

Wikis are easy to use hypertext systems that can be read and modified by its online users through their browsers. The most prominent example of a wiki is probably Wikipedia. Wikipedia has several thousand registered users and an unknown number of anonymous users from all over the world. The users edit existing articles or add completely new texts. Millions of articles in different languages were authored by volunteers in this way (Danowski and Voss, 2004).

Numerous companies, research institutes and private persons rely on wikis to manage their business-related knowledge (Danowski and Voss, 2004; Rohs, 2007). While most CSCW literature is focusing on analyzing user behavior and cooperation in Wikipedia (Antin and Cheshire, 2010; Kittur and Kraut, 2010), our work concentrates on work environments. The *Mobile Knowledge* (MoKnow¹) and *Context and Attention in Personalized Learning Environments* (CAPLE²) work groups at Fraunhofer FIT share a wiki — called the MOKNOWPEDIA — for documentation and knowledge management. Typical topics of articles are news, projects and persons, technical documentation and tutorials, trade fair exhibitions, seminars and workshops, publication planning, processes, interesting links, or IT infrastructure.

However, first experiences with MOKNOWPEDIA showed that participation and contribution was not as active as hoped. In order to improve knowledge exchange in the groups, the articles' quality and quantity needed to be improved. Although resilient numbers are hard to come by for work environments, MOKNOWPEDIA is surely not an exception (see also Hoisl et al. (2007)). Indeed, Kraut et al. (2010) find that success stories like that of Wikipedia are rare. Out of 6000 installations of the MediaWiki software, not even every second one has eight users or more.

A major problem is that users of public and corporate wikis have to make personal efforts and invest their own precious time in creating and improving content that is often far more useful to others than themselves: an article's author already has the knowledge he writes about. To him, the article is of less value because it provides only few benefit. Rational beings like humans will clearly prefer consuming content over producing it. But a wiki with few contents and poor articles fails to support knowledge management as there is nothing to consume. The wiki is not attractive, and the few contributors are further demotivated as they feel that nobody else is contributing. In a closed community, where the number of potential authors is limited and everyone is a specialist with specialist knowledge, a small percentage of volunteers is not enough. Instead, active participation of all is necessary.

We make the following contributions that help to increase contribution quantity (article count and size) and quality (readability and timeliness) in a corporate wiki:

- a concept that motivates personnel by means of reputation,
- an implementation of our concept that amplifies contributing behavior,
- an evaluation based on a four months field-test in productive use,
- a comparison of the three rewarding mechanisms integrated in our solution,
- and we inform the future design of corporate knowledge repositories.

¹ http://www.fit.fraunhofer.de/projects/mobiles-wissen_en.html

² http://www.fit.fraunhofer.de/services/mobile/caple_en.html

Motivation

This section explains the motivation behind our concept. Starting from a view of the wiki's deployment history, we interviewed users for the problems they see with the wiki, and discuss theoretical backgrounds of the mentioned problems.

History of the MOKNOWPEDIA

A few years ago, MoKnow and CAPLE decided that a new form of managing knowledge was needed. In a dedicated discussion, it was decided to set up the MOKNOWPEDIA based on the MediaWiki³ software. One advantage that favored the MOKNOWPEDIA approach was that a free wiki software is cost-saving compared to other commercial software. Furthermore, there is no need for an explicit editorial control. The wiki structure can evolve freely and all kinds of articles are possible. As MediaWiki puts only few restrictions on users and what they write, it provides a work space that enables the vertical and horizontal exchange of knowledge between group leaders and equal colleagues. Additionally, entrance barriers to wiki use are low because users can easily contribute to collaboratively written articles, just by using their browsers. Finally, the success stories of Wikipedia or wikis used by open source software projects clearly voted for the wiki approach. In practice, however, it turned out that MOKNOWPEDIA had some troubles taking off:

Only a few members of the group actively participated in the wiki. Not everybody seemed equally motivated to contribute. There was a higher number of people willing to consume than those willing to produce. "You should write an article about this!" was often heard but rarely done. In the end, the total number of articles remained low; leading to a limited exchange of knowledge through MOKNOWPEDIA. Most information was still conveyed in the old ways like hallway conversations or explicit requests for documentation. Instead, we wanted to achieve that information was provided pro-actively through a central and well-known repository.

Our observations with MOKNOWPEDIA confirm the results in a similar study of knowledge sharing tools and processes by Reichling and Veith (2005).

An investigation into the users' problems

To qualitatively investigate into the reasons for low contribution, we interviewed all members of the work group. We wanted to know why they themselves or someone else was possibly not actively contributing to MOKNOWPEDIA, and what could be changed to improve this situation and motivate them to contribute. Our interviewees reported various problems. Saying that a good tool would be reward enough to use it, many named mostly shortcomings of the wiki as problems:

Every fifth interviewee answered that the wiki's syntax was too complex to use and learn. This stopped them from using it. Also, it was missing a clear structuring. Half of the interviewees mentioned that there should be someone who creates the

³ <http://www.mediawiki.org>

structure. Adding own contributions would be much easier then. Some interviewees stated that there were so many other group tools requiring their attention that they could not additionally contribute to MOKNOWPEDIA. Others said that they preferred to have a semantic wiki, instead, because it would help them to better structure the wiki. No opportunity should be left out to promote the wiki, too. A few persons were frustrated because they felt they were the only ones to write and tend to articles in a certain domain. Especially student employees were missing an instructor that regularly reminded them to contribute, or told them to write a certain article. Students were also afraid that they were not competent enough to make professional and quality contributions. Finally, most interviewees stated that they did not have enough time to write articles, although they would really like to do so.

