work is the expansion of recommendation capabilities towards larger datasets in art-related applications. The currently existing systems that work with recommendations in museums and art galleries focus on the content of the exhibition for which they were designed. This difference also marks the entire workflow of the application and the way in which recommendations are made. ArtVis can be implemented in several museums and art galleries around Europe and be useful for users with different profiles and backgrounds. The combination between semantic user profiling and recommendation capabilities for information systems has been explored in the past, nevertheless in ArtVis such techniques are implemented in a way that the algorithm itself chooses the better suited combination of techniques to be applied for each profile.

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. *Knowledge and Data Engineering, IEEE Transactions on*, 17(6):734–749, 2005.

[2] Michelle Anderson, Marcel Ball, Harold Boley, Stephen Greene, Nancy Howse, S McGrath, and Daniel Lemire. Racofi: A Rule-Applying Collaborative Filtering System. 2003.

[3] Plamen P Angelov and Xiaowei Zhou. Evolving Fuzzy-Rule-Based Classifiers from Data Streams. *Fuzzy Systems, IEEE Transactions on*, 16(6):1462–1475, 2008.

[4] Angeliki Antoniou and George Lepouras. Modeling Visitors' Profiles: A Study to Investigate Adaptation Aspects for Museum Learning Technologies. *Journal on Computing and Cultural Heritage (JOCCH)*, 3(2):7, 2010.

[5] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. Dbpedia: A Nucleus for a Web of Open Data. In *The Semantic Web*, pages 722–735. Springer, 2007.

[6] Chris Baber, Huw Bristow, Sean-Lee Cheng, Anna Hedley, Yuri Kuriyama, Marc Lien, James Pollard, and Phil Sorrell. Augmenting Museums and Art Galleries. In *13rd Conference on Human-Computer Interaction*, pages 439–447, Tokyo, Japan, 2001.

[7] Sung Min Bae, Sang Chan Park, and Sung Ho Ha. Fuzzy Web Ad Selector Based on Web Usage Mining. *Intelligent Systems, IEEE*, 18(6):62–69, 2003.

[8] Andrew D Bagdanov, Alberto Del Bimbo, Lea Landucci, and Federico Pernici. MNEMOSYNE: Enhancing the Museum Experience through Interactive Media and Visual Profiling. In *Multimedia for Cultural Heritage*, pages 39–50. Springer, 2012.

[9] Nicholas J Belkin and W Bruce Croft. Information Filtering and Information Retrieval: Two Sides of the Same Coin? *Communications of the ACM*, 35(12):29–38, 1992.

[10] Nicholas J Belkin and W Bruce Croft. Information Filtering and Information Retrieval: Two Sides of the Same Coin? *Communications of the ACM*, 35(12):29–38, 1992.

[11] Daniel Billsus and Michael J Pazzani. Learning Collaborative Information Filters. In *International Conference on Machine Learning*, volume 98, pages 46–54, 1998.

[12] Daniel Billsus and Michael J Pazzani. User Modeling for Adaptive News Access. *User Modeling and User-Adapted Interaction*, 10(2-3):147–180, 2000.

[13] Alberto Del Bimbo, Giuseppe Lisanti, Iacopo Masi, and Federico Pernici. Person Detection Using Temporal and Geometric Context with a Pan Tilt Zoom Camera. In *Pattern Recognition (ICPR), 2010 20th International Conference on*, pages 3886–3889. IEEE, 2010.

[14] Stephen Bitgood. An Overview of the Methodology of Visitor Studies. *Visitor Behavior*, 3(3):4–6, 1988.

[15] Stephen C Bitgood and Ross J Loomis. Environmental Design and Evaluation in Museums. *Environment and Behavior*, 1993.

[16] Paul S Bradley and Usama M Fayyad. Refining Initial Points for K-Means Clustering. In Jude W. Shavlik, editor, *International Conference on Machine Learning*, volume 98, pages 91–99, 1998.

[17] John S Breese, David Heckerman, and Carl Kadie. Empirical Analysis of Predictive Algorithms for Collaborative Filtering. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial i Intelligence*, pages 43–52. Morgan Kaufmann Publishers Inc., 1998.

[18] Erich Bruns, Benjamin Brombach, Thomas Zeidler, and Oliver Bimber. Enabling Mobile Phones to Support Large-Scale Museum Guidance. *MultiMedia, IEEE*, 14(2):16–25, 2007.

[19] Jay Budzik, Kristian J Hammond, and Lawrence Birnbaum. Information Access in Context. *Knowledge-based Systems*, 14(1):37–53, 2001.

