Analysis of Tag-Based Recommendation Performance for a Semantic Wiki

Frederico Durao and Peter Dolog

IWIS — Intelligent Web and Information Systems,
Aalborg University, Computer Science Department
Selma Lagerlöfs Vej 300, DK-9220 Aalborg-East, Denmark
{fred,dolog}@cs.aau.dk

Abstract. Recommendations play a very important role for revealing related topics addressed in the wikis beyond the currently viewed page. In this paper, we extend KiWi, a semantic wiki with three different recommendation approaches. The first approach is implemented as a traditional tag-based retrieval, the second takes into account external factors such as tag popularity, tag representativeness and the affinity between user and tag and the third approach recommends pages in grouped by tag. The experiment evaluates the wiki performance in different scenarios regarding the amount of pages, tags and users. The results provide insights for the efficient widget allocation and performance management.

Key words: wiki, recommendation, tags, performance, adaptation

1 Introduction

Wiki is a collaborative knowledge space that can be edited by anybody who is granted permission [11]. Due to its simple usage, the wiki adoption is less about learning new technology and more about changing habits. A part from the complexity of existing Web solutions, wiki has instituted a new and democratic way of usage with simple text syntax for creating pages and cross links between internal pages on the fly. Although wikis provide an easy way for editing content pages, user interaction is still on "one way" i.e. users have to look at wiki pages to find interesting content to them. In the other direction, wikis could notify users about what they hide behind the currently viewed page. In this sense, recommendations can be utilized to lead users to unknown pages and reveal related topics addressed in the KiWi (KiWi - Knowledge in a Wiki). Furthermore, the recommendations can be tailored to user tastes and adaptively configured depending on system needs such as performance. The KiWi system addressed in this paper is a social semantic wiki in which individuals work collaboratively by editing content items and sharing knowledge. In serves as a platform for implementing and integrating many different kinds of social software services by allowing users to connect content in new ways that go beyond the level of the user interface, e.g. through semantic annotation [14].

In this work, we extend the KiWi system with three tag-based recommender approaches, which suggest links to wiki pages based on the similarity of their tags. The first approach recommends pages which share tags, the second approach takes into account external factors such as tag popularity, tag representativeness and the affinity between user and tag and finally the third approach groups recommendations by tags. The performance of the approaches is compared in different scenarios, which varies in terms of amount of pages, users and tags. The outcome from this analysis provides insights for widget allocation (where the recommendations are placed) and subsequent performance optimization. Our development is placed at KiWi [14], a semantic wiki for knowledge management built on previous experience in areas such as semantic web [5], semantic wiki [13] and personalization [6].

The paper is organized as follows: In Section 2 we discuss related work. In Section 3, a motivation scenario is presented. Section 4 introduces the recommendation approaches. Section 5 presents the experimental evaluation and results. A discussion about the results from the previous section is presented in Section 6 and finally in Section 7, we conclude the study and also point to future works.

2 Related Work

A number of semantic wiki applications have been explored over the last years and most of them utilize annotations to contextualize the content presentation and improve the navigation throughout all existing pages. *SemperWiki* is a semantic personal wiki developed for the Gnome desktop in which users can edit and annotate pages semantically [12]. In order to navigate through the wiki pages, users have to query pages containing certain annotation statements. The retrieval brings a list of links to the existing pages in the system. In addition, SemperWiki provides a history navigation section that allows users to go back and forth in their navigation history. The navigation support provided by SemperWiki is enhanced by a search and retrieval mechanism. We observe that the discovery of new pages in SemperWiki depends more on user's curiosity whereas the recommendations in KiWi are always displayed without imposing any additional work on the users. The history navigation however can be considered as a positive feature in SemperWiki because it is very practical for rapid navigation between visited pages. This feature can be adopted in KiWi to generate a new sort of recommendation triggered by history log of visited pages. Already *IkeWiki* [13] as a predecessor of KiWi provided a "references box" containing related pages triggered by annotation in the wiki pages. Similarly, *Semantic MediaWiki* [10] suggests related pages which share similar instances. Recommendations in KiWi are less formal than Semantic MediaWiki and IkeWiki since they are triggered by tags which are not bounded to any ontology. On the other hand, the flexibility of tags allows users to spill their personal feelings to a wiki page so that this generates more personalized recommendations.

