An empirical user profile adaptation mechanism that reacts to shifts of interests

Jérôme Picault[1] and **Myriam Ribière**[2]

Abstract. User modelling is a very difficult problem as it deals with complex, changing user behaviours. Lot of research has been done to understand long-term user preferences about e.g. movies, television, music, books and has lead to successful systems (such as TiVo, Amazon, etc.). But as stated in [3], *"many user interests are ephemeral and transient"*; however relatively few research work address this particular topic.

In our paper, we present a novel empirical framework to handle the problem of the adaptation of the user profile in an efficient manner to different usage patterns, which characterize different evolutions of the user's interests, such as persistent interests, periodic interests or shift of interests.

1 INTRODUCTION

One key aspect for successful user profile adaptation resides in the ability of detecting shifts of interests, i.e. quite rapid or fast emerging changes in the behaviour of users that may lead to the apparition of new long-term interests, or disappear as fast as they appear. Shift of interests may be induced for many reasons such as professional change, external events (e.g. elections or the Olympics). If undetected over a long time, abrupt interest changes may result in significant degradations in the quality of the filtering and loss of satisfaction among users during the period of non-detection. It is extremely difficult to find a good solution to that problem: if introduced too early, "new" interests may not be relevant enough from the user's point of view; if introduced too late, the new interests may already be obsolete. This problem has already been described for many years. Thus, in [2] is described a hybrid user model consisting in both a short-term and long-term model of the user's interests for news filtering. The method employs the short-term interests first (based on most recent observations). If news cannot be classified with the short-term model, the long-term model is used. This model introduces a clear separation between the different kinds of preferences, but lack of flexibility, as there is no direct impact from one type of preference to the other one. In [9] the approach consists in using clusters of past user queries to infer long-term user preferences. Each cluster is associated to a time stamp and infrequently used clusters are removed from the set of clusters in order to effectively model the user interests. If this method may be relevant for long-term evolution of preferences, it is not appropriate enough for rapid evolution of preferences that will remain undetected (or over-detected).

The approach we present in this paper consists in a mechanism for lifelong user profile adaptation, i.e. for adaptive personalised information filtering systems (such as personalised news) or personalised content push systems, (such as targeted advertising) which allows adapting the user profile to very changing users' behaviours and where the adaptation of the user profile relies on a combination of several complementary mechanisms:

- a concept history stack to introduce new preferences in the user profile;
- a preference update formula to update the weight of preferences according to collected implicit feedback on consumed content;
- a personalised decay factor parameter to remove concepts that are not of interest any more.

This paper is organized as follows: section 2 presents our general approach to adaptation of users' preferences, section 3 describes an empirical approach to the problem of shifts of interests. In section 4 we present a set of early experimental results. Finally, some conclusions are drawn in section 5.

2 ADAPTATION OF SEMANTIC PREFERENCES

In a typical user modelling process (fig. 1) are defined three main tasks: the acquisition of the user profile, the evolution (or *adaptation*) of the user profile according to his use of the system and the changes in user interests, and the exploitation of the user profile.

The method we propose allows the creation and update of the user profile by exploiting the usage logs of each user. As most of methods exposed in various recent research [4][5][6], we are exploiting data such as vector of keywords/concepts, user feedback and some indicators to create and update user preferences, such as the time spent on the document, its length, the mean number of documents read by a user.

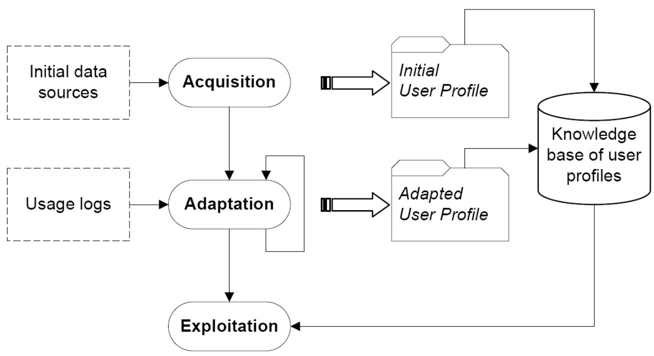

Figure 1. Typical user modelling process

[1] Motorola Labs, France, email: jerome.picault@motorola.com
[2] Motorola Labs, France, email: myriam.ribiere@motorola.com

The user profile adaptation process is done off-line.When triggered, a long-term adaptation session is composed of two main sub-processes: (1) insertion/removal of concepts in the user profile, and (2) update of concept weights in the user profile. Both sub-processes (1) and (2) are using some data sources, which are collected inside a Content Consumption Analysis (CCA) database. Then, for each line in the CCA database, and for each concept related to a content item: if the concept does not already exist in the user profile then we apply the process "insertion/removal of concepts in the user profile" (section 2.1) else we apply the process "update concept weight" (section 2.2).

