

WorldFinder: A tool for finding virtual worlds

James Walkerdine, Tom Rodden
Computing Department
Lancaster University
Lancaster,
LA1 4YR, UK
{walkerdi, tam}@comp.lancs.ac.uk

ABSTRACT

The past few years has seen a rapid growth of virtual worlds available over the Internet. However, this growth has also resulted in it being much more difficult for a user to find relevant worlds. Virtual worlds contain a much richer set of information than ordinary web pages and standard web search engines do not sufficiently capture this. In this paper we present a system that aids in the finding of virtual worlds. It makes use of various information filtering techniques including collaborative filtering, and is overlaid with a simple to use direct-manipulation user interface.

Keywords

Recommender Systems, Information Filtering, Collaborative Filtering, Virtual Worlds, direct-manipulation, customisation, Java

INTRODUCTION

The past few years have seen a dramatic increase in Internet usage. Not only are the numbers of people browsing it increasing, but also the amount of information stored within it. The types of information held range from your typical Web pages, files on an FTP server, through to complex Virtual Worlds. A virtual world is a computer-generated environment in which users can navigate, interact with their surroundings and, in some cases, interact with each other. With the advent of more powerful and affordable personal computers, 3D modelling languages like VRML[18] and Java 3D[Sun], creating and using virtual worlds has become a much more realistic possibility.

However along with the increase of information on the Internet, comes the disadvantage of the increased difficulty in finding it. This is especially the case with respect to virtual worlds. Virtual worlds can range greatly in not only their content but also their physical characteristics. Apart from the environment they were created for (ie MUDS[5], VRML, etc), there are also features such as size, or graphical quality that can be used to distinguish a world. These additional characteristics are not taken into account by current search engines.

At present Internet search engines such as Yahoo[Yahoo] or Alta Vista[Alta Vista] provide the main way of finding a virtual world. However, although they are quite successful with locating web pages, their performance is poor when searching for relevant virtual worlds. Typically, the returned results offer only a title and, if they are very lucky, a brief description. Users either have to then make a judgement based on these, or examine each world in turn until they find one that meets their requirements.

The Information Retrieval world has produced a substantial amount of research over recent years to aid in the location of information. In particular is the work based around the idea of information filtering; the filtering out of information that is not of interest to the user. This paper describes the World Finder tool, a system that attempts to overcome the problem of finding virtual worlds by drawing on various aspects of this research. The paper will first endeavour to explain some of the techniques used in information filtering, before finally describing the World Finder tool.

FILTERING INFORMATION

Voss and Kreifelts[17] highlight two general problems applying to search engines. Firstly, searching is often an exploratory activity; users do not always know what exactly they are looking for. This poses a problem because the large number of results that are often produced by a search engine frequently means that the user has to reformulate their query. A query that the user is not entirely sure about in the first place. With virtual worlds this might be less of a problem, a user will likely have an idea of what characteristics he or she requires, but it does still exist to a certain extent. Secondly, there is the uncertainty of quality in search engine results. The quality of virtual worlds, as with the quality of web pages, can vary greatly and there is no way of discovering this until the user actually visits the world.

Both of these problems can be somewhat reduced by applying various information filtering techniques that have been developed. Filtering is the process of selecting items from a larger set of items. Generally this filtering is designed to separate information that might prove valuable from that which would not. The two main approaches used to achieve this are Cognitive and Social Filtering [12].

COGNITIVE FILTERING

Cognitive filtering (also known as content-based filtering) is based on the contents of the documents being searched. Techniques range from using a simple query through to creating a profile of the users likes and dislikes that are then compared with the documents content. In addition a user profile can attach weightings to certain attributes that are automatically modified when a user provides some feedback; thus hopefully making the filtering more accurate over time. Examples of such systems are InfoFinder[9] and NewsWeeder[10].

However, relying solely on content possesses some problems as described by Balabanovic and Shoham [1] Firstly, the kinds of content analysis performed are generally very high level and simplistic, a lot of useful information is not taken into account, for example the quality of the document. Secondly, with a system that makes use of user profiles, once one has been created the system will only return results that compare well against it. This presumes that the user is always going to be looking for the same types of document, which is not necessarily going to be true. Finally it is pointed out that the actual task of making users provide feedback is a difficult one. Users are going to be reluctant to constantly spend time rating documents, and so making the frequency of the rating activity and the actual effort that is required as minimal as possible is important.

Another problem mentioned by Maltz and Ehrlich [11], is that content-based filters suffer from a 'cold-start'. A new user profile is going to contain very little information about a users preferences and therefore a period of training is required before it can reflect, to a certain degree of accuracy, their likes and dislikes. This problem would again resurface if the user tried to explore previously un-visited content areas.