Implications, target group and goals

Although there is truth in many of the users' complaints, we think that the main problem is that they are not motivated sufficiently to try to overcome the hurdles by themselves: The reason that they do not have the time — to learn the wiki syntax, to create a structure, have their attention drawn to other tools, have no one to tell them what to do, felt incompetent, and have other important things to do — is that they do not see enough value for themselves in contributing.

For rational, thinking beings, a clear disproportion exists: Writing articles costs them their own time, in the extreme case even making themselves replaceable, while others profit from easily accessible information. While the group as a whole would profit, every single person is tempted to free ride on the group's achievements. Reducing the extent of this disproportion is key to increasing contribution.

The target group in our study are self-organizing project teams at a research institute with flat hierarchies. Individuals have far-reaching freedom and autonomy in organizing their work. Explicit instruction or punishment are, at most, rare. Intrinsic motivation is valued very much. Staff members are mostly researchers, one third working students, and two executive managers. The vast majority has a computer science background with a few exceptions that still have good computer knowledge. In this environment, we design our solution to the following goals:

- G1: Increase content volume to preserve a broader spectrum of information.
- G2: Improve article quality to make information more easy to grasp.
- G3: Involve deniers to preserve the unique expert knowledge of everyone.
- G4: Keep the cost of operation low.
- G5: Do not force contribution and avoid punishment (self-organizing groups).
- G6: Do not cause destructive phenomena like rat races or bullying.

The reputation model

According to the Oxford Dictionary, *reputation* is what is generally said or believed about the abilities or qualities of somebody or something. A reputation system is a

software that determines a means of a user's reputation from his actions. The computed reputation scores can be used to predict future user behavior or be published. A reputation system is a core component in many modern web-based communities, where it serves to promote well-behaving and trust (Jøsang et al., 2005).

In our concept, users earn reputation points by serving their work group with valuable contributions to MOKNOWPEDIA. For example, writing high quality articles earns users a certain amount of reputation points. The users' reputation points are published resulting in appreciation by their peers. This gives users something in exchange for their efforts, and motivates them to systematically collect more points.

Consider the following scenario as a simplified example: Sonny creates a new article in MOKNOWPEDIA. Rico reviews the article with a rating of "average". For his review, Rico receives one point, while Sonny receives ten points for his average-quality article. Later, Gina edits Sonny's article, replacing about one third of the text. The article's ten points are now split proportionally between Sonny (seven points) and Gina (three points). In addition, Rico's review loses timeliness because it applies to a quite different, older revision of the article. Rico only retains 0.7 points for his old review. Gina adds a new review of "good" to the article, giving her an extra point for the review itself. When the voted quality of the article is determined from reviews, Gina's new review weighs more than Rico's old review. Therefore, the average rating of the article is almost "good". Because of its quality, fifteen points (instead of ten) are now distributed among contributors of the article. This leaves Sonny with ten points, Gina with five points plus one for her review, and Rico with 0.7 points. When Rico refreshes his review, he gets the full point again.

In the example, you see that there are two ways for users to collect reputation points: The first one is to provide a quality assessment of an article in form of a rating and a review comment. The quality assessments are used to democratically determine the quality of an article. A fresh review earns the user one point. However, this worth decreases as the review ages due to later changes to the article. The reviewer has to refresh his review to regain the full point. A reviewer cannot have more than one active review per article, but can revise his review at any time. And, to avoid that reviewers are influenced by fear of bad (or expectation of good) consequences for themselves (Elster, 1989), reviews are always submitted anonymously.

The second way to earn reputation points is by contributing to MOKNOWPEDIA articles. The amount of reputation points that are awarded for a contribution to an article depends on three criteria: quality of the article as determined by above democratic review process, importance of the article (e.g. size, page views), and ratio of the contributor's contribution to all contributors' contributions to the article.

While users contribute to articles in the wiki and review them, a reputation system collects information about reputation-relevant interactions between users and MOKNOWPEDIA in the background, gauging the social value of their interactions. The collected reputation data is visualized in three different ways, to make users themselves and other users aware of how much they have done for their group:

- **Levels** — similar to Internet forums, every user is assigned to a level in a hierarchy. Users have to collect a certain amount of points before being promoted

to a new level. This mechanism addresses the users' drive for achievement.

- **Ranking** – the total number of points of each user is displayed in a ranking table, with top users at the top of the table. This mechanism motivates users through social comparison and competition.
- **Awards** – once per week the not too seriously meant title “MOKNOWPEDIA Held der Woche” (Hero of the Week) is awarded to the user who contributed most during the week. This mechanism is a chance for newcomers to quickly achieve something. It rewards short-term contribution.

Metrics and calculating reputation

To compute the reputation that is awarded to user u for contributing to articles $a \in A$, CollabReview depends on article importance $w(a)$, quality $q(a)$ and contribution ratios $C(u, a)$ that are determined by an authorship function. With this data, the reputation score s_u of every user u is defined as the average of twice-weighted article quality scores (plus the score obtained from their reviews s_u^r)

$$s_u = \frac{\sum_{a \in A} q(a)w(a)C(u, a)}{|A|} + s_u^r.$$

Article importance depends on two factors: a logarithmic function of article size in characters, and the average number of article views per week. Size is included to represent the fact that larger articles potentially contain more information because they aggregate more contributions. The page view frequency is included to accommodate the fact that the article probably contains more interesting information, and is referenced more often in other articles. It should therefore have a higher weight. Multiplying the two numbers results in an article's estimated importance $w(a)$.