2.1 Concept history stack

In our personalisation model, a preference of a user is described by a weighted concept. The analysis of the consumed content for a user helps determining from the content annotations which concepts appear and how often they occur during a period. We are introducing a *concept history stack* as a mechanism to keep track of all *potential* interests of a user and to detect if an interest is confirmed over time and can be considered as a preference. Thus, the history stack contains all the concepts that are currently in use in the profile as preferences and all concept candidates to be part of the user profile preferences. An object in the history stack is described by: the concept URI; the number of occurrences N_{occ_c} of the concept c in the consumed content; the date of first appearance of the concept d_c. These data enable the computation of the *concept persistence* P_c, at a given time D (when the adaptation process is triggered), for each concept c:

$$P_c \quad = \quad N_{occ_c}/(D - d_c) \tag{1}$$

The computation of P_c is naturally reflecting a decay factor - or gradual forgetting of preferences which is sufficient to handle gradual interest changes [8]. A main issue concerns user shift of interest [12], since an important interest of one day can potentially create a new preference in the profile that will take a week to disappear from the profile (based on the natural decay factor).

This metric is not impacting directly the weight of the user preferences, but the decision for keeping or not a preference in the profile.The following heuristics should apply:

- A concept occurs once and its number of occurrence is confirmed with time with roughly the same level, this concept can be introduced after a period as a long-term preference of the user.
- A concept occurs once and its occurrence is very high on a short period, and then disappears very quickly. Even if this concept can be considered as a preference during a period of time, it must be removed very fast from the preferences, once the interest is over.
- A concept occurs once but the persistence value is not very high and even if confirmed in time, it does not constitute a significant interest for the user. In that case this concept will never become a user preference.
- A concept occurs and becomes a preference as in the first case, but disappears with time. It must at a some point of time be removed from the preferences.

The introduction and removal of concepts in the user profile is decided by comparing the value of the concept persistence with thresholds determined empirically.

2.2 Preference weight update

The adaptation of semantic user preferences does not only consists in adding or removing preferences, but also in updating the concept weights of newly introduced preferences in the user profile, based on the analysis of consumed content. A possible model, based on the work of [10] is:

$$w_{new} \quad = \quad w_{old} + f * r * e^{-\beta * x * y} * \log \frac{t}{l} \tag{2}$$

The factor w_{old} represents the current weight of the concept. f is the relevance feedback factor given through an analysis of the content consumption. r is the rank assigned to the content by the personalised retrieval system . $\log \frac{t}{l}$ incorporates the time t spent reading or watching an item of length l and operates as a normalizing factor. $e^{-\beta * x * y}$ is used to follow the personalised non-linear change of the concept's weight according to user usage history data. x represents the mean number of content that the user is consuming per update period; y represents the number of consumed content where the concept appears in the set of metadata. β is a positive constant, which may take different values in case of positive or negative implicit feedback.

We conducted some simple initial qualitative simulations of evolution of the weight of semantic interests, in order to validate a correct behaviour of the update formula, on simple use-cases, such as the impact of the variations of the user feedback, of the content consumption or of repeated negative feedback on on the concept weight. In all cases, we obtained expected and coherent results.

Although this mechanism was quite interesting as first approach, it did not fit all our requirements in terms of changes of user behaviours. In the next section, we introduce an aditional mechanism that combines these two independent and separate mechanisms (use of concept history stack, concept weight update) into a new one, which has for effect to better adapt to changing patterns of content consumption.

3 IMPACT OF SHIFT OF INTERESTS

The notion of shift of interest can be described as the rapid change of content consumption associated to a concept. This definition enables testing the impact of shift of interests in the mechanism of long-term adaptation described above. Two cases may be studied: the impact on a semantic interest which is already in the user profile; and the impact for concepts that were not yet introduced as semantic interests in the user profile. When the semantic interest is already in the user profile, a simple simulation allows to observe what happens when:

- The number of occurrences of a concept increases suddenly during a short time frame (e.g. sudden interest for a new concept or for a concept which importance was low beforehand). This rapid increase of concept consumption can be confirmed with time or not.
- The number of occurrences of a concept decreases suddenly during a short time period (e.g. short high interest in elections).