With respect to locating virtual worlds cognitive filtering techniques could play a significant role. A virtual world possesses certain properties that are forced onto them by their software environment (for example its dimensionality). These along with a text description of a world could provide a foundation for the 'content' of the world. However as the system would be designed for finding rather than recommending worlds, the filtering used would be of a more querying style in nature (keyword or attribute type filtering). The user states which attributes they are looking for, and worlds that do not possess any of these are filtered out. Because of this a lot of the problems related to cognitive filtering would be avoided, with the exception of the high level analysis problem.

SOCIAL FILTERING

Social filtering relies on the subjective views of other users rather than solely sticking to a documents content. The techniques used can range from users casting votes on certain attributes of the document, to actually providing a written comment to supplement it.

Collaborative filtering is the most common example of social filtering and was first introduced in the Tapestry project [6]. Users would express their opinions on incoming information such as emails and Usenet news, and then these views would be made accessible to other users. In particular, users could

generate queries that would, for example, return all the documents that a certain user liked. Since Tapestry numerous other collaborative filtering systems have been developed. The GroupLens system [13], Ringo [16] and the Bellcore video recommender [8] are all examples of ratings-based systems.

Other examples of social filtering include Maltz and Ehrlich's Active Collaborative Filtering [11], in which users can send pointers (essentially hypertext links) as recommendations to other users. In addition implicit information can be used to form the basis of recommendations, for example a web page that has been visited by a lot of users is likely to be useful to others [7]. Similar work has been carried out by Chalmers et al. [3] in which the consistencies of web use have been used to build up connections and parallels, without relying on the content.

Although Collaborative systems do solve the majority of Content-based filtering problems, they do introduce some new ones [1]. Firstly, when a new document is created a vicious circle situation develops; a new document initially has no recommendations and will not receive any until a user views it, but a user is unlikely to view it unless it has been recommended to them. In a sense this is a collaborative systems own version of the 'cold start' problem.

Another problem is the assumption that users tastes show some form of correlation. If a user possess tastes that differ from the norm, then the recommendations they receive would not be of much use. On top of this there is the fact that people are likely to possess contradictory opinions which again would not help a browsing user. Resnick and Varian [14] also point out that users are unlikely to spend time providing a recommendation unless there is some kind of incentive. In addition, they point out if anyone can provide recommendations then there is nothing to stop the owners of documents repeatedly recommending their own work and providing only negative recommendations for their rivals.

Capturing information on various non-quantitative properties of virtual worlds, such as quality or feel, can not be effectively achieved with simple attributes, especially as people will possess differing views. Making use of a form of social filtering would provide an invaluable aid in finding virtual worlds. One of the main problems was that of user feedback. The system should therefore be designed to be simple, keeping the amount of user input down to a minimal is important. Using simple social filtering techniques, such as voting, would be more suitable rather than expecting users to provide written reviews.

HYBRID – A MIXTURE OF BOTH

Some systems have been developed which make use of both types of filtering techniques. The main advantage this possesses is the reduction in the problems experienced when the two techniques are used separately. For example the 'no recommendations' problem experienced in social filtering. Using a hybrid approach, even if no other user has recommended the document by voicing his/her opinion, it can still be recommended based on its content. Examples are the Fab system [1] or the Filterbot model [15].

Though a hybrid system would likely be more successful in filtering information than the separate techniques, to make full use of the two techniques the resulting system would have to be much more complex. Users would have to provide feedback for both types and this could make the system un-useable. One solution to this would be to only take the simplest ideas from each technique, producing a simpler system but still possessing a certain degree of filtering. Ultimately it will be the amount of filtering required that makes the decision.

DESCRIBING A VIRTUAL WORLD

A virtual world contains a very diverse set of information that can not be realistically captured in a short text description. Imagine the difficulty in trying to describe part of the real world or even a simple room, in only a few words. A virtual world is subject to more constraints than the real world and some attributes of a world can therefore be defined objectively, for example what dimensionality it exists in, or whether more than one user can access it at once. However, like the real world, there are various attributes that can not be easily described or would be recapitulated differently depending on the person.

Our solution is to use four types of descriptor; Objective attributes, Subjective attributes, a text description and the browsing users. Objective attributes are used to capture the elements of the world that are either true or false (e.g. Can you fly in this world?). Subjective attributes capture the elements that are likely to be interpreted differently by different people (e.g. the size of the world). The text description is used to describe what the world is actually about and its purpose. Finally the browsing users provide non quantifiable information about the world, for example, is the world actually any good?

Table 1 illustrates the attributes that have been proposed to describe a virtual world with, what category they fall in, and their range of possible values. Although subjective attributes possess the advantage of providing a far greater range of information, they also have the drawback of being difficult to map across different users. For example, two people could rate the size of the same world differently.