The contribution ratio of a user for an article is determined by his authorship in the article, i.e. it depends on the ratio of characters contributed by that user. A line of text is considered as being contributed by a user, when that user last edited the line. We use an adaptation of the algorithm by Prause (2009): each new revision of an article is compared to several of its probable ancestors. When an optimal candidate has been found, the algorithm determines an author for the newly added or modified lines; unmodified lines will retain their author. If a non-trivial line (e.g. one that contains only whitespace) is found to be a clone of an already existing line, authorship information is copied from there. Otherwise, the creator of the revision is considered as the line's author.

Implementation

This section describes how we implemented our concept in MOKNOWPEDIA. We describe the software architecture and present user interfaces.

Software architecture

The software architecture of MOKNOWPEDIA combines two distinct softwares into one integrated platform: CollabReview and MediaWiki.

MediaWiki is one of the most widely used wiki softwares. It is written in PHP and published free under the GNU General Public License (GPL). MediaWiki is easy to use, but has diverse functionalities and is very customizable (Barrett, 2008). The online encyclopedia Wikipedia is operated by the MediaWiki software.

CollabReview (Prause and Apelt, 2008) is a Java-based web application for reputation management in collaboratively written software source code. With multiple developers engaged in developing software code, responsibility for a specific piece of code is difficult to assign. Nonetheless, responsibility is a major factor in achieving quality and preventing code from being developed carelessly. CollabReview statistically acquires per developer per document accountabilities and enables learning and self-monitoring processes within a development team while maintaining anonymity to a certain degree in order to not endanger team spirit.




Figure 1. Architecture of the integrated MOKNOWPEDIA platform.

Figure 1 depicts how the MOKNOWPEDIA platform is composed of the sub-systems MediaWiki and CollabReview. Users access MOKNOWPEDIA with their browsers. When logging on to MOKNOWPEDIA, a single-sign-on mechanism ensures that a user is automatically logged on to both sub-systems. All interaction of the user with the MOKNOWPEDIA happens through the user interface of MediaWiki. However, the MediaWiki user interface has been extended with widgets that pass through CollabReview functionality (see Section “User interfaces”).

Apart from the user interface, some necessary communication happens directly between MediaWiki and CollabReview. In order to compute reputation scores, all articles and their revisions are run through an additional processing step by CollabReview. It analyzes articles for authorship and responsibility information, and stores review and quality data in its own database. Additionally, CollabReview’s own user management was removed. CollabReview, directly accesses the MediaWiki user database, instead.

User interfaces

The user interface of MOKNOWPEDIA is a mash-up of the original MediaWiki interface and additional services provided by CollabReview. MediaWiki's menu on the left is augmented with a widget that allows to submit or update a review for the current article. It is visible on every page of MOKNOWPEDIA (see Figure 2).




Figure 2. Review widget in the menu bar, left of a MOKNOWPEDIA article.

When entering a review for an article, a user can select a rating from one of “sehr gut” (very good), “gut” (good), “befriedigend” (satisfactory), “ausreichend” (sufficient), or “mangelhaft” (poor). He can also enter an optional review comment in the text box below. The comment is meant as justification of the rating, and as a hint on how to improve the page. The submit button (“Senden”) will enter the review into the database, and is send to the authors of the article via email. Below the button, there are three indicators: the stars represent the average quality rating as determined by the users’ reviews, the gauge in the middle shows the article’s importance level, and the traffic light shows how up-to-date the user’s review is. In this case, the review is outdated and should be refreshed. The link at the bottom of the widget (“Alle Bewertungen”) lets users view all reviews for the article without revealing reviewer names (see Figure 3).

A screenshot of the starting page of MOKNOWPEDIA can be seen in Figure 3. Users typically arrive at this page first when opening MOKNOWPEDIA. Two more widgets are embedded in the article text: the left one is presenting the reputation ranking list of users. The user with the most reputation points has 100%, all other

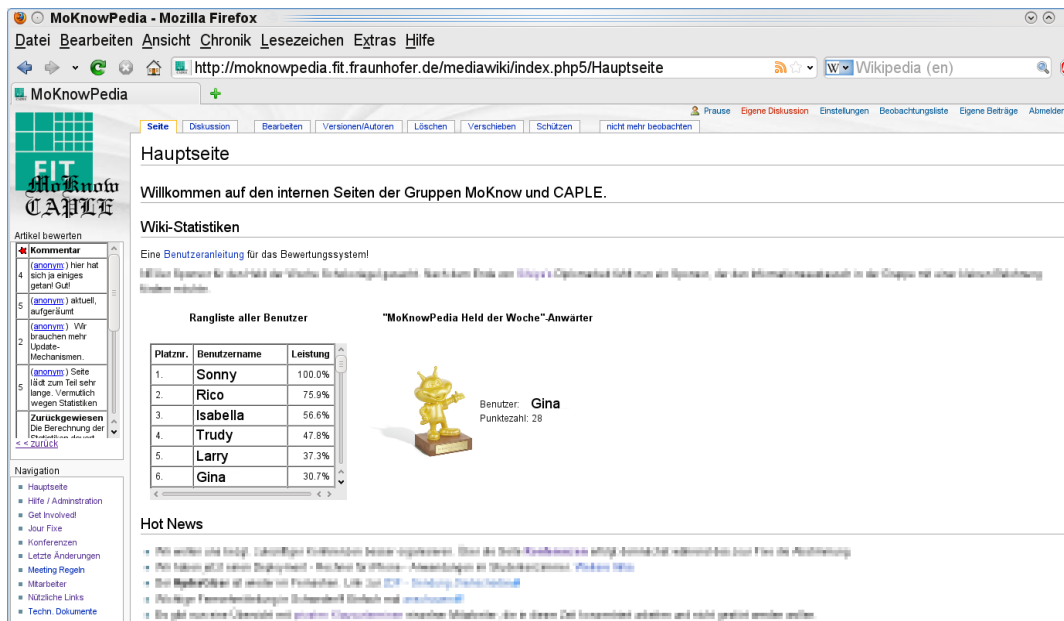


Figure 3. Main page with ranking and hero of the week; review widget shows reviews of article.

