

Coordinating Tasks on the Commons: Designing for Personal Goals, Expertise and Serendipity

Michel Krieger, Emily Stark, Scott R. Klemmer
Stanford HCI Group
Computer Science Dept, Stanford, CA 94305
[mkriegerlestarklsrk]@cs.stanford.edu

ABSTRACT

How is work created, assigned, and completed on large-scale, crowd-powered systems like Wikipedia? And what design principles might enable these federated online systems to be more effective? This paper reports on a qualitative study of work and task practices on Wikipedia. Despite the availability of tag-based community-wide task assignment mechanisms, informants reported that self-directed goals, within-topic expertise, and fortuitous discovery are more frequently used than community-tagged tasks. We examine how Wikipedia editors organize their actions and the actions of other participants, and what implications this has for understanding, and building tools for, crowd-powered systems, or any web site where the main force of production comes from a crowd of online participants. From these observations and insights, we developed *WikiTasks*, a tool that integrates with Wikipedia and supports both grassroots creation of site-wide tasks and self-selection of personal tasks, accepted from this larger pool of community tasks.

Author Keywords

Wikipedia, Crowdsourcing, Task management, Social software

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

On the community-written encyclopedia Wikipedia, there are no formal requirements for participation, and no participant is held accountable for his or her level of work. However, work still gets done, and Wikipedia recently surpassed the 2.5 million article mark [11]. The growth of Wikipedia is even more surprising given that the site does not follow typical practices for large-scale coordination such as GANTT charts and top-down organization.

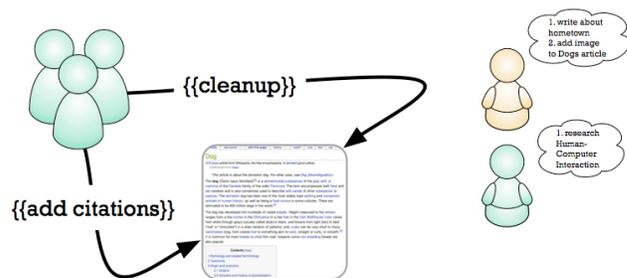
Indeed, Wikipedia's success stands in contrast to how one

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2009, April 3 - 9, 2009, Boston, MA, USA.

Copyright 2009 ACM 978-1-60558-246-7/07/0004...\$5.00.

Figure 1. Current task model on Wikipedia. Tasks are assigned on a page-level to the entire community, while users have their own set of goals.

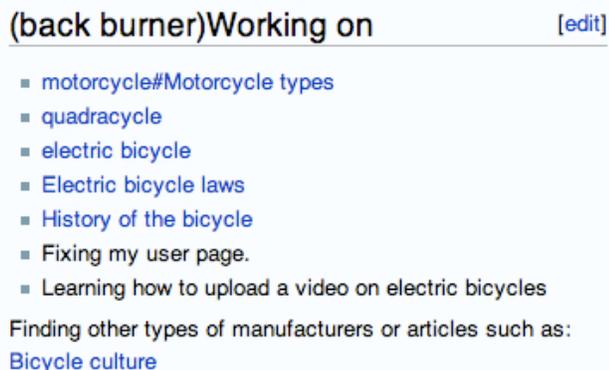


might picture a traditional organization; it is most closely matched to volunteering efforts and open-source projects in its structure [20, 24]. There are no managers and chains of responsibility; administrative duties are earned, and seeking to “rise in the ranks” is discouraged [23]. There is no single deadline; the encyclopedia lives in a double state of having the deadline of “right now” – after all, millions of people visit Wikipedia every month, and the community would like to present as complete a product as possible to them—and of “probably never” – collecting all the World’s encyclopedic knowledge is an infinite task.

Since traditional coordination mechanisms are not present at the macro-level on Wikipedia, individual participants (editors) are typically not given daily guidance as to the structure of their work at the micro-level. Editors can flag pages as needing attention or work through “template tags”, which signal to the entire community that a particular page needs additional citations, better copywriting, or other attention. However, no single person (or even group) has that task assigned to them (Figure 1).