Subjective Attributes Rating range of 1 - 5	Objective Attributes Rating range of Yes or No
Size	3D ¹
Graphics	2D ¹
Realism	1D ¹
Complexity	Video
	Picture
	Text
	Sound Effects
	Music
	Speech
	Walk
	Fly
	Single User ²
	Multi User ²
	Collision Detection
	Gravity
	Remembrance
	Network Links ³

Table 1 - World Attributes

Notes

1, 2. Although it might seem that only one of these attributes could exist at a time for a given world, they are represented as individual characteristics because a user might want to search for worlds, that possess one or the other. So for example they might want to search for both 3D and 2D worlds.

3. Upon creation of a world description object, the author enters the exact number of network links it possesses. When a query is created the different range of network links available are represented by five different groups; No links, 1 – 2 links, 3 – 5 links, 6 – 10 links, and 11+ links.

WORLD FINDER

World Finder is a collaborative tool for finding virtual worlds. Written in Java, it is a client-server based system in which a user at the client can query a database of virtual world description objects held at the server. The use of user profiles allows a query to be composed of both objective and subjective attributes (which are then normalised) and also provides support for various tailorability functions. The search engine caters for different weightings of attributes, and all information that is held at the server, is stored as objects. Finally the software supports various forms of collaboration to aid in the finding of relevant worlds.

As the information used to describe a virtual world is content and social based, both types of filtering are used. However unlike as in some of the examples mentioned earlier, only simple techniques will be used that do not require much user feedback. The task of finding a virtual world needs to be as straightforward and easy as possible. Complex filtering techniques would only hinder this goal and provide unnecessary functionality.

The cognitive filtering viewpoint deals with the objective and subjective attributes and the text description. For the objective attributes and text description it simply involves filtering out any worlds that possess none of the required attributes or keywords. Subjective attributes are more complicated due to the fact that users

will rate them differently. To deal with this we propose the use of a user profile that attempts to capture user bias with such attributes, ie whether they underrate/overrate a worlds size, etc. These profiles can then be used to modify subjective information that is passed to and from the user. The number of objective, subjective and keyword matches can be used to formulate a percentage match for a world.

The social filtering viewpoint deals with the information provided by browsing users. In order to keep user input to a minimal, a simple voting mechanism is used to allow users to express their opinions on a world. In addition non-implicit information can be derived by noting the number of visits to a world and the most 'liked' worlds (drawn from the voting results). This information can be compiled into lists to provide addition information to the user.

By using a hybrid filtering approach a much broader range of information can be presented to browsing users. This additional information would hopefully make it much easier for a user to find a relevant virtual world.

Figure 1 shows a simplified architecture of the software and what objects are stored within it.

World Finder Client – This is the application that users use to browse and search for virtual worlds. It communicates with the World Finder Server by making use of an RMI network connection.

Profile Objects – Objects that store information about the users who use the World Finder client. The information stored includes their bias in regards to the subjective attributes and also their customisation details.

World Finder Server – The server which all the clients connect to. The server manages the various objects that it stores and also contains the search engine that matches a users query against the World objects.

User Objects – These store a user's username and password. The user has to logon before they are granted access to the server.

World Objects – All the information that World Finder uses to describe a world is stored in a World Object. This includes content information and social information (votes etc).

Query Objects – A query object represents the query that the user has created to be processed. It is originally produced by the client, and is passed to the server to be processed and also stored. The stored queries can then be later reused to provide awareness information such as 'What are the most used queries?'

Template Objects – Templates are queries that can be accessed or modified by other users. Apart from containing query information they also contain the access rights for other users, a modification history and which query category they have been placed in (see later).

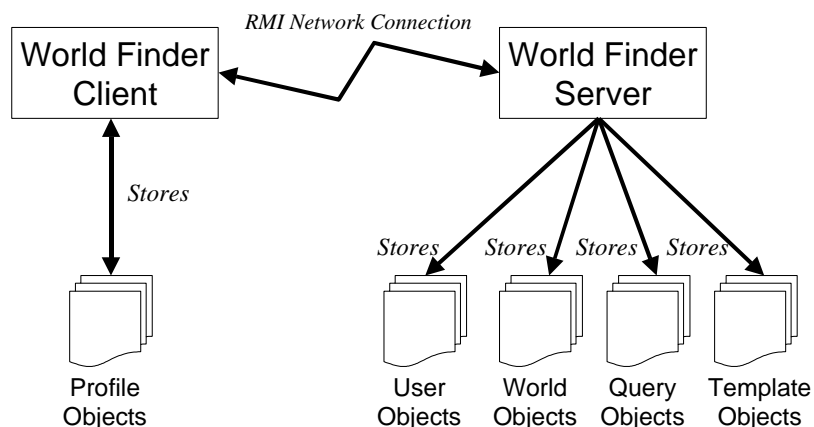


Figure 1 - Simple overview of the World Finder architecture

CREATION OF WORLD OBJECTS

World Finder contains two separate packages for browsing and for registering virtual worlds. The main browser software deals with the searching and collaboration aspect of the software (the bulk of this document describes this part), and the registry software deals with the actual creation and registration of

world objects. The registry software is small and simple, and so to provide ease of use it was decided to make it an applet that would be run from a web browser.