[20] Robin Burke. Hybrid Web Recommender Systems. In Peter Brusilovsky, Alfred Kobsa, and Wolfgang Nejdl, editors, *The Adaptive Web*, volume 4321 of *Lecture Notes in Computer Science*, pages 377–408. Springer Berlin Heidelberg, 2007.

[21] Stuart K Card, Jock D Mackinlay, and Ben Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, 1999.

[22] B Chandra and P Paul Varghese. Fuzzy SLIQ Decision Tree Algorithm. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 38(5):1294–1301, 2008.

[23] Liren Chen and Katia Sycara. WebMate: A Personal Agent for Browsing and Searching. In *Proceedings of the second international conference on Autonomous agents*, pages 132–139. ACM, 1998.

[24] Jie Cheng and Russell Greiner. Comparing Bayesian Network Classifiers. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, pages 101–108. Morgan Kaufmann Publishers Inc., 1999.

[25] Thomas Cover and Peter Hart. Nearest Neighbor Pattern Classification. *Information Theory, IEEE Transactions on*, 13(1):21–27, 1967.

[26] Navneet Dalal and Bill Triggs. Histograms of Oriented Gradients for Human Detection. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 886–893. IEEE, 2005.

[27] Judith E Dayhoff and James M DeLeo. Artificial Neural Networks. *Cancer*, 91(S8):1615–1635, 2001.

[28] Joaquin Delgado and Naohiro Ishii. Memory-based Weighted Majority Prediction. In *SIGIR Workshop Recommender Systems: Algorithms and Evaluation*, volume 2, page 3, 1999.

[29] William H DeLone and Ephraim R McLean. Information Systems Success: the Quest for the Dependent Variable. *Information Systems Research*, 3(1):60–95, 1992.

[30] Thierry Denoeux. A K-Nearest Neighbor Classification Rule Based on Dempster-Shafer Theory. *Systems, Man and Cybernetics, IEEE Transactions on*, 25(5):804–813, 1995.

[31] Andrew Dillon. The Evaluation of Software Usability. 2001.

[32] Susan Dumais, Edward Cutrell, Jonathan J Cadiz, Gavin Jancke, Raman Sarin, and Daniel C Robbins. Stuff I've Seen: A System for Personal Information Retrieval and Re-use. In *Proceedings of the 26th annual international ACM SIGIR Conference on Research and Development in Informaion retrieval*, pages 72–79. ACM, 2003.

[33] Bruno Dumas, Bram Moerman, Sandra Trullemans, and Beat Signer. ArtVis: Combining Advanced Visualisation and Tangible Interaction for the Exploration, Analysis and Browsing of Digital Artwork Collections. In *Proceedings of AVI*, 2014.

[34] Yoav Freund, Raj Iyer, Robert E Schapire, and Yoram Singer. An Efficient Boosting Algorithm for Combining Preferences. *The Journal of Machine Learning Research*, 4:933–969, 2003.

[35] Nir Friedman et al. Learning Belief Networks in the Presence of Missing Values and Hidden Variables. In *International Conference on Machine Learning*, volume 97, pages 125–133, 1997.

[36] Nir Friedman, Dan Geiger, and Moises Goldszmidt. Bayesian Network Classifiers. *Machine Learning*, 29(2-3):131–163, 1997.

[37] Lise Getoor and Mehran Sahami. Using Probabilistic Relational Models for Collaborative Filtering. In *Workshop on Web Usage Analysis and User Profiling (WEBKDD'99)*. Springer-Verlag, 1999.

[38] Ken Goldberg, Theresa Roeder, Dhruv Gupta, and Chris Perkins. Eigentaste: A Constant Time Collaborative Filtering Algorithm. *Information Retrieval*, 4(2):133–151, 2001.

[39] Saul Greenberg and Chester Fitchett. Phidgets: Easy Development of Physical Interfaces through Physical Widgets. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology*, pages 209–218. ACM, 2001.

[40] Nicola Guarino. *Formal Ontology in Information Systems: Proceedings of the First International Conference (FOIS'98), June 6-8, Trento, Italy*, volume 46. IOS press, 1998.

[41] Micheline Hancock-Beaulieu and Stephen Walker. An Evaluation of Automatic Query Expansion in an Online Library Catalogue. *Journal of documentation*, 48(4):406–421, 1992.

[42] John A Hartigan. Clustering Algorithms (Probability & Mathematical Statistics), 1975.

[43] Jonathan L Herlocker, Joseph A Konstan, Loren G Terveen, and John T Riedl. Evaluating Collaborative Filtering Recommender Systems. *ACM Transactions on Information Systems (TOIS)*, 22(1):5–53, 2004.

[44] Will Hill, Larry Stead, Mark Rosenstein, and George Furnas. Recommending and Evaluating Choices in a Virtual Community of Use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 194–201, 1995.