Equally to KiWi, *OntoWiki* interface is surrounded by widgets that provide meta-information from semantic annotations and navigation support. Although

OntoWiki does not use tags for processing recommendations, it contains a particular widget for related pages categorized by the *Most Popular* and *Most Active* [1]. *HyperDEWiKi* is a semantic wiki intended to support domain ontology evolution [15]. It allows the user to define specific pages for instances of formally described types. In this sense, users can create a dynamic page that is better suited to support his tasks. We observe that the end view of KiWi and HyperDEWiKi can be fully customizable and render various set of information in different places and layouts. Besides the common wiki style editing with annotations, both systems provide personalized features tailored to user's tastes.

In general, the semantic wiki applications analyzed utilize annotations in the pages for navigation, rendering and search purposes. However, we observe that navigation is still centered on dynamically generated lists of related content with no or few personalized information. Personalization will drive the system features in accordance with individual tastes and preferences [7]. In addition, it is observed that adaptive techniques can be more explored in order to support wikis to present their content more intelligently [4]. Following these premises, we introduce three personalized tag-based recommendations and evaluate their allocation in KiWi interface aiming at performance optimization.

3 Motivating Scenario

In general, tags are assigned to Web resources in order to conceptualize, categorize and organize them in a way that users can be reminded later about the tagged content [9]. Invariably, tags represent some sort of affinity between user and the page that is being assigned. Users label pages freely and subjectively, based on their sense of values. This information provides useful hints about what a user thinks about the pages [8]. In this sense, we utilize tags to compute similarity between wiki pages and generate personalized recommendations without imposing any extra work on the users. We credit personalized recommendation as an important feature for supporting the main activity in wiki systems. For instance, when users are reading or editing a wiki page, recommendations of similar pages can be processed simultaneously and exhibited so that users can navigate through wiki pages following the topic that is being addressed. According to [3], recommendations have a significant importance because they expose alternative ways to the users fulfill their goals. In this sense, we provide recommendations in KiWi whereby users can follow links related to their interests, which assist them to achieve their tasks or bring further information for what they are looking for.

Figure 1 shows the current development in KiWi system in which tags assigned to the currently viewed page are located on the bottom widget on the left side (1). The recommendation widgets are highlighted on the right side: the widget number (2) contains the *standard recommendation*, the widget number (3) contains the *multifactor recommendation* and the widget number (4) contains the *recommendations grouped by tags*. The recommendations expose a variety of options for a user to visit just on a single click. If this activity is designed well,

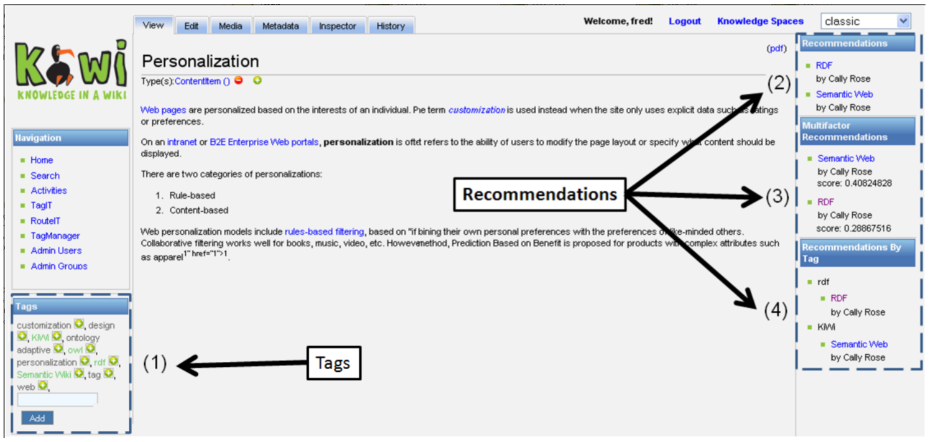

Fig. 1. Recommendations in KiWi

then the choice is easy, and the user keeps interacting with the system by visiting the related pages or adding new content. Although each recommendation approach has its own particularities (See in Section 4), very hardly the three solutions will run in parallel in real life scenario because they occupy much space in the user's interface and occasionally issue the same information (at the same time). In addition of being useless from the usability perspective, to have all three widgets running together compromise the system performance at all. As known, wiki is a collaborative space utilized for multiple interactions and any performance concern is always advisable. Based on these premises, the performance analysis is undertaken in order to find out widget combinations so that the overall performance is optimized and users take advantage of better widget arrangements.