Before worlds can be registered a user profile has to be created to capture user bias in regards to the subjective attributes; this process is described in more detail in a later section. Once the profile has been set up a user can then register his or her worlds by filling in a straightforward form. The user is encouraged to provide as much information about the world as possible; ultimately it will be this information that will decide if a user will visit his or her world. Any subjective data is then normalised in accordance with their profile and finally all the details of the world are sent to the World Finder server where they are created as a World Object and stored. Figure 2 shows an example of a world being registered.

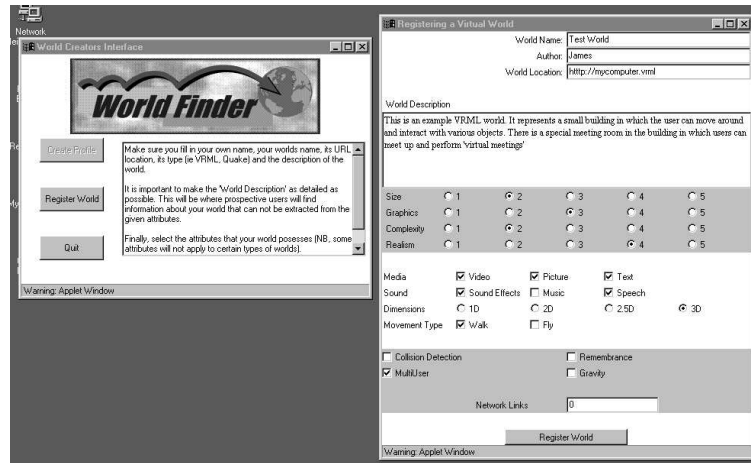


Figure 2 - Registering a virtual world

FORMLESS QUERIES

As a query could be potentially complex it was important to provide a simple but powerful way of creating them; endless form filling would not be appealing to your average user. The creation of queries in World Finder is carried out through direct-manipulation. In World Finder all the possible types of attributes are displayed to the user as icons which can then be dragged into a query window.

Once an icon has been placed in the query window it can be negated (a user wants to find a world which 'does not' possess this attribute) or have its priority set. World Finder makes use of six priority levels that are calculated from the vertical position of the attribute in regards to the query window. Negated attributes have a small red cross in the corner of their corresponding icon.

Obviously text keywords have to be typed in, but after this they can then be selected and converted into icons. These can then be treated exactly in the same way as the standard physical attribute icons.



Figure 3 - Creating a query

THE SEARCH ENGINE

The World Finder search engine is of a simple design. In its most basic form the engine merely matches the attributes specified in the query with the attributes contained within a world object. However factors such as an attributes priority or whether a range of values are used to represent it (ie it is subjective), also have to be taken into account.

Priority Levels

Queries in World Finder have six priority levels, each one possessing a percentage weighting, Figure 4 illustrates this. The first step when a query is processed involves finding the total possible percentage that can be obtained if the query is matched perfectly. This is simply calculated by adding up each attributes percentage value in the query.

For the query shown in Figure 4 the total would be:-

Level 1 there are 3 attributes: $3 \times 40 = 120$

Level 2 there is 1 attribute: = 25

Level 3 there is 1 attribute: = 15

Level 4 there is 1 attribute: = 10

Level 5 there is 1 attribute: = 7

Total: = 177

As worlds are examined they achieve a score based on how well their attributes match the queries, the points they get for each matching attribute is again dependent on the priority level. For example, suppose the query in Figure 3 is used and one particular world possesses the 3D and Multi-User attributes. This world would obtain a total score of $40 (3D) + 7 (Multi-User) = 47$. Resulting in a percentage match of 27% ($47/177 \times 100$).