[45] Katja Hofmann, Catherine Reed, and Hilary Holz. Unobtrusive Data Collection for Web-Based Social Navigation. In *Workshop on the Social Navigation and Community based Adaptation Technologies*, 2006.

[46] Frank Hopfgartner and Joemon M Jose. Semantic User Profiling Techniques for Personalised Multimedia Recommendation. *Multimedia systems*, 16(4-5):255–274, 2010.

[47] Cezary Z Janikow. Fuzzy Decision Trees: Issues and Methods. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 28(1):1–14, 1998.

[48] Dietmar Jannach. Finding Preferred Query Relaxations in Content-based Recommenders. In *Intelligent Techniques and Tools for Novel System Architectures*, pages 81–97. Springer, 2008.

[49] Rong Jin, Luo Si, and ChengXiang Zhai. Preference-based Graphic Models for Collaborative Filtering. In *Proceedings of the Nineteenth Conference on Uncertainty in Artificial Intelligence*, pages 329–336. Morgan Kaufmann Publishers Inc., 2002.

[50] Tapas Kanungo, David M Mount, Nathan S Netanyahu, Christine D Piatko, Ruth Silverman, and Angela Y Wu. An Efficient K-Means Clustering Algorithm: Analysis and Implementation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(7):881–892, 2002.

[51] Svebor Karaman, AndrewD. Bagdanov, Gianpaolo D'Amico, Lea Landucci, Andrea Ferracani, Daniele Pezzatini, and Alberto Bimbo. Passive Profiling and Natural Interaction Metaphors for Personalized Multimedia Museum Experiences, 2013.

[52] Svebor Karaman, Giuseppe Lisanti, Andrew D Bagdanov, and Alberto Del Bimbo. From Re-identification to Identity Inference: Labeling Consistency by Local Similarity Constraints. In *Person Re-Identification*, pages 287–307. Springer, 2014.

[53] Leonard Kaufman and Peter J Rousseeuw. *Finding Groups in Data: An Introduction to Cluster Analysis*, volume 344. John Wiley & Sons, 2009.

[54] Patricia Kearney, Sarabjot Singh Anand, and CM Shapcott. Employing a Domain Ontology to Gain Insights into User Behaviour. In *Working Notes of the IJCAI Workshop on Intelligent Techniques for Web Personalization*, pages 25–32, 2005.

[55] Daniel A Keim. Information Visualization and Visual Data Mining. *Visualization and Computer Graphics, IEEE Transactions on*, 8(1):1–8, 2002.

[56] Georgia Koutrika and Yannis Ioannidis. A Unified User Profile Framework for Query Disambiguation and Personalization. In *Proceedings of Workshop on New Technologies for Personalized Information Access*, pages 44–53, 2005.

[57] Leah S Larkey. Automatic Essay Grading Using Text Categorization Techniques. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 90–95. ACM, 1998.

[58] Peter M Lee. *Bayesian Statistics: an Introduction*. John Wiley & Sons, 2012.

[59] Pasquale Lops, Marco Gemmis, and Giovanni Semeraro. Content-based Recommender Systems: State of the Art and Trends. In Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor, editors, *Recommender Systems Handbook*, pages 73–105. Springer, 2011.

[60] Tariq Mahmood and Francesco Ricci. Improving Recommender Systems with Adaptive Conversational Strategies. In *Proceedings of the 20th ACM Conference on Hypertext and Hypermedia*, pages 73–82. ACM, 2009.

[61] David Maltz and Kate Ehrlich. Pointing the Way: Active Collaborative Filtering. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 202–209. ACM Press/Addison-Wesley Publishing Co., 1995.

[62] Sean M McNee, Istvan Albert, Dan Cosley, Prateep Gopalkrishnan, Shyong K Lam, Al Mamunur Rashid, Joseph A Konstan, and John Riedl. On the Recommending of Citations for Research Papers. In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work*, pages 116–125. ACM, 2002.

[63] Manish Mehta, Rakesh Agrawal, and Jorma Rissanen. SLIQ: A Fast Scalable Classifier for Data Mining. In *Advances in Database Technology—EDBT'96*, pages 18–32. Springer, 1996.

[64] Stuart E Middleton, David C De Roure, and Nigel R Shadbolt. Capturing Knowledge of User Preferences: Ontologies in Recommender Systems. In *Proceedings of the 1st International Conference on Knowledge Capture*, pages 100–107. ACM, 2001.

[65] Stuart E Middleton, Nigel R Shadbolt, and David C De Roure. Ontological User Profiling in Recommender Systems. *ACM Transactions on Information Systems (TOIS)*, 22(1):54–88, 2004.