3.1 Tags as semantic annotations for personalization

In this work we are mapping semantics of user activities based on tagging activity. Using uncontrolled tags, users are able to annotate pages without any restriction constrained by ontology vocabularies. In this sense, users are free to express their feelings about the page as they like on any purpose. The outcome of this tagging activity is a relation between user and a page through a tag property, as seen in the Figure 2.

Fig. 2. Relation between user and a page through a tag property

From the tag-based relationships, we are deriving personalized recommendations by computing similarities between tags, however, in later stages, other relevant information can be derived and utilized to annotate the wiki pages using RDF properties such as *ont:mostFrequentTag*, *ont:userMostInterested* and *ont:mostSimilarPage*. These properties would create a semantic network between content items in KiWi, and also could be utilized for other personalization goals such as group formation, semantic search and creation of link structures.

4 The Recommendation Approaches

This section depicts the standard, multifactor and recommendation grouped by tags addressed in this work.

Standard Tag-based Recommendation. In this approach, all pages that share tags with the currently viewed page are recommended. In this standard approach no further similarity processing is carried out therefore the list of recommendation is not ranked. The advantage of this approach is the performance since the recommendations relies simply on a data retrieval task. On the other hand, a single tag shared by pages may not be sufficient means to determine a similarity between pages. This approach however cannot be discarded without analyzing its applicability in the different possible KiWi scenarios. Figure 3 shows standard

Fig. 3. Widget with Standard Recommendations

recommendations in KiWi with their respective authors.

Multifactor Recommendation. The multifactor recommendation approach computes similarity between pages considering multiple factors. The recommendations rely on calculus of *cosine similarity*, *tag popularity*, *tag representativeness* and *affinity user-tag*. We utilize a cosine similarity measure between tag vectors to calculate basic similarity of the pages. We measure tag popularity as a count of occurrences of a certain tag in total number of wiki pages. The term frequency measure is used to compute tag representativeness for a certain wiki page. The tag affinity between a user and a tag is calculated as a count of how many times the user utilized the tag at different web pages. We propose a formulae which consider all these factors in a normalized way and gives a ranking of pages for particular user.

We define a page score as:
$Ps = \sum_{i=1}^{n} weight(Tag_i) + \sum_{i=1}^{n} representativness(Tag_i)$, where n is the total number of existing tags in the repository.

We define the tag user affinity as:
$Affinity_{(u,t)} = card\{p \in Pages \mid (u,t,p) \in P, P \subseteq U \times T \times P\} / card\{t \in T \mid (t,u) \in P_u, P_u \subseteq U \times T\}$, where t is a particular tag, u particular user, U is a set of users, P set of pages and T set of tags.

Finally, similarity is computed as:
$Similarity_{(P_i,P_{ii})} = [Ps_{P_i} + Ps_{P_{ii}} * cosine\_similarity(P_i, P_{ii})] * Affinity_{(u,t)}$.

Informally, each one of the factors in the above formulas is calculated as follows:

i. **Cosine Similarity** — Our tag similarity is a variant on the classical cosine similarity from the text mining and information retrieval [2] whereby two items are thought of as two vectors in the m dimensional user-space. The similarity between them is measured by computing the cosine of the angle between these two vectors.
ii. **Tag Popularity** — Also called *tag weight*, is calculated as a count of occurrences of one tag per total of resources available. We rely on the fact that the most popular tags are like anchors to the most confident resources. As a consequence, it decreases the chance of dissatisfaction by the receivers of the recommendations.
iii. **Tag Representativeness** — It measures the relation between the tag and the resource it belongs to. It is believed that the most frequently occurring tags in the document can better represent the document. The *tag representativeness* is measured by the *term frequency*, a broad metric also used by the Information Retrieval community [2].
iv. **Tag Weight** — it is calculated as a number of occurrences of a tag divided by the overall number of tags in a repository.
v. **Affinity between user and tag** - It measures how often a tag is used by a user. It is believed that the most frequent tags of a particular user can reveal his/her interests. This information is regarded as valuable information for personalization means. During the comparison of two resources, the similarity is boosted if one of the resources contains top tags of the author from the other resources around.

Figure 4 shows the same recommendations as Figure 3 however sorted differently due to the quality factors calculus. In terms of performance, the multifactor approach is worse than the standard one however the ranked list provides credibility to the recommendations. Pages ranked higher are personalized to user's tastes and closer to the content discussed in the currently viewed page. Although the recommendations are more effective than the standard approach, its applicability also depends on further performance analysis in different KiWi scenarios.