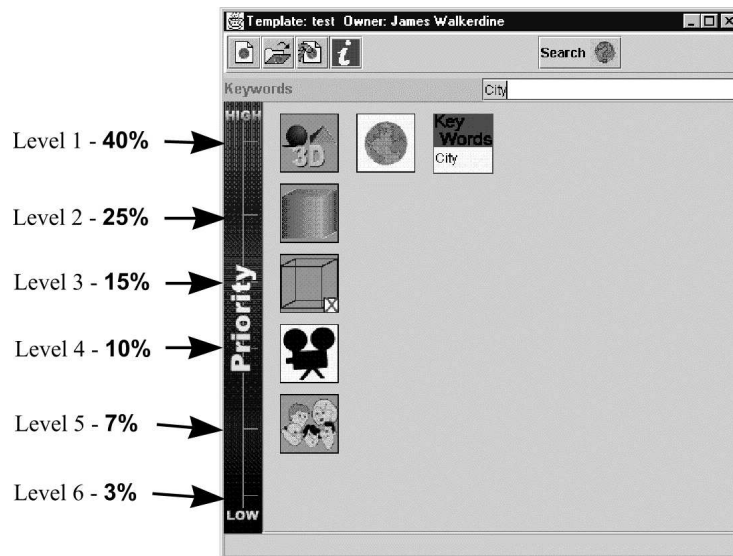


Figure 4 - Differing Priority levels of a Query

Spread of Subjective Attributes

Subjective attributes are hard to map across different users. Even though World Finder makes use of User Profiles to capture user bias, there is still the possibility that the system might regard an attribute as one thing, but the user see it as another. For example, the system might predict that a user would rate the size of a given world as 4, but the user actually rates it as a 5. To take into account the possibility that the system might predict wrongly an attribute value, the search engine calculates the attributes score based on the spread of possible values.

DISPLAYING THE RESULTS

Once a query has been created and searched with, the list of search results is displayed along with their percentage match. A user can then make a selection from this list in order to find out more information about the world. Figure 5 illustrates this.

The newly displayed window contains all the information about the world that was entered in its initial registration. Rather than displaying it in text form, a row of icons is used to capture the main characteristics of the world. The idea behind this is that a browsing user can get an easy 'at a glance' idea of what the world is like without actually having to read any text. Below this bar is the description of the world and a fuller text version of the worlds attributes.

In addition a user can access the opinions of other users from this window and, in turn, add their own views upon visiting the world. Voting is described in more detail in the next section.

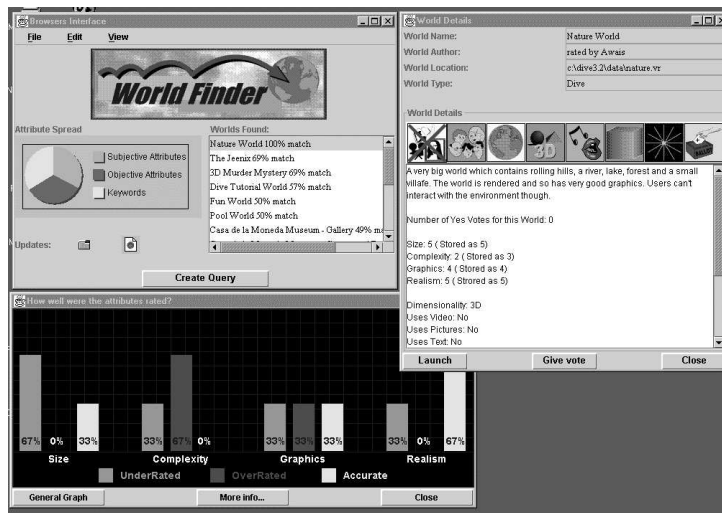


Figure 5 - Displaying the results

SOCIAL BROWSING FOR VIRTUAL WORLDS

Collaboration in World Finder is represented in four main areas: -

- The making and usage of templates
- The creation of personal profiles that incorporate differing 'views' of created categories.
- User awareness, for example a list of the most popular worlds or most frequently used templates.
- The ability for users to voice their opinion on worlds which can then be taken into account by other prospective users.

TEMPLATES

In simple terms a template is merely a World Finder query that can be accessed, modified and used by other users. It provides a mechanism that allows more than one user to jointly create a query, or alternatively provides users with access to a range of pre-created queries.

The original creator always owns a template, and thus only the owner can delete it from the World Finder server or modify user access rights. By making use of access rights, the owner can specify which users can access the template and which can modify it. Typical cases of this in use would be when a template is designed to be accessible to all users but only modifiable by the owner (i.e. providing a template for all users to search with); or restricting read and write access to those users who are collaboratively creating a template.

For those users who do have write access but do not own the template, only the attributes and keywords that make up the query can be modified. This means that they can alter what the template/query is composed of, but can not alter the access rights of the template.

Each time a change, be it made by the owner or other users with write access, their name is stored. In addition any comments that relate to the change can also be stored, and so in this way a history of changes is kept which other users can then later view. Because templates are ultimately only search engine queries, it was decided that to include a versioning system for template modifications would only result in lowering the usability of the software.

USER PROFILES

Each user of World Finder has his or her own personal profile. This not only allows the software to be customised to individual needs, but also makes up an important part of the search process. When a user first uses the World Finder software they have to 'rate' the subjective attributes of a test virtual world. From this the software can generate a profile which attempts to capture user bias. The profile is then used to modify any subjective information that is passed to or from the user. This means that any world which the user creates and then registers is modified to a standard norm Any world description object that the user looks at is modified from the standard norm to the users norm. Figure 6 and the example below illustrate this point.