users' scores are given as percental value of the top contributor. The right widget shows who has achieved the most points in this week so far, and how many points he has. This is the Hero of the Week aspirant. Any other user that surpasses the aspirant will automatically become the aspirant himself. As long as no user has made at least one point, there is no aspirant.

Once a week, right before the weekly group meeting, the current aspirant becomes the Hero of the Week. The new hero is automatically announced via email on the group's mailing list. Additionally, he receives a chocolate bar during the meeting. After this, the aspirants are reset, so that for the next week everybody has the same chance of being the aspirant and becoming Hero of the Week.

Figure 4 displays a cutout of the MOKNOWPEDIA's user page. This page lists all members of the work group and links to personal introductory pages. Every user entry uses a MediaWiki template to standardize the presentation form. The template has been extended to include a widget that will show each user's reputation level.

Evaluation

We evaluated our concept in a field-test that would last for several months. This section describes the approach, reports our findings, and interprets results.

Approach and threats to validity

The evaluation serves the purpose to find out, if enhanced MOKNOWPEDIA satisfies our six goals. However, the focus of our evaluation is explorative, and on qualitative data and human aspects. We obtained a good impression of the impact of the

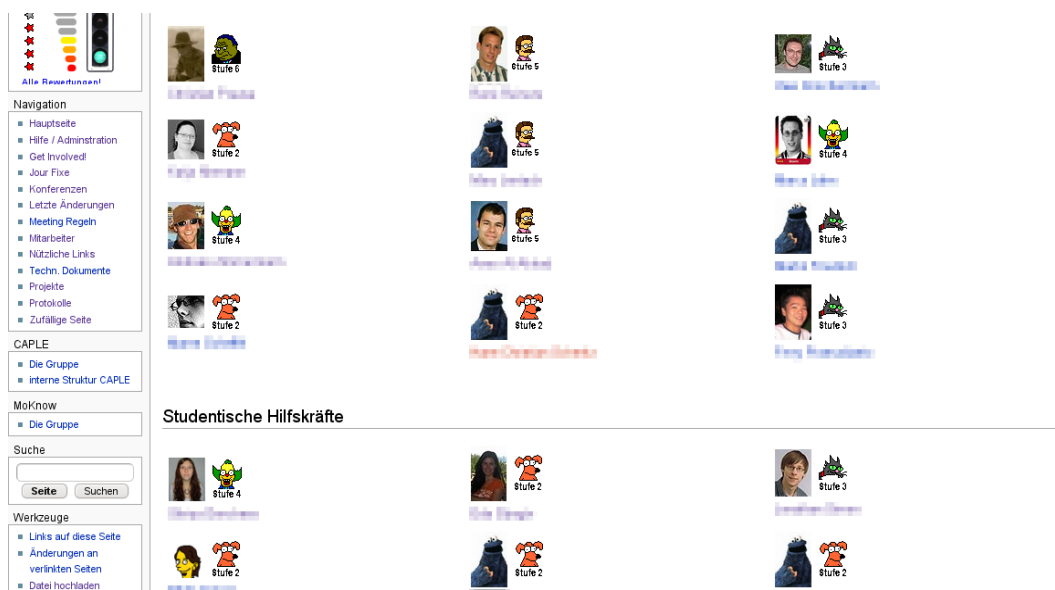


Figure 4. Cutout of MOKNOWPEDIA's user overview article and individual user reputation levels.

reputation mechanisms on a personal level, albeit at the cost of only having a small number of probands and limited statistical significance. At start and end of the evaluation, the work group had 18 and 16 members, respectively. Due to fluctuation in personnel, only 15 employees were able to participate for the full evaluation period.

Before making our modifications to MOKNOWPEDIA, we interviewed all members of our work groups and had them fill out questionnaires. This was to capture the initial spectrum of opinions as a base line. A second round of interviews was executed after the evaluation phase. The results of both interviews could then be compared. Yet we did not tell probands that we would interview them twice. A possible social desirability bias would then be in effect during both interviews and cancel itself out. Additionally, the time between the two interviews is probably too long for probands to remember earlier answers.

The evaluation phase itself — between the two interviews — ran for four months. During this time, we recorded any incidents and observations regarding the modified MOKNOWPEDIA. In addition to the interviews, we collected actual usage data, and compared it to data of the four months right before the evaluation. Although the evaluation phase of about four months of productive use is long, our results might still be influenced by seasonal events like holidays.

Socio-technical systems are highly complex due to the involved human beings that tend to be unpredictable. Changing only a minor factor in the environment or the people themselves, could lead to very different results. For example, if the work groups were not built on friendly cooperation and intrinsic motivation, if the environment were more competitive, or if the majority were not computer scientists, then results might have been different. Also, when the evaluation started, the MOKNOWPEDIA was about two years old. Usage behavior had had its time to establish and stabilize, with several attempts to improve contribution. Yet we cannot fully

preclude that the change in behavior we observed during the evaluation, is not due to some other effect.