A simple experiment (fig. 2, 1st curve) shows that the model reacts quite well in a shift of interest which is confirmed with time: in case of rapid increase of concept consumption during one update period, the weight of the concept increases rapidly too. However, if this increase of concept consumption is not confirmed with time (or if there is a suddenly decrease of consumption of an interest that was high), the weight of the concept in the user profile remains too high and does not decrease unless there is some repeated negative implicit feedback. This phenomenon is due to the fact that once a concept has been introduced as a semantic user interest (because it was exceeding the persistence value in the concept history stack), the evolution of weight is not submitted to any decay factor based on the variation of consumption: if no new concept c is consumed, its weight in the semantic user profile will remain constant (which may generate unwanted recommendations related to this concept!), but its persistence value will decrease. Later on, this preference will be removed, because the persistence value will be too low, but this will be sudden, which is not appropriate either.

So, the weight update formula should be modified: if a user is showing a particular interest for a concept during a 1-day period, this concept will be quickly taken into account in preferences (correctly taken into account by the formula), but if this interest is not confirmed the following day, this preference is expected to disappear quickly in the user profile. Thus, a "fast enough" decay factor has to be added to the concept weight, so that very temporary, although high interests are penalized.

In order to do so, some solutions to adapt the size of the time window which is considered for profile adaptation have been proposed [7]. We adopted another method: an interesting indicator for detecting shift of interest could be the analysis of the variation of the concept persistence over time. $\frac{\Delta P_c}{\Delta t}$ computed at time t gives the local tendency of the evolution of the interest for concept c at time t. If $\frac{\Delta P_c}{\Delta t} \approx 0$ the user interest for concept c tends to be stable during the time window; if $\frac{\Delta P_c}{\Delta t} > 0$ the user interest for concept c tends to increase during the period; if $\frac{\Delta P_c}{\Delta t} < 0$ the user interest for concept c tends to decrease during the period. The precision of detection of shifts if interest is governed by the choice of the time window chosen for computing the variation of concept persistence. The intensity of interest shift is determined by the amplitude of variation of $\frac{\Delta P_c}{\Delta t}$. A real shift of interest happens when $\frac{\Delta P_c}{\Delta t} \ll 0$ or when $\frac{\Delta P_c}{\Delta t} \gg 0$.

This can help define a refined update formula, which adds an additional personalised decay factor taking into account the diminution of content consumption:

$$w'_{new} \quad = \quad w_{old} + f * r * e^{-\beta * x * y} * \log \frac{t}{l} + \lambda \frac{\Delta P_c}{\Delta t} \qquad (3)$$

where λ is a positive constant.

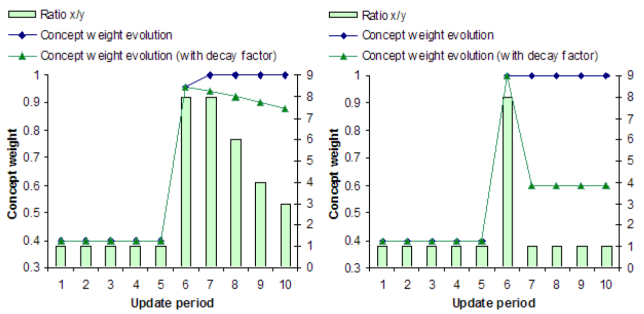

Figure 2. Impact of shift of interest on the concept weight, confirmed over time (left) or not (right), without (1st curve) and with (2nd curve) the introduction of a personalised decay factor

The curves on fig. 2 illustrate the impact of formula (3). Compared to formula (2), the new system behaves better when a concept appears suddenly but its number of occurrences decreases more slowly, or when a concept appears suddenly and its number of occurrences decreases quickly.

4 EXPERIMENTAL RESULTS

The user profile adaptation mechanism described in this work has been implemented within a proof-of-concept prototype. In order to study the validity and the soundness of the proposed model, tune parameters and observe the behaviour of the user adapation model, we have tested our prototype on a real data. We present here early results of our approach.