John is a user of the World Finder system. When he first starts the software he examines the given test world and rates the subjective attributes as follows: *Size* = 2, *Graphics* = 3, *Realism* = 2 and *Complexity* = 3.

As mentioned earlier, subjective attributes are rated on a range of 1 to 5, where the software regards the value of 3 to be the norm. So John believes that the test worlds size and realism are both less than average. The system notes this and creates Johns user profile with the values: *Size* = -1/+1, *Graphics* = 0, *Realism* = -1/+1 and *Complexity* = 0. Unless the value is equal to 0, the first part represents the modifier that is used when John examines a world, so for *Size* this is -1. The second part represents the modifier that is used when John creates a world, so for *Size* this is +1. If the value is 0, then no modifications are made either way.

John now decides to register his own virtual world (we are only using subjective attributes in this example). He gives it the following attribute ratings: *Size* = 3, *Graphics* = 4, *Realism* = 2 and *Complexity* = 2. Before his world is registered it is modified in accordance with his user profile. So his *Size* rating will be increased to 4, *Realism* will be increased to 3, and *Graphics* and *Complexity* will stay the same at 4 and 2 respectively. The now 'modified' ratings are stored in the world database. Note if John examines his own world, his profile will subtract one from the stored value of *Size* and *Realism* and thus they will be the same as his original values.

Jane has also started using the World Finder system. She rates the test world as follows: *Size* = 4, *Graphics* = 4, *Realism* = 3 and *Complexity* = 2. Her generated profile would be: *Size* = +1/-1, *Graphics* = +1/-1, *Realism* = 0 and *Complexity* = -1/+1. She believes the *Size* and *Graphics* of the test world are above average, *Realism* is average and *Complexity* is below average.

Jane notices that John has registered a virtual world and so decides to examine it. Remember that the values of Johns attributes stored in the database are: *Size* = 4, *Graphics* = 4, *Realism* = 3 and *Complexity* = 2. These values are modified by Jane's profile before they are displayed to her. So the *Size* rating will be increased to 5, *Graphics* will be increase to 5, *Realism* will stay at 3 and *Complexity* will be reduced to 1.

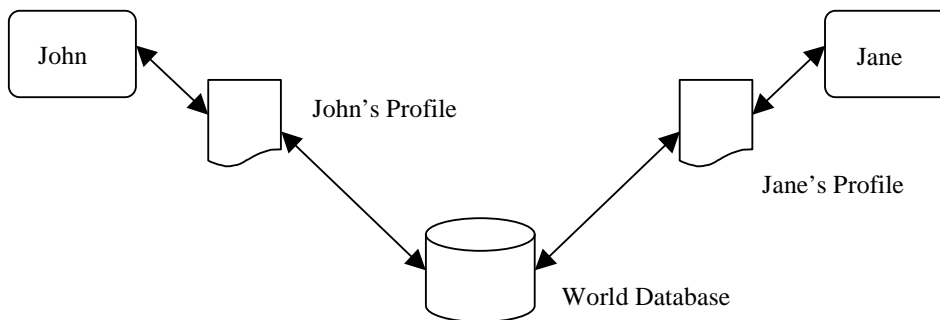


Figure 6 - User Profile System


Apart from modifying subjective ratings, the profile also stores the users view of the category tree (as described later).

CATEGORISATION AND DIFFERING VIEWS

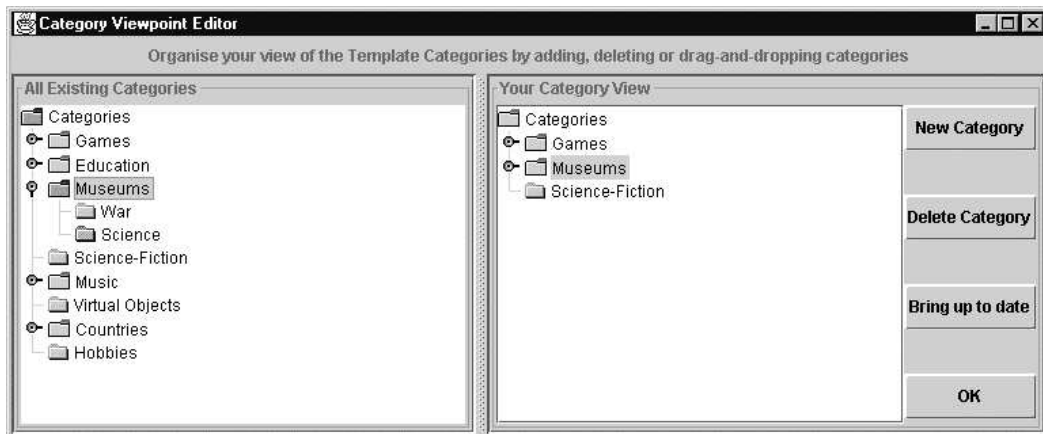
In the likely event that hundreds, if not thousands, of templates are created, World Finder makes use of categories to provide some form of organisation. When a template is created the user assigns it to a category; this can be one already in existence or be newly created by the user. A tree structure is used to store these categories and to facilitate browsing. For example, John creates a template that he believes helps in searching for city style 3D virtual worlds. He wants it to be available to other users and so looks through the existing categories to find a suitable place to locate it. Unfortunately he does not find one and so, instead, decides to create his own called ‘City Worlds’, in which he then stores his template.