It is therefore an open question how Wikipedia editors decide to structure their work on a day-to-day basis. Collaborative efforts such as the Wikipedia have been often examined through the lenses of motivation and degrees of participation [9, 14, 20], with less attention paid to how individuals in the community plan and execute actionable tasks within the scope of the project. Prior work has posited the question: “How is it that the self-allocation of effort by individuals on Wikipedia can be efficient?” [28]. This paper seeks to examine how exactly this self-allocation occurs, and where break-

Figure 2. Currently, users resort to maintaining ad hoc lists of current tasks on their own user pages. These lists are independent from the overall community’s tasks.



downs occur. Further, it seeks to resolve the tension, also raised in prior work, between Benkler’s assertion that peer-production – such as the production on Wikipedia – benefits from small, organized, individually assigned tasks [7], and the current state of free-form task tags on Wikipedia.

Wikipedia provides an intriguing space for tool design. The existing flexibility of the MediaWiki platform has allowed practices to emerge independently of tool constraints. Now, like the proverbial University paving foot paths after the behavior of students, user-centered design can surface these practices and ponder how to develop more focused designs around them. To this end, this paper examines how Wikipedia participants organize their actions and the actions of other participants, and what implications this has for understanding, and building tools for, crowd-powered systems, or any website where the main force of production comes from a crowd of online participants.

The remainder of the paper is structured as follows: we present prior work on Wikipedia and individual task management; we describe our paper’s qualitative research method; we explore task management for each Wikipedia role (and the insights gained from studying this management); and finally, present WikiTasks, our prototype social software task management system.

Related Work

There are two principal areas of research that inform our work: research on Wikipedia, and research on task management. Our work leverages the body of research on Wikipedia to understand the general patterns of participation of Wikipedia editors, and seeks to lend qualitative explanations to quantitative patterns previously found through Wikipedia visualizations. Our system’s design is strongly informed by prior work on task list manager design.

On Wikipedia

Due to its transparency (complete XML dumps of the Wikipedia are available for researchers) and popularity, Wikipedia has been the subject of several academic studies in the

past few years. Bryant et al. provide a seminal paper that describes the structure and basic participation patterns on Wikipedia [9]. This work was informative in establishing our framework of roles on Wikipedia.

Other Wikipedia research can be broken down into five primary categories: quantitatively understanding the participation and production of Wikipedia, connecting Wikipedia users to work, improving trust and article quality, analyzing the policy and structure of Wikipedia, and using Wikipedia as a corpus towards another goal (Natural Language Processing, for example). Since the first two of these are the categories most relevant to our own work, the following section addresses both. We suggest [1], [8], [15] as exemplary papers in the remaining three categories.

Wikipedia’s massive article base, as well as the rich edit history that is available for open perusal by researchers, has naturally led to several quantitative, large-scale data analyses efforts. Almeida et al. [3] examined 48 million edits on Wikipedia and described the evolution of Wikipedia, finding that the number of edits increases exponentially due to an increasing number of individual contributors. Another evolution observed by Kittur et al. [16] is the transition of contributions from mostly administrator-based to a rising “bourgeoisie” on Wikipedia of non-administrator power users. Kriplean et al. studied patterns of between-member encouragement through the practice of awarding “barnstars”, or badges of honor [19]. These studies’ strength is the breadth of their coverage, and the insights they provide at the macro level.

Visualization techniques have been developed for insight into this large Wikipedia corpus, including Viegas et al.’s “history flow” [26] and “chromograms” [28] visualizations. These studies have characterized interesting patterns on Wikipedia – for example, history flow helps show that article lengths do not stabilize over time, but instead fluctuate significantly, in an article’s unique pattern. Most relevant to our studies are the insights gained from the work on chromograms. This visualization technique surfaced particular editor behaviors: finding a particular task (or task type) and then performing a series of edits within this task “space” – other edits to the same article, or other similar articles. This prior work’s quantitative methods contributed descriptions of overall editor behavior that we found merited further qualitative study.

Also relevant to our work is the set of literature that attempts to connect users with work in social software collaborative systems. Cosley et al. developed the intelligent task routing technique [12, 13], which (in the general case) leverages a user’s behavior and demonstrated preferences and matches these contributors to work that would be valuable to the community. In the Wikipedia case [13], users’ edit histories, co-editing histories with other users, and lexicographically similar articles were used to form a user profile to match them to work. This work showed a four-fold improvement in recommendations when compared to randomly assigning articles. Our work explores how this “tell me what to do” behavior matches users’ actual Wikipedia activity, and what other methods users employ to find and track work.

On Task Management

Since our research focuses on the task practice of Wikipedia editors, the research on task management practice, and task management tool design, is relevant.