Usage statistics

By analyzing the data in databases and server logs, we compare the four months before the review extensions were added to the four months of the evaluation period. Start of the evaluation period was February 18th, 2010.

The number of new article revisions raised from 320 to 517, which is an increase of about 62%. As the four pre-evaluation months span Christmas time, we also compared our evaluation phase to the same period in 2009. Here we found 363 new revisions, which still means an increase by 42%. The average number of characters per article in MOKNOWPEDIA grew from 5324 to 6134.

In addition to an increased number of contributions, we also counted an increased number of viewed articles. MediaWiki counted 176 page views in the four months before the evaluation phase, and 258 views during the evaluation. This is an increment of 47%. The number of page views is lower than that of changes because MediaWiki does not count page views that recur in a short period. We attribute this increment to advertising effects of the MOKNOWPEDIA extensions.

Finally, we wanted to know if users used the review tool actively. After the evaluation phase, we counted 237 reviews in the database. 155 reviews (more than half of all reviews) additionally contained sensible comments. In total, these reviews were submitted by 16 different users. Only one user who had participated for the evaluation's full duration had not submitted a single review.

Even after the evaluation phase is over, MOKNOWPEDIA's review facilities keep being used, and new reviews are being added continuously.

Results from pre- and post-evaluation interviews

Before and after the evaluation period, the users were interviewed with questionnaires. The next two sub-sections discuss the results of the interviews to evaluate

1. if contents improved, and if the reviewing extensions were accepted,
2. and to gather future design recommendations.

Quantitative feedback: effects on MOKNOWPEDIA and acceptance

The pre- and post-test questionnaires contained twelve statements to which users should express their degree of agreement on a Likert scale: "fully agree", "agree", "neutral", "disagree", "absolutely disagree". Items 1 to 8 concerned the effect that reviewing extensions had on the Wiki contents. The remaining four items related to the reviewing extensions themselves. Detailed agreement is presented in Figure 5.

Assigning one ("absolutely disagree") to five points ("fully agree") to each user's agreement, an average agreement score is computed. The average agreement of first and second interview can then be compared to see how agreement has




Figure 5. Results of pre- and post-evaluation interview.

#	Statement	Pts.	Chng.
1	MOKNOWPEDIA is generally attractive	4.0	+0.73
2	The articles in the wiki are substantial	3.6	+0.53
3	The total number of articles is satisfactory	3.3	+0.27
4	The articles are interesting to read	3.5	+0.60
5	The wiki contains important and useful information	4.5	+0.20
6	Article core information is comprehensible and easy to find	3.9	+0.53
7	The treated topics are diverse	4.0	+0.60
8	It is worth using MOKNOWPEDIA	4.1	+0.47
9	A reward motivates me to use the wiki more actively	3.0	+0.47
10	A mechanism to rate the quality of articles is valuable	3.9	+0.53
11	It makes sense to invest time in reviewing and commenting	3.6	+0.80

Table I. User agreement to items in post-test questionnaire.

changed. The texts of the different statements, translated into English, and the associated agreement scores — absolute and relative change between first and second interview — are listed in Table I. The answers substantiate that MOKNOWPEDIA and its contents improved during the evaluation. In the second interview, users disagreed with the positive statements about MOKNOWPEDIA in only seven cases, as compared to 22 before. Many users that were previously undecided, later agreed to positive statements.

Four statements regard the reviewing extensions themselves (see Table I). Before the test, there were several users that were not sure of whether rewards would motivate them (Item 9). Through the test, these users changed to agreement. Also, the number scepticists decreased so that afterwards a slight majority was in favor of the reviewing extensions. However, this statement is still the one with most objection. Similarly, most users initially did not think that they would invest time in writing reviews (Item 11). This has clearly changed towards acceptance: the number of agreeing users grew from four to ten, and disagreeing users decreased from six to two. Although the majority of users agreed from the start that reviewing would make sense, after the test even more agreed (Item 10).

We wanted to know, if the acceptance of the reviewing system was especially high among profiteers. But we found only weak correlations (Spearman rank) between the Items 9 to 11, and the user's position in the ranking list. The strongest correlation we found was $r_{Item11} = 0.25$ for Item 11 ($r_{Item9} = 0.14$, $r_{Item10} = -0.06$): Profiteers might be more willing to do reviewing work than others. Yet none of the correlations is significant due to the size of the test set. It could be that there is no relation between profiting from the review extensions and advocating them.

The last depicted item was mixed into the statements as an indicator of a possibly changing response bias. It was intentionally formulated diffuse (*The integration of a reward system into the wiki is interesting.*). The intention was that it could reveal a bias towards giving more positive feedback in the second interview. We observed only a small change in feedback behavior (3.7 points, +0.2), so it seems

that there is few social desirability bias, thus supporting that observed improvements in the other items are real.

Qualitative feedback: future design recommendations

In the second interview, there were three additional statements about the rewarding mechanisms. All three rewarding mechanisms were not perceived as very effective at motivating users. On the average, they were all weakly rejected. The ranking list was perceived as most effective (2.9 points), followed by the Hero of the Week (2.7 points), and finally the level hierarchy (2.6 points).

The questionnaire invited users to provide further remarks of what had or would have motivated them, what they had liked or disliked, and if it changed their behavior. In the following, we provide the major design recommendations:

A few users said that they would not feel motivated by the rewards in any way. Instead, they answered it would have required much larger rewards, or coercion. They had a distrust against the reviewing system, and considered reviewing articles as distracting and a waste of time that they could rather spend on writing articles. Some users bore concern that the extensions could be manipulated, and be attacked systematically to harm targeted people.