Fig. 4. Multifactor Recommendations

Recommendations Grouped by Tags. In this approach, the recommendations are grouped by the tags which are assigned to the currently viewed page. Similarly to the standard approach, no further similarity processing is undertaken and the list of recommendation is not ranked. On the other hand, the user can go directly to the recommended wiki page just following the tag he/she is interested. The tag-based distribution explicitly provides a justification why the recommendations were generated and assist users to find related specific wiki pages.

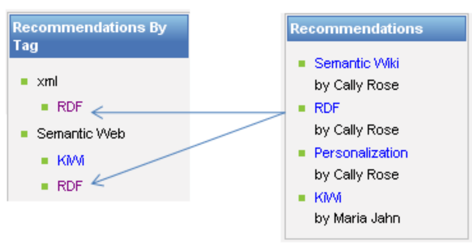

Fig. 5. Widgets with recommendations grouped by tags

The disadvantage however is the possibility of existing duplicated recommendations since two different pages can share two distinct tags as well. Figure 5 shows two tags *xml* and *Semantic Web* with their respective linking recommendations. The duplication problem is outlined since both tags recommend a link to the *RDF* wiki page. The performance analysis therefore will answer whether this duplication problem affects the performance to the point of discarding the applicability of this approach.

Scenarios	Pages	Tags	Users
1º	20	225	5
2º	50	500	10
3º	100	700	15

Table 1. KiWi Scenarios assembled for the experiment

5 Experimental Evaluation

The experimental evaluation assessed the performance of the recommender approaches by simulating a mix of scenarios regarding the amount of pages, users and tags. In addition, we discussed which widgets of recommendations should be displayed or suppressed in accordance with the performance findings. Although having in mind that standard approach is theoretically the most advisable approach in terms of performance, we evaluated when and in which conditions the other approaches become suitable and necessarily required. The goal therefore is to propose insights for widget adaptation aligned with running scenarios without compromising the system performance.

The proposed scenarios were created aiming at simulating realistic usage of KiWi. Nevertheless there is no "pattern" or "standard" about wiki activity. This is more about politics from where the system is deployed, maturity of whom is using it and time of activity. Although understanding that building wiki scenarios is a little subjective, we proposed three growing scenarios that are likely envisioned in KiWi life cycle as described in Table 1.

The variables addressed by each scenario are:

- *Amount of Pages* – each page has a set of tags that are compared for processing the recommendations. Therefore the more pages exist, the more time will be spent to calculate the similarity between the pages.
- *Amount of Tags* – the similarity of the pages is given by their tags. The whole set of tags of each page must be compared to verify which ones are similar. In particular, the amount of tags of an user impacts directly in the time for the computation of the affinity between user and tag.
- *Amount of Users* – the more users KiWi contains, the more are the chances of tagging activity. As a consequence, more time will be necessary to process the personalized recommendations.

These variables were chosen justified because they are row material for calculating the recommendations. This process is time consuming and invariably affects the system performance however it does not mean that other factors such as page size should not be considered. Finally, the scope of this work only comprises these three variables.

5.1 Methodology

Initially, the KiWi system was populated with pages and tags using a random generator to produce sufficient amount of content. The users were created manually since they were only 15 at most. The content of the pages and tags were extracted from Web sites on the Internet and local documents. Similarly, we utilized a random generator to assign tags to wiki pages tagged by a particular user. In this study, it is not important to speculate about the random function we have utilized as we look at overall performance of the system and not the method for distributing the content. We needed just to generate adequate number of satisfactorily different pages and sufficiently different assignment of tags to them. To each scenario created, we collected the response time necessary to load the recommendations. We repeated this procedure for 10 times in order to have a more real and democratic results. In KiWi, the recommendations are processed every time a page is called, then we tested the performance after the user login, accessing a page from page link and from the own recommendations. The tests ran in machine equipped with processor Intel (R) Core (TM) 2 Duo CPU T7500 @2.20Ghz.

Expected Results Our assumption is that at some point due to the high amount of recommendations, the standard approach become ineffective although the performance continues better than the other approaches. The quality achieved with the multifactor recommendation will be needed even though the performance is decreased. Moreover, we believe the group recommendation is always useful due to its facility of identifying similar pages by tags however its performance may discourage its adoption due to the high number of duplicated pages.