Tasks and task management are frequent topics of business management books, as well as receiving some coverage in the HCI community. Online communities have taken up Allen's *Getting Things Done* methodology [2] as a definitive work on task management, and online productivity sites frequently cite Allen's core principles: breaking down tasks into actionable next items; frequently pruning and triaging tasks; and keeping tasks associated with contexts.

The CSCW community has done work on distributed task management, including the *Task Manager* system [18]. The Task Manager extends the traditional to-do list to a large organization, allowing for asynchronous, spatially distributed collaboration. We take inspiration from the flexibility of Task Manager, as well as its user-centered design in the context of the EuroCoOp project. HCI work on task management has also been associated with its relationship to e-mail. *Taskmaster* implements a "thrask"-based system that creates threads of tasks that span e-mail messages, drafts, and links [6]. Our largest inspiration from Taskmaster is the contextual display of tasks in the e-mail inbox, rather than in a separate application.

Seeking to extend some of these tool-building insights into generalized principles, Belloti et al. [5] performed a study of task management practice, and suggested a set of design principles in task list management design. These include: faceted views into tasks; capturing task history; capturing time constraints; tying tasks to context; displaying social relations; and capturing task information away from the system itself. The design of WikiTasks closely follows these principles.

Recently, *Jourknow* has explored light-weight capture of information scraps—including tasks—in a light-weight, context-linked manner [17], which provided us inspiration on how to improve the current Wikipedia task capture system, that uses user-contributed template tags to mark tasks.

Method

Our investigation began with three principal guiding questions regarding the everyday task practices of Wikipedia editors:

1. On a daily basis, how do Wikipedia editors select and track tasks to perform on the site?
2. How does this vary across participation levels on Wikipedia, from beginning editor to administrator?
3. What implications does this have for the design of tools for Wikipedia?

To obtain a rich understanding of the participation patterns on Wikipedia, we undertook a three-month qualitative study with three primary components: we involved ourselves with

the community over the course of two months, joining "Recent Changes" patrol, participating in WikiProjects, and making edits through out the site. We interviewed 15 current editors (through e-mail, IRC, and phone). Finally, we prototyped ideas for task management on Wikipedia that we validated with editors early in the design process. Editors were predominantly from the English-language Wikipedia, though one editor was also an administrator for the Belarusian and Russian Wikipedias. Participants were recruited through direct contact on Wikipedia, as well as calls for participation on the English-language Wikipedia IRC channel.

Participants were recruited through three venues: initiating contact through their Wikipedia Talk pages, initiating contact through the #wikipedia-en IRC channel, and the Bay Area Wikipedians mailing list. The editors we contacted through their Talk pages were WikiProject participants or administrators. 2 interviews were conducted over the phone with Bay Area Wikipedians, both with Advanced Editors. 9 interviews were conducted through IRC, with 4 Advanced Editors and 5 editors who we would characterize as Beginning Editors. The remaining interviews were conducted through e-mail, with 2 Advanced Editors (and WikiProject maintainers) and 2 Beginning Editors. We conducted a brief exchange with Jimmy Wales, Wikipedia founder and "God"-level participant on Wikipedia, which provided some high-level overview of current Wikipedia monitoring tools.

Questions to participants were tailored to their specific role. For example, when interviewing an admin for the Beer WikiProject, we probed the role he played in delegating and organizing tasks; his own editing pattern (spurts or continuous); and what characterized successful collaborations on Wikipedia.

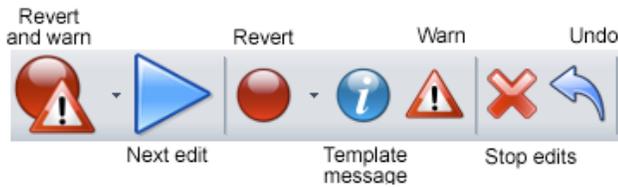
Results

Roles on Wikipedia

Participation on Wikipedia has previously [9] been characterized as transitioning from beginning editors to full-blown advanced editors ("Wikipedians"). Through our interviews, we sought to separate these broad categories into buckets ranging from minimal contact with the encyclopedia to heavy involvement through the parent Wikimedia Foundation. From our interviews, the data suggests this list of roles:

1. Passers-by (arrive at WP through Google search, read, and then leave, potentially without even understanding the way WP is written and produced)
2. Readers (informed consumers of Wikipedia)
3. Lurkers (are familiar with WP's system and policies, but perform few edits or contribution; sometimes are interested