Positive remarks were that the design was pleasant, the integration fitted well into the existing MOKNOWPEDIA. It was simple to access, and easy to use and understand. The comments were regarded as very useful because they motivated and helped with improving articles. Users felt that the reviews made them read and tend to more articles. Reviews also gave them a feeling of trust. But also the anonymity of writing reviews was appreciated. Several users valued that they could see their own activity, or compare it to the others' activity. Some users noted a positive fun factor in the rewarding mechanisms, which they regarded as entertaining.

One user testified that the rewards had motivated him to start contributing to MOKNOWPEDIA. Others browsed more vigilantly to be able to give reviews, or wrote articles more carefully. A major contributor, however, said that he had reduced his contributions to give others a chance. A couple of users felt appealed by the rewards and answered that it had changed their behavior, but noted that the motivation was not sustainable, and could not withstand other pressing tasks. The lack of time was given as the major reason for not contributing to articles as well as not reviewing. A few users told us that they would not care for the rewards and, in fact, were using the wiki because they were convinced by it, but they also said that the rewards made them think more about the wiki.

Wishes for future improvements were to harden MOKNOWPEDIA against manipulation, to provide a user with hints how to earn his next points, a way to immediately respond to reviews, and to have computer-generated initial quality assessments. The value of rewards should be increased (e.g. an additional day off), and getting a certain amount of reputation points per week should be compulsory. Even punishment like paying money into a thank-you box was proposed. Another wish was to have user groups so that teams could compete against each other.

Observations during evaluation phase

This section describes observations that either we ourselves or designated confederates made during the evaluation phase. It is a collection of episodes that clarify the reception of MOKNOWPEDIA in the work group.

Curiosity Trudy had missed the meeting in which the MOKNOWPEDIA extensions were introduced. She was curious and went to her colleagues to ask them how she could try the extensions. As the extensions were already installed in MOKNOWPEDIA, she could start using them without problems.

Missing comment Gina goes to one of the administrators of MOKNOWPEDIA and complains that somebody gave her article a bad review rating without writing a comment. The administrators themselves could not do anything about it, but after this incident we updated MOKNOWPEDIA to allow rejecting of seemingly unjustified reviews. Yet there were no further incidents like this one.

First activity Lou, who was registered in MOKNOWPEDIA but had never written an article before, started to regularly write and review articles. Within one week, he made it from zero points to the fifth position in the rankings. He announced that the extensions motivated him to start contributing to MOKNOWPEDIA.

Frustration During the evaluation, Izzy became Hero of the Week several times in a row. This frustrated other users. They complained that there must have been something wrong with MOKNOWPEDIA. As a consequence of this, Izzy deliberately chose to contribute less to give others a chance to earn the title, too.

Playful competition Incited by the fact that a colleague had overtaken him in the rankings, Larry went to several colleagues, asking them for hints on how he could quickly gather some more points. They gave him the advice that an easy way to get some points was to write reviews. Within a few hours, Larry reviewed dozens of articles until he had again overtaken said colleague. For the next weeks, he kept an eye on the rankings, watching to remain ahead of that colleague. Although he reviewed articles primarily to get the points, his reviews were fair and contained valuable review comments.

Anti-hero Caitlin had been a minor contributor, only contributing what she had explicitly been requested to. She never submitted a review and for a long time none of her articles had received a review, leaving her with zero points. She was proud of being the last one in the reputation ranking with zero points. One day, however, she received a number of points; not because she had done something recently, but because one of her older articles had received a good review. So she came to our office and asked if she could rate her own articles down to get her old position back. We told her that she then would get points for the reviews, what she did not like.

Unique positions — even if not directly associated with a positive merit — have a certain attraction to users.

Improvement Through a review comment, Martin was adverted to the template that is recommended for creating user profiles. He had not used the template before, so he asked his colleagues about the profile and had them help him to change his profile page. Feedback from the reviews was helpful to him and for improving article quality.

Manipulation To quickly gain a considerable amount of points and become Hero of the Week, Stan hacked the MOKNOWPEDIA by writing a script that rated all articles “good”. Although he immediately received lots of points, he still did not win the title because the articles’ authors received more points than him. Additionally, hundreds of emails for the reviews were sent within seconds. The incident shows that — given enough incentive — attacks against the system become probable. We removed all reviews that were created through the attack and negotiated that no one would attempt such an attack anymore.

Teasing From his office, Rico saw Sonny at the end of the hallway, he laughed and jokingly called: “Hey Sonny, I rated your last meeting minutes ‘poor’!” Yet a few minutes later Rico changed the review to a better and more appropriate rating. The users used MOKNOWPEDIA to tease each other, leading to a form of viral advertising.

Interpretation of results

The users expressed the wish to have more time to make their contributions and reviews, or even want to be forced by their managers. With our concept, we take a different approach. We think that if users have enough motivation to contribute, then they themselves will find the time to do so. Although purely intrinsic motivation would be preferable, our concept has some extrinsically motivating effects that lead to the desired effects.

Earlier in this paper, we defined six goals that we want to achieve with our concept. Our first goal (G1) is to increase the volume of contents in MOKNOWPEDIA. The evaluation shows that we have reached this goal: Users feel that the number of articles increased. We measured an increase in content size, too. The same holds true for quality, where users think that interestingness and understandability increased. Additionally, reviews provide confidence in article content, further adding to their quality (G2). MOKNOWPEDIA became more attractive through this.