5.2 Results from the First Scenario

The first scenario was setup with 20 pages, 5 users and 225 tags. Figure 6 shows 10 time stamps collected from KiWi system to calculate the recommendations for the three approaches. The average column from Figure 6 shows that standard recommendations were computed in 28 ms; multifactor recommendation in 77 ms and recommendations grouped by tags in 29.1 ms.

Approach	10 Timeshots(ms)										Average (ms)
CLASSICAL	15	15	15	31	47	47	15	31	32	32	28
MULTIFACTOR	15	16	109	93	172	109	6	110	62	78	77
GROUPED	47	10	16	78	16	16	15	15	63	15	29.1

Fig. 6. Results from the First Scenario

The standard and grouped approaches had better performance than the multifactor approach. While standard and grouped approaches achieved rates quite close to one another (about 28ms), the multifactor approach spent 175% more time than both approaches to calculate the recommendations. As already known, the multifactor recommendation generates a ranked list of recommendations more personalized and reliable. The point to be assessed therefore is whether the ranking compensates this time necessary for ranking the recommendations. In this particular case where only 20 wiki pages are considered, the multifactor approach may be ignored depending on the visibility of the recommended pages to the users. If the widgets of standard and grouped recommendations are able to expose the whole recommendations with easy access, they will be continuously exposed and hardly forgotten. In this sense no ranked recommendation is necessary.

5.3 Results from the Second Scenario

The second scenario was setup with 50 pages, 10 users and 500 tags. Figure shows 10 time stamps collected from Kiwi system to calculate the recommendations for the three approaches. The average column from Figure 7 shows that standard recommendations were computed in 97 ms; multifactor recommendations in 121 ms and grouped recommendations in 45.3 ms.

Approach	10 Timeshots(ms)										Average (ms)
CLASSICAL	94	78	78	47	63	78	125	110	125	172	97
MULTIFACTOR	103	92	112	78	62	93	124	212	166	168	121
GROUPED	93	15	15	16	54	78	93	31	47	11	45.3

Fig. 7. Results from the Second Scenario

The grouped approach achieved the best performance followed by standard and finally by the multifactor approach. The multifactor approach lasted approximately 25% more than the standard approach to generate the recommendations. Comparing to the first scenario, it is observed a considerable approximation between standard and multifactor approach (from 125% to 25%). The issue to be discussed therefore is whether the standard recommendation approach is still desired due to the low performance variation between the standard and multifactor approach. In this scenario, 50 wiki pages are addressed and the recommendation widgets tend to enlarge with the increase of recommendations. In this case, therefore, a ranking approach becomes useful since the most similar pages are placed at the top of the list of recommendations. In addition, very hardly the recommendation widget will provide an ample visibility of whole set of recommendations. In this sense, the 24 additional milliseconds to generate the multifactor recommendations compensate the probability of having ineffective recommendations. The high performance of the group based approach can be justified by the low amount of duplicate recommendation. Furthermore, it is important to observe that although the amount of pages of this scenario is the double comparing to the first one, the current performance was decreased of only 55.6% (or 16.2 ms) from the first measure. This approach therefore becomes a strong candidate to be utilized in this particular scenario even in combination with the multifactor recommendations.

5.4 Results from the Third Scenario

The third scenario was setup with 100 pages, 15 users and 700 tags. Figure 8 shows 10 times stamps collected from KiWi system to calculate the recommendations for the three approaches. The average column from Figure 8 shows that the standard recommendations were computed in 98.1 ms; multifactor recommendation in 145 ms and grouped recommendations in 53.1 ms.

Approach	10 Timeshots(ms)										Average (ms)
CLASSICAL	93	87	94	63	141	93	94	110	125	81	98.1
MULTIFACTOR	109	113	141	163	188	141	140	187	156	112	145
GROUPED	32	46	62	63	52	43	52	31	78	72	53.1

Fig. 8. Results from the Third Scenario

Similarly to the second scenario, the grouped approach achieved the best performance followed by standard and finally by multifactor recommendation. In this turn the standard approach was approximately 50% faster than the multifactor approach. This is a considerable difference that stands face-to-face two goals: performance and effectiveness. On one side, KiWi process lots of recommendations very quickly however unsorted, on the other hand the same recommendations last at least 50% more but they are ranked. A qualitative experiment in which the users show their satisfaction in terms of performance and effectiveness would indicate in which direction to go. A possible solution however is to combine the multifactor approach with the fastest group recommendation, which attenuates the overall loss of performance for this scenario. The group approach performance decreased only 17% from the second scenario (from 45.3 ms to 53.1 ms) and still provide a special distribution of the recommendations.