[66] Clifton Nock. *Data Access Patterns: Database Interactions in Object-Oriented Applications*. Addison-Wesley Boston, 2004.

[67] J Ross Quinlan. Generating Production Rules from Decision Trees. In John P. McDermott, editor, *International Joint Conferences on Artificial Intelligence*, volume 87, pages 304–307. Morgan Kaufmann, 1987.

[68] John Ross Quinlan. *C4. 5: Programs for Machine Learning*, volume 1. Morgan kaufmann, 1993.

[69] Al Mamunur Rashid, Istvan Albert, Dan Cosley, Shyong K Lam, Sean M McNee, Joseph A Konstan, and John Riedl. Getting to Know You: Learning New User Preferences in Recommender Systems. In *Proceedings of the 7th International Conference on Intelligent User Interfaces*, pages 127–134. ACM, 2002.

[70] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. GroupLens: An Open Architecture for Collaborative Filtering of Netnews. In *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative work*, pages 175–186. ACM, 1994.

[71] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. GroupLens: An Open Architecture for Collaborative Filtering of Netnews. In *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work*, pages 175–186. ACM, 1994.

[72] Paul Resnick and Hal R. Varian. Recommender Systems. *Commun. ACM*, 40(3):56–58, March 1997.

[73] Elaine Rich. User modeling via stereotypes*. *Cognitive science*, 3(4):329–354, 1979.

[74] Joseph John Rocchio. Relevance Feedback in Information Retrieval. 1971.

[75] Gerard Salton. *Automatic Text Processing: The Transformation, Analysis, and Retrieval of*. Addison-Wesley, 1989.

[76] Gerard Salton and Michael J McGill. Introduction to Modern Information Retrieval. 1986.

[77] Fabrizio Sebastiani. Machine Learning in Automated Text Categorization. *ACM computing surveys (CSUR)*, 34(1):1–47, 2002.

[78] Shokri Z Selim and Mohamed A Ismail. K-Means-Type Algorithms: A Generalized Convergence Theorem and Characterization of Local Optimality. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, (1):81–87, 1984.

[79] John Shafer, Rakesh Agrawal, and Manish Mehta. SPRINT: A Scalable Parallel Classifier for Data Mining. In T. M. Vijayaraman, Alejandro P. Buchmann, C. Mohan, and Nandlal L. Sarda, editors, *Proc. 1996 Int. Conf. Very Large Data Bases*, pages 544–555. Morgan Kaufmann, 1996.

[80] Guy Shani, Lior Rokach, Amnon Meisles, Lihi Naamani, Nischal Piratla, and David Ben-Shimon. Establishing User Profiles in the MediaScout Recommender System. In *Computational Intelligence and Data Mining, 2007. CIDM 2007. IEEE Symposium on*, pages 470–476, Honolulu, Hawaii, USA, 2007.

[81] Upendra Shardanand and Pattie Maes. Social Information Filtering: Algorithms for Automating “Word of Mouth”. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 210–217, 1995.

[82] Ahu Sieg, Bamshad Mobasher, and Robin Burke. Representing Context in Web Search with Ontological User Profiles. In *Modeling and Using Context*, pages 439–452. Springer, 2007.

[83] Michael Steinbach, George Karypis, Vipin Kumar, et al. A Comparison of Document Clustering Techniques. In *KDD Workshop on Text Mining*, volume 400, pages 525–526. Boston, 2000.

[84] Dorothea Tsatsou, Lyndon JB Nixon, Matei Mancas, Miroslav Vacura, Rüdiger Klein, Julien Leroy, Jaroslav Kuchar, Tomás Kliegr, Manuel Kober, Maria Loli, et al. Contextualised User Profiling in Networked Media Environments. In *UMAP Workshops*, 2012.

[85] Motohide Umano, Hirotaka Okamoto, Itsuo Hatono, HIROYUKI Tamura, Fumio Kawachi, Sukehisa Umedzu, and Junichi Kinoshita. Fuzzy Decision Trees by Fuzzy ID3 Algorithm and Its Application to Diagnosis Systems. In *Fuzzy Systems, 1994. IEEE World Congress on Computational Intelligence., Proceedings of the Third IEEE Conference on*, pages 2113–2118. IEEE, 1994.

[86] Lyle H Ungar and Dean P Foster. Clustering Methods for Collaborative Filtering. In *AAAI Workshop on Recommendation Systems*, volume 1, 1998.

[87] Yufei Yuan and Michael J Shaw. Induction of Fuzzy Decision Trees. *Fuzzy Sets and Systems*, 69(2):125–139, 1995.

[88] Jacek M Zurada. Introduction to Artificial Neural Systems. 1992.