With our extensions, we were able to draw more users to MOKNOWPEDIA. Except for very few users that denied the review concept outright (one of them was a quite active contributor of articles), there was no user that did not at least contribute some reviews. At least one user admitted to have started contributing because of the rewarding system. Although we were not able to achieve active contribution by

everyone without exceptions, our concept can be considered successful at drawing additional users to MOKNOWPEDIA (G3).

The additional cost of operation of our extensions is negligible. All that is needed is a little more computing power, and a small reward for the Hero of the Week (G4). By its design, our reputation mechanism does not prescribe users what to do; in our implementation, no one is forced to comply and there is no threat or punishment involved (G5). Similarly, we did not record any events that would suggest the existence of undesirable phenomena like rat races, or bullying (G6). Instead, we observed that the reputation system was perceived as fun.

We acknowledge that all six goals are closely related to the subjective value of reputation for the individual. The more reputation is worth, the more efforts will users invest to attain a high reputation (Jøsang et al., 2005). For example, by relying on coercion or higher rewards, and thereby giving up on goals G4 or G5, the value of reputation can be increased. This will lead to greater successes in goals G1, G2 and G3, but at the cost of risking destructive effects (G6).

Related Work

In social media and the Internet, *quality* is often interchangeably used with trust and credibility. Reputation systems are central to these communities (Jøsang et al., 2005; Resnick et al., 2000). Ordinary wiki software like MediaWiki offers no integrated reputation system. However, the demand for such systems has pushed researchers and practitioners alike to research and develop extensions. Freely available extensions for the MediaWiki software are, for instance, *Review*⁴, which lets logged in users decide if an article is ready to be printed, or *AjaxRatingScript*⁵, where quality ratings from different users are collected and averaged. This approach is quite similar to how we determine article quality, except that it does not consider ageing of reviews.

Wiki extensions that compute reputations for authors are much fewer. The social rewarding extension of Hoisl et al. (2007) is similar to our approach. They determine a score for every author in order to recompense users investing their time and effort contributing to the wiki. However, their approach differs from ours in the way how reputation is computed. They do not make a distinction between article quality and article importance. Also, they calculate responsibility for articles from size differences and passed time instead of considering the authorship of the article's text. Next, they rely on only one social reward: displaying reputation in a ranking list. In their concept, reviewing is not a valued contribution. However, as they have automatic means of assessing article quality, it is not so important for them, too. Hoisl et al. did not test their tool with users, and do not provide experience or evaluation data.

Adler et al. (2008) present a system that computes trust values for articles of the

⁴ <http://www.mediawiki.org/wiki/Extension:Review>

⁵ <http://www.mediawiki.org/wiki/Extension:AjaxRatingScript>

online encyclopedia Wikipedia. Readers can use WIKITRUST to identify words that are deemed trustworthy or not trustworthy, respectively. The higher an author's reputation, the higher is the predicted quality of his contributions. Consequently, such content should remain unchanged for a longer time. They predict the imminence of deletion of a word in the next revision of an article by means of recall and precision. Yet WIKITRUST does not publish the reputation scores it computes. It does not want to change contributor behavior but rather help readers to trust in content.

In their study, Holtzblatt et al. (2010) similarly report that the cost of contributing to articles is an important impeding factor to wiki contribution. They conclude that social incentives are worth exploring in corporate environments.

Muller et al. (2009) argue in their paper on metrics for enterprise social software "Social media applications should record and analyze the activities of information-consumers, ... to understand what information and knowledge is proving to be valuable to employees ...". We believe that our approach provides a valuable contribution towards that direction. Our evaluation demonstrates that participation in our wiki was increased both from contributors as well as commentators and evaluators due to the provision of statistical data about user activities.

Farzan et al. (2008) use social incentives (levels and ranking list) to motivate contributions to a corporate social network. They find that social rewards can effectively increase contribution and that different kinds of incentives are needed to motivate different people. A major recommendation from their study is, however, that points should decay. The hero of the week is a similar mechanism in MOKNOWPEDIA. Yet we found no indication that points should generally decay over time; at least not in wikis.

Conclusion

We implemented an extension to MOKNOWPEDIA that lets users review, comment on and rate articles. For each article, every user can provide exactly one review, so that potentially there is a set of reviews for the article. While the review comment is plain text that helps authors with how to improve their article, the review ratings are aggregated to obtain an average quality assessment. This helps users identify high and low quality articles. In this context, quality can be defined as the readability, understandability and information payload of an article. However, in all comparable solutions it is normally up to the community to define their notion of quality.

We extend the basic concept of reviewing tools by adding the CollabReview reputation system that tracks and reconstructs how much individual users do for the wiki. In order to calculate personal reputation scores, the reputation system determines who has contributed to what articles, how important these articles are, and what the quality of the articles is. Authors that have contributed much to high quality articles will have a high reputation score, while authors that have made only few or low quality contributions will have low reputation scores. Additionally, we also regard the users' reviews as valuable contributions to articles that earns the reviewers a few more reputation points.

The idea is that if authors are given something in exchange for high reputation scores, then they will — to further raise their reputation score — try to contribute high quality articles, increase the quality of existing low quality articles, or provide hints for improvement. For the MOKNOWPEDIA, three different social rewarding mechanisms are simultaneously fueled by the reputation scores: a ranking list, a Hero of the Week award, and a level hierarchy. Of these three, the ranking list seems to be the most effective rewarding mechanism.