5.5 Overall Results

The performance of the approaches in each scenario analyzed are presented in the Figure 9.

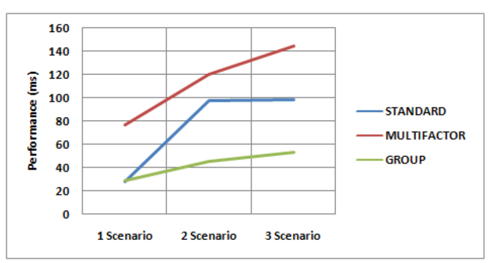

Fig. 9. Evolution Graph

In the first scenario, the multifactor approach was considerably more expensive (in terms of performance) than the others. This significant cost for generating recommendations discourages its adoption. In the second scenario, the standard approach reduces significantly its performance, which encourages the use of the multifactor approach. From the first to the second scenario, the group approach maintains a satisfactory level of performance. In the third scenario, standard and grouped approaches keep their performance approximately equal to the previous scenario, whereas multifactor approach presents a significant loss of performance.

6 Discussion

The performance outcomes from the scenarios analyzed allow us to suggest intelligent widget allocations without compromising the system performance. For the first scenario, the standard and group approaches are the most advisable to run together spending in total about 57.1 ms to calculate the recommendations. For the second scenario, we suggest the multifactor and group approaches running in parallel spending in total 166.3 ms and for the third scenario, due to the need of quality, again the multifactor and group approaches are the most advisable spending together about 198.1 ms to calculate the recommendations. Figure 10 shows the performance with the suggested combination to each scenario analyzed.

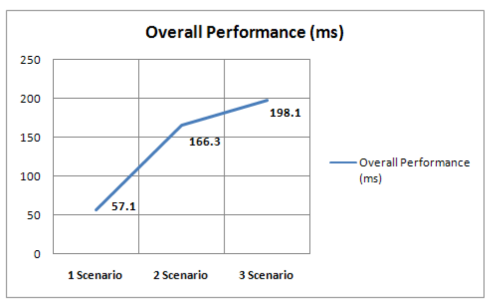

Fig. 10. Suggested Performance Graph

According to Figure 10, the performance for the first scenario is 57.1 ms, which is significantly low since two approaches of recommendation are being computed. In the second scenario, it is observed an expressive loss of performance (from 57.1ms to 166.3 ms) mainly due to the utilization of the multifactor approach. In the third scenario however the loss of performance was 31.8 ms, which is much less than the loss of performance from the first to the second scenario (109.2 ms). In fact, the multifactor approach utilized in both second and third scenarios reduced the performance, however, we assure that the best recommendations will be placed in the top of the list of recommendations. In general, the overall result obtained shows the tendency of the performance whenever the amount of user, page and tag grows. Although the more scenarios are advisable for confirming this tendency, this preliminary outcome can be employed for predicting the system performance in emergent scenarios.

Other advantage from the results obtained, besides the widget allocation, is that they show a tendency of the performance whenever the amount of user, page and tag grows. The achieved numbers can be utilized for predicting the system performance in emergent scenarios. On the other hand, the amount of the scenarios analyzed is still low to confirm this tendency. More experiments, with a bigger setup and with multiple users using the system at the same time would provide a more realistic feedback about the performance.

7 Conclusion and Future Works

This paper analyzes the performance of three tag-based recommender approaches for a semantic wiki. Three different scenarios were assessed varying in terms of number of pages, amount of tags and users. To each scenario assembled, it was analyzed which recommendation approaches could be more appropriate taking account the system performance and user needs. The results showed that the combination between standard and group approach is feasible for scenarios up to 20 pages, which are constantly accessed. For scenarios with 50 and 100 pages with more than 10 users, the multifactor and group approaches are more advisable in spite of being more expensive in terms of performance. The grouped recommendation approach is always adequate since it provides justification for recommendation and visual support for navigating among the recommendations.

As future works, first, pre computation of some factors in multifactor recommendation will be studied to further increase performance of recommendation computing. Furthermore, semantic relatedness between tags must be considered since current recommendations in KiWi only consider the tag syntax to identify similarities. The next step therefore is to combine the tag algorithms with some reasoning on the annotations to provide more efficient recommendations. Another direction is to annotate tags with their role (or purpose) in order to formalize the relationship between tags and pages. The semantic recommendation will likely be more efficient by capturing precise needs of the users expressed by the annotations.