4.1 Dataset description

In order to evaluate the pertinence of our approach, we performed several tests on a real dataset. We tried to choose a dataset appropriate to the characteristics of our problem: in particular, with known dates of consumption of content, and with implicit indicators for computing implicit feedback. The BARB (Broadcasters' Audience Research Board) [1] dataset used in these experiments seemed relevant; it consists of 6 months of viewing data, from January to June 2005, for approximately 3000 households (5000 viewers) in UK. Such data is usually used to provide minute-by-minute estimates of the number of people watching television overall UK and is therefore very accurate. Users belong to panels of television owning households representing the viewing behaviour UK households. Each viewing session indeed corresponds to a period where at least one person from the household was actually watching television. And if multiple users were watching television together, all of their identities are known.

This viewing information was combined with program schedule information for the same period to produce viewing sessions which indicate how much of a given program a given individual has watched. For performance reasons, experiments were conducted with a subset of 30 "typical" users (i.e. with no abnormal viewing patterns) selected for their tendency to watch TV alone (to limit the influence of others in their program choices). The idea is that these users can potentially provide more accurate implicit preferences. Each user had between 2000 and 3000 viewing session. Each user's viewing session in the dataset describes: the identifier of the TV programme; the percentage watched; the percentage of the beginning and of the end that has been missed; its title; its duration; its date and time and a list of weighted metadata describing the programme, obtained computing the TF-IDF score [11] which is used to discriminate between different metadata for the same content.

We performed a set of test to validate the behaviour the user profile should have in the following situations:

- *Gradual preference update*: overall change of preferences' weights based on content consumption (increase and decrease of weights using the formula);
- *Gradual forgetting*: usage of decay factor to remove concepts, which are not of user interest anymore;
- Usage of insertion threshold, and decay factor (based on persistence variations) to support *shift of interests*;
- Usage of insertion and removal thresholds and decay factor to support *periodic interests*.

The results we obtained are illustrated for a few specific users (randomly chosen) in the next subsections.

4.2 Gradual preference update

Fig. 3 and 4 represent the evolution of the number of occurrences of concept `games` for user 54 and its associated concept persistence and the resulting weight of the preference for `games` in the user profile.

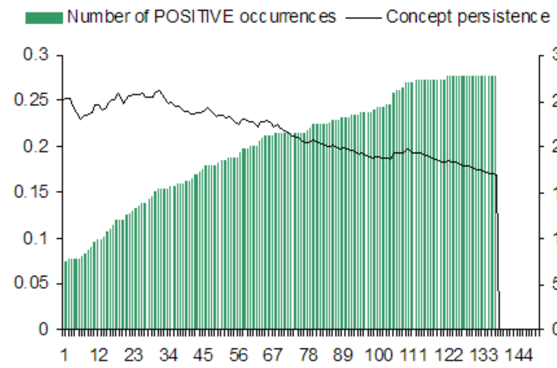

Figure 3. Gradual preference update: occurrences and persistence of concept `games` for user 54

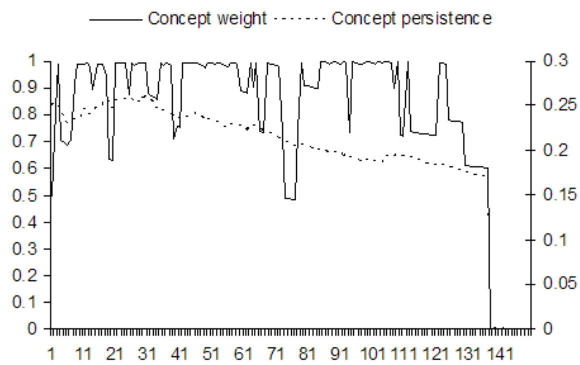

Figure 4. Gradual preference update: persistence and weight of concept `games` for user 54

Based on the user's implicit feedback on consumed content, the weight of the concept changes over time, we can observe two important behaviours in this example:

- When the concept persistence curve is stable or decrease slowly, the concept weight vary depending on the positive or negative user feedback;
- At update phase 133, the preference disappears from the user profile, due to a sufficient decrease of the concept persistence: when the persistence curve is decreasing, an additional decay factor is added to the update formula, making the concept weight decrease quicker. In addition, we introduced in the concept stack a removal threshold (which value in this example is 0.171) to remove from the preferences a concept which persistence is too low.

4.3 Gradual forgetting of preferences

Fig. 5 represents the evolution of the persistence of concept `cookery` for user 54 and the resulting weight of `cookery` in the user profile.