Due to the fact that different categories will be of interest to different people, World Finder incorporates user views. Each user can modify the tree by renaming, adding, deleting or moving a category. These different views are stored locally in the user profile and are then mapped onto the raw category data stored at the World Finder server. Figure 7a illustrates the theory behind this and Figure 7b shows an actual screenshot. In Figure 7b the tree on the left represents all the categories that have been created, whereas the tree on the right represents the users view.

One issue that had to be dealt with, are updates. If a user adds his/her own category how are other users informed of this? World Finder overcomes this by highlighting the tree path to any new categories the next time a user logs on. In this way a user can see which categories are new, drag them across to their view if needs be, or simply reset the update information.



(a)



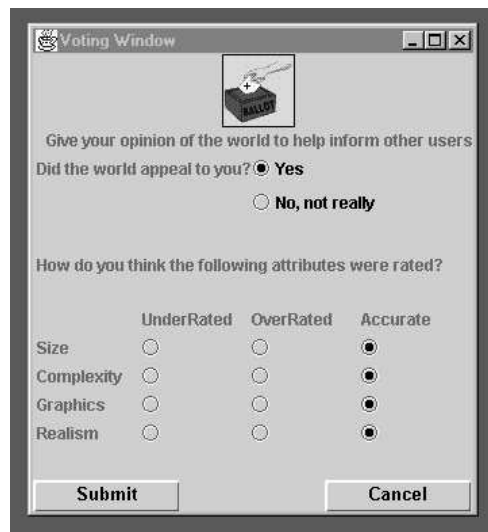
(b)

Figure 7 - (a) Differing views, (b) Category views in World Finder

VOTING AND AWARENESS

As previously mentioned, World Finder also captures social information about virtual worlds. To do this a facility is provided to allow users to 'vote' on whether or not they like the world, and whether the subjective attributes in the world description accurately reflected the world. Figure 8a shows how a user can cast their vote. This information is represented via bar charts, providing an immediate at a glance group opinion of the world. (See Figure 5)

When a query is sent to the World Finder server apart from being processed, it is also stored. These queries, along with the voting information, can then be later reused to generate asynchronous awareness information for other users. Currently World Finder allows users to view the ten most popular virtual worlds (the worlds with the highest proportion of yes votes), and the ten most visited worlds (the worlds that have been launched the most). In addition, World Finder makes a note of the worlds a user has visited so that a user history can be developed. Figure 8b shows an example Top 10.



(a)

World Name	World Author	World Type	World Location	Viewings
Nature World	rated by Awais	Dive	c:\dive3.2\data\...	3
Virtual Lounge	rated by James	VRML	http://www.iun...	2
Casa de la Mon...	rated by James	VRML	http://www.fnmt...	1
Music Box - Da...	rated by James	VRML	http://www.geoc...	1
Casa de la Mon...	rated by James	VRML	http://www.fnmt...	0
Crystal Beads	rated by James	VRML	http://www.mar...	0
Dive Tutorial W...	rated by Awais	Dive	C:\dive3.2\data\...	0
Enterprise	rated by James	VRML	http://www.mar...	0
Fun World	rated by Walter	Dive	c:\dive3.2\data\...	0
Hina dolls	rated by James	VRML	http://www.mar...	0

(b)

Figure 8 – (a) Casting your vote, (b) Top Ten most visited worlds

SECURITY IN WORLD FINDER

With the inclusion of collaboration and the use of user profiles another issue that raises its head is providing system security. Templates, World Objects and User Profiles should only be altered or deleted by their respective owners.

When a user first uses World Finder they are provided with a unique login name which is generated from their surname. After the user has then in turn provided a password, a user object is created on the World

Finder server which holds the users full name, login name and password. Each time a user intends to use World Finder they have to correctly logon (this also allows the system to set up the software in accordance to the users profile). As all login names are unique and ownership of an object is represented by a login name, security in the system is insured.