8 Acknowledgment

The research leading to these results is part of the project "KiWi - Knowledge in a Wiki" and has received funding from the European Communitys Seventh Framework Programme (FP7/2007-2013) under grant agreement No. 211932.

References

1. S. Auer, S. Dietzold, and T. Riechert. OntoWiki - A Tool for Social, Semantic Collaboration. In I. F. Cruz, S. Decker, D. Allemang, C. Preist, D. Schwabe,

P. Mika, M. Uschold, and L. Aroyo, editors, *The Semantic Web - ISWC 2006, 5th International Semantic*, volume 4273 of *Lecture Notes in Computer Science*, pages 736–749. Springer, 2006.
2. R. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison Wesley, May 1999.
3. K. Bischoff, C. S. Firan, W. Nejdl, and R. Paiu. Can all tags be used for search? In J. G. S. et al., editor, *Proc. of the 17th ACM Conference on Information and Knowledge Management, CIKM 2008*, pages 193–202, Napa Valley, California, USA, Oct. 2008. ACM.
4. P. Dolog. Designing adaptive web applications. In V. Geffert, J. Karhumäki, A. Bertoni, B. Preneel, P. Návrat, and M. Bieliková, editors, *Proc. of Sofsem 2008: The 34th Intl. Conference on Current Trends in Theory and Practice of Computer Science*, volume 4910 of *LNCS*, pages 23–33, Novy Smokovec, Slovakia, Jan. 2008. Springer Verlag.
5. P. Dolog, B. Simon, T. Klobucar, and W. Nejdl. Personalizing access to learning networks. *ACM Transactions on Internet Technologies*, 8(2), Feb. 2008.
6. P. Dolog, H. Stuckenschmidt, H. Wache, and J. Diedrich. Relaxing rdf queries based on user and domain preferences. *Journal of Intelligent Information Systems*, 2009. published online.
7. F. A. Durao, P. Dolog, and K. Jahn. State of the art: Personalization. Technical report, 2008. KIWI project FP7, Project Number: ICT-2007.4.2-211932 Document number: ICT211932/SFRG/D2.7/D/PU/b1.
8. H. Halpin, V. Robu, and H. Shepherd. The complex dynamics of collaborative tagging. In C. L. Williamson, M. E. Zurko, P. F. Patel-Schneider, and P. J. Shenoy, editors, *Proc. of the 16th Intl. Conference on World Wide Web, WWW 2007*, pages 211–220, Banff, Alberta, Canada, May 2007.
9. R. Jschke, L. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme. Tag recommendations in folksonomies. In A. Hinneburg, editor, *Workshop Proceedings of Lernen - Wissensentdeckung - Adaptivitt (LWA 2007)*, pages 13–20. Martin-Luther-Universitt Halle-Wittenberg, sep 2007.
10. K. Lassleben, D. Vrandecic, and M. Vlkel. Semantic mediawiki. In *The Semantic Web - ISWC 2006*, pages 935–942, 2006.
11. A. Majchrzak, C. Wagner, and D. Yates. Corporate wiki users: results of a survey. In *WikiSym '06: Proceedings of the 2006 international symposium on Wikis*, pages 99–104, New York, NY, USA, 2006. ACM.
12. E. Oren. SemperWiki: a semantic personal Wiki. In S. Decker, J. Park, D. Quan, and L. Sauermann, editors, *Proceedings of the 1st Workshop on The Semantic Desktop, 4th International Semantic Web Conference*, Galway, Ireland, November 2005.
13. S. Schaffert. Ikewiki: A semantic wiki for collaborative knowledge management. In *WETICE '06: Proceedings of the 15th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises*, pages 388–396, Washington, DC, USA, 2006. IEEE Computer Society.
14. S. Schaffert, F. Bry, P. Dolog, J. Eder, S. Grünwald, J. Herwig, J. Holý, P.-A. Nielsen, and P. Smrž. The kiwi vision: Collaborative knowledge management, powered by the semantic web. Technical report, August 2008. FP7 ICT STREP KiWi project Deliverable D8.5.
15. D. Schwabe and M. R. da Silva. Unifying semantic wikis and semantic web applications. In C. Bizer and A. Joshi, editors, *International Semantic Web Conference*, volume 401 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2008.