In this example we observe two main user behaviours:

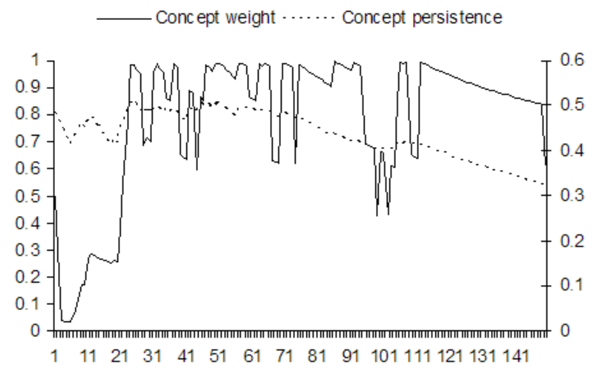

Figure 5. Gradual forgetting: persistence and weight of concept `cookery` for user 54

- The user is interested during some time in the concept `cookery`. Based on positive or negative user feedback the weight of the concept varied, as in our first example;
- After update period 111, the number of positive occurrences of the concept `cookery` is stable and the persistence begins to decrease. If the user does not provide new negative feedback on this concept, the weight of the preference will stay stable (according to update formula (2)). The introduction of a decay factor (based on the variation of persistence) allows to gradually forget the concept (decrease of the preference weight, except local perturbations due to negative implicit feedback).

4.4 Shift of interests

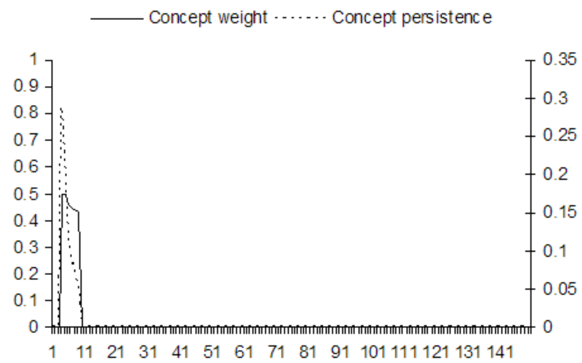

Figure 6. Shift of interest: persistence and weight of concept `games` for user 121

Fig. 6 represents the evolution of the persistence of concept `games` for user 121 and the resulting weight of the preference for `games` in the user profile.

This graphs show a sudden interest for the `games` concept, but this interest is not confirmed in time and the persistence decreases very fast. In that case using both the decay factor and the threshold for removal, the concept is removed quickly from the user profile, which avoid affecting the recommendations.

4.5 Periodic interests

The phenomenon of periodic interest happens often for event-related interests - typically sport events, where e.g. the interest of a user in football is correlated to the frequency of championship. In the case of periodic interest, the decay factor and the insertion/removal threshold help to deal with situations where the user consumed a concept during a certain amount of time and forgot about it during a big period and then returns back to this interest. If the personalised decay factor is not applied, only negative feedbacks decrease the preference weight in the update formula. If the user is stopping during a period of time the consumption of a concept, the preference weight (as exposed in gradual forgetting example) is staying stable.

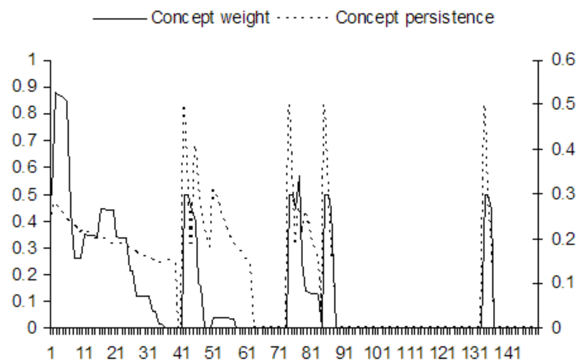

Figure 7. Periodic interest: persistence and weight of concept `cartoons` for user 496

Fig. 7 shows the evolution of persistence of concept `cartoons` for user 496 and its resulting weight in the user profile. In this particular case of periodic interest, the use of the decay factor helps the interest decrease when the consumption stops (update periods 71-91). Besides, if there is no mean to remove from the preferences a concept that is not consumed during an amount of time, the preference will decrease in time, and once the user gets back to its interest and consuming related content, then it takes some time for the learning mechanism to increase the preference weight, so that the recommender can finally propose it.