CONCLUSION

This paper has presented World Finder, a tool to aid the user in locating virtual worlds. Virtual worlds contain a vast diversity of information, a spread of data that is not effectively captured and represented by traditional Web search engines. World Finder has been designed to resolve this and makes use of various information filtering techniques that have been developed in the Information Retrieval field.

In World Finder, the contents of virtual worlds are described by a set of attributes. A query is also constructed from these attributes and then a simple form of Cognitive filtering is used for when worlds are being searched. In order to capture non-quantifiable information such as world quality, Social filtering is also supported with the inclusion of a virtual world voting system.

By providing this additional information and by supplying ways in which to filter out unwanted material it is hoped that World Finder would provide the user with a much more powerful and successful tool for locating virtual worlds.

There are various aspects of the World Finder system that could be further expanded. For a start the system could be made more generic so that it does not just have to deal with virtual worlds. To achieve this however, the attributes used to construct and describe information would also have to be generic. To devise a set of generic attributes to describe all types of information would not be a simple task.

Another area of interest surrounds the actual storage of queries and the possibilities this might offer. It is perfectly feasible for the stored queries themselves being recommended by other users. In addition facilities could be provided to actually query the stored queries themselves.

REFERENCES

- [1] Balabanoic, M., Shoham, Y., "Content-Based, Collaborative Recommendation" in *Communications of the ACM*, March 1997, Vol 40, No 3, pp. 66-72.
- [2] Belkin, N. J., Croft, W. B., Information Filtering and Information Retrieval: Two Sides of the Same Coin?. In *Communications of the ACM*, December 1992, vol.35, No.12.
- [3] Chalmers, M., Rodden, K., Brodbeck, D., The Order of Things: Activity-Centred Information Access. In *Proceedings of the 7th International Conference on the World Wide Web*, Brisbane, April 1998, pp. 359-367.
- [4] Cousins, S.B., Paepcke, A., Winograd, T., Bier, E.A., Pier, K., "The Digital Library Integrated Task Environment (DLITE)" in *Proceedings of DL 97*, Philadelphia, USA, pp. 142-151.
- [5] Curtis, P., Dixon, M., MUDS grow up: Social Virtual Reality in the Real World. In *Proceedings of IEEE Computer Conference*, January 1994
- [6] Goldberg, D., Oki, B., Nichols, D., Terry, D.B., "Using Collaborative Filtering to Weave an Information Tapestry" in *Communications of the ACM*, December 1992, Vol 35, No 12, pp. 61-70.
- [7] Hill, W. C., Hollan, J. D, Wroblewski, D., McCandless, T., Edit Wear and Read Wear. In *Proceedings of CHI'92*, Monterey. 1992, Addison-Wesley, pp. 3-9.
- [8] Hill, W., Stead, L., Rosentain, M., Furnas, G., Recommending and Evaluating Choices in a Virtual Community of Use. In *Proceedings of CHI'95*. Denver, CO. May 1995. ACM.
- [9] Krulwich, B., Burkey, C., Learning user information interests through extraction of semantically significant phrases. In *Proceedings of the AAAI Spring Symposium on Machine Learning in Information Access* (Stanford, Calif., March 1996)
- [10] Lang, K. Newsweeder:, Learning to filter netnews. In *Proceedings of the 12th International Conference on Machine Learning* (Tahoe City, Calif.) 1995

- [11] Maltz, D., Ehrlich, K., Pointing The Way: Active Collaborative Filtering. In Proceedings of CHI'95, Denver, CO. May 1995
- [12] Malone, T. W., Grant, K. R., Turbak, F. A., Brobst, S. A., Cohen, M. D., Intelligent Information-Sharing Systems. In Communications of the ACM, May 1987, Vol. 30, No. 5, pp.390 – 402.
- [13] Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Reidl, J., GroupLens: An Open Architecture for Collaborative Filtering of Netnews. In Proceedings of CSCW'94, NC. 1994
- [14] Resnick, P., Varian, H.R., "Recommender Systems" in *Communications of the ACM*, March 1997, Vol 40, No 3, pp 56-58.
- [15] Sawar, B. M., Konstan, J. A., Borchers, A., Herlocker, J., Miller, B., Reidl, J., Using Filtering Agents to Improve Prediction Quality in the GroupLens Research Collaborative Filtering System. In Proceedings of CSCW'98, Seattle, Washington. 1998
- [16] Shardanand, U., Maes, P., Social Information Filtering: Algorithms for Automating "Word of Mouth". In Proceedings of CHI'95. Denver, CO. May 1995
- [17] Voss, A., Kreifelts, T., SOaP: Social Agents Providing People with Useful Information. In Proceedings of Group 97, Phoenix, Arizona, 1997. ACM.
- [18] VRML97, International Standard ISO/IEC 14772-1:1997, "the Virtual Reality Modeling Language"