After installing the reputation system extension in MOKNOWPEDIA, we evaluated its use for several months with two work groups. The goal of the evaluation we conducted was to gather first experiences with such a system. The focus was on qualitative feedback from observations in the field and interviews with users. However, we also collected some hard usage numbers.

All in all, the evaluation results are promising. Not only have the employees accessed the wiki more often, but they have also read more articles, and made more contributions. The quality and quantity of the articles in MOKNOWPEDIA increased. Most users accepted our extensions and enjoyed using them. Some users even made their first contributions because of our system. The trial allayed our initial fear that the system might be misused as a tool for bullying, or be confronted with total rejection.

Besides showing the overall good success of our concept, we inform the design of similar future systems with the findings from our evaluation. The usage observations provide a vivid picture of the social processes. These are complemented by improvement hints gathered through the final interviews.

Acknowledgments

This paper is based on Dencheva (2010). The research was supported by the Hydra (IST-2005-034891) and BRIDGE (project number 261817) EU projects.

References

- Adler, B. T., K. Chatterjee, L. de Alfaro, M. Faella, I. Pye, and V. Raman (2008): ‘Assigning Trust to Wikipedia Content’. In: *International Symposium on Wikis (WikiSym)*.
- Antin, J. and C. Cheshire (2010): ‘Readers Are Not Free-riders: Reading as a Form of Participation on Wikipedia’. In: *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW)*. New York, NY, USA, pp. 127–130, ACM.
- Barrett, D. J. (2008): *MediaWiki*. USA: O’Reilly Media.
- Danowski, P. and J. Voss (2004): ‘Bibliothek, Information und Dokumentation in der Wikipedia’. *Information, Wissenschaft und Praxis*, vol. 55, no. 8.
- Dencheva, S. (2010): ‘Entwicklung und Evaluation einer soziotechnischen Komponente zur dynamischen Selbstmoderation in einem Firmeninformationssystem’. Diplomarbeit, Fraunhofer FIT, Universität Bonn.

- Elster, J. (1989): *The cement of society: A study of social order*. Cambridge University Press.
- Farzan, R., J. M. DiMicco, D. R. Millen, C. Dugan, W. Geyer, and E. A. Brownholtz (2008): 'Results from Deploying a Participation Incentive Mechanism within the Enterprise'. In: *26th Annual SIGCHI Conference on Human Factors in Computing Systems*. New York, NY, USA, pp. 563–572, ACM.
- Hoisl, B., W. Aigner, and S. Miksch (2007): 'Social Rewarding in Wiki Systems — Motivating the Community'. In: *Online Communities and Social Computing*. Berlin Heidelberg New York, p. 362–371, Springer.
- Holtzblatt, L. J., L. E. Damianos, and D. Weiss (2010): 'Factors Impeding Wiki Use in the Enterprise: a Case Study'. In: *28th International Conference Extended Abstracts on Human Factors in Computing Systems*. New York, NY, USA, pp. 4661–4676, ACM.
- Jøsang, A., R. Ismail, and C. Boyd (2005): 'A survey of trust and reputation systems for online service provision'. *Decision Support Systems*, vol. 43, no. 2, pp. 618–644.
- Kittur, A. and R. E. Kraut (2010): 'Beyond Wikipedia: Coordination and Conflict in Online Production Groups'. In: *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW)*. New York, NY, USA, pp. 215–224, ACM.
- Kraut, R., M. L. Maher, J. Olson, T. W. Malone, P. Pirolli, and J. C. Thomas (2010): 'Scientific Foundations: A Case for Technology-Mediated Social-Participation Theory'. *Computer*, vol. 43, no. 11, pp. 22–28.
- McAdam, R. and S. McCreedy (2000): 'A critique of knowledge management: using a social constructionist model'. *New Technology, Work and Employment*, vol. 15, no. 2, pp. 155–168.
- Muller, M. J., J. Freyne, C. Dugan, D. R. Millen, and J. Thom-Santelli (2009): 'Return On Contribution (ROC): A Metric for Enterprise Social Software'. In: I. Wagner, H. Tellioğlu, E. Balka, C. Simone, and L. Ciolfi (eds.): *Proceedings of the 11th European Conference on Computer-Supported Cooperative Work (ECSCW)*. pp. 143–150, Springer London.
- Prause, C. R. (2009): 'Maintaining Fine-Grained Code Metadata Regardless of Moving, Copying and Merging'. In: *Proceedings of the IEEE International Working Conference on Source Code Analysis and Manipulation*. Los Alamitos, CA, USA, pp. 109–118, IEEE Computer Society.
- Prause, C. R. and S. Apelt (2008): 'An approach for continuous inspection of source code'. In: *Proceedings of the 6th international workshop on software quality*. Leipzig, Germany, pp. 17–22, ACM New York, NY, USA.
- Reichling, T. and M. Veith (2005): 'Expertise Sharing in a Heterogeneous Organizational Environment'. In: H. Gellersen, K. Schmidt, M. Beaudouin-Lafon, and W. Mackay (eds.): *Proceedings of the Ninth European Conference on Computer-Supported Cooperative Work (ECSCW)*. pp. 325–345, Springer-Verlag.
- Resnick, P., K. Kuwabara, R. Zeckhauser, and E. Friedman (2000): 'Reputation systems'. *Communications of the ACM*, vol. 43, no. 12, pp. 45–48.
- Rohs, M. (2007): 'Qualitative Forschung auf dem Wiki Way'. *Technologie verändert die Bildungsforschung*, vol. 2, no. 4, pp. 1–17.