The use of an insertion/removal threshold in that particular case of periodic interest helped to remove quickly a preference that is not any more included in the consumed content, and reintroduced it with a medium weight of 0.5, when the consumption is back to a good level (reintroduction of concept `games` at update period 131). If the decay factor and the removal threshold would have not removed the concept from the preference and the concept stack, it would have taken much more time to acquire a sufficient level of weight for the preference to be considered in recommendation, because the persistence of the concept is calculated according to the first appearance in the stack, meaning that the persistence would have been too low to reconsider the concept as a new preference.

5 CONCLUSION

Reliability of user models is a key for the success of personalised systems. This implies to build models that are able to react in a flexible way to very different user behaviours, not limited to the inference of long-term interests, but also more complex content consumption schemes.

Thus, we presented in this paper a unified empirical model for transient and persistent user preferences, more specifically a mechanism to take care (acquire and forget) of punctual preferences based on a personalised decay factor which is applied in lifelong user profile update methods. This approach has been tested on a real dataset, which tends to validate our thoughts. Although we presented in this paper a very limited set of use-cases, we observed similar behaviours on all the users data we tested. However, we only obtained qualitative results. The evaluation of adaptative systems is known to be a difficult task [13]. In this initial experiments, we used a data-driven approach, but we were confronted with a lack of adequate evaluation metrics: indeed, common metrics such as learning rate, recommender accuracy or precision are not really appropriate in this case. We are also planning to incorporate this proof-of-concept into a broader prototype that will deliver multimedia personalised news to the users and which will provide means to evaluate more in depth our approach through a user-study.

ACKNOWLEDGEMENTS

The research leading to this document has received funding from the European Community's Sixth Framework Programme (IST-FP6-027685 - MESH). However, it reflects only the authors' views, and the European Community is not liable for any use that may be made of the information contained therein.

REFERENCES

[1] BARB - Broadcasters' Audience Research Board. http://www.barb.co.uk.
[2] D. Billsus and M. J. Pazzani, 'A Hybrid User Model for News Classification', in *7th International Conference on User Modeling*, pp. 99–108, (1999).
[3] E. H. Chi, 'Transient user profiling', in *Workshop on User Profiling (CHI2004)*, (2004).
[4] S. Gauch, J. Chaffee, and A. Pretschner, 'Ontology-Based Personalised Search and Browsing', *Web Intelligence and Agent Systems*, 219–234, (2003).
[5] S. Gauch and J. Trajkova, 'Improving Ontology-Based User Profiles', in *RIAO (Computer assisted information retrieval)*, pp. 380–389, (2004).
[6] H. R. Kim and P. K. Chan, 'Learning Implicit User Interest Hierarchy for Context in Personalisation', in *8th International Conference on Intelligent User Interfaces*, pp. 101–108, (2003).
[7] I. Koychev and R. Lothian, 'Tracking Drifting Concepts by Time Window Optimisation', in *AI-2005, the 25th SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence*, pp. 46–59. Springer, (2005).
[8] I. Koychev and I. Schwab, 'Adaptation to Drifting User's Interests', in *Workshop on Machine Learning in New Information Age (MLNet 2000) at the 11th European Conference on Machine Learning (ECML 2000)*, pp. 39–45, Barcelona, Spain, (2000).
[9] J. Lu and J. Callan, 'User Modeling for Full-Text Federated Search in Peer-to-Peer Networks', in *SIGIR*, Seattle, Washington, USA, (2006).
[10] M. Papadogiorgaki, V. Papastathis, E. Nidelkou, S. Waddington, B. Bratu, M. Ribière, and Y. Kompatsiaris, 'Distributed User Modeling for Personalized News Delivery in Mobile Devices', in *SMAP 2007, 2nd international workshop on Semantic Media Adaptation and Personalization*, London, United Kingdom, (2007).
[11] G. Salton and M. J. McGill, *Introduction to modern information retrieval*, McGraw-Hill, 1983.
[12] A. Tsymbal. The problem of concept drift: definitions and related work. Technical report TCD-CS-2004-15, Department of Computer Science, Trinity College Dublin, Ireland, 2004.
[13] S. Weibelzahl, 'Problems and Pitfalls in Evaluating Adaptive Systems', in *Fourth Workshop on the Evaluation of Adaptive Systems, held in conjunction with the 10th International Conference on User Modeling (UM'05)*, pp. 57–66, Edinburgh, United Kingdom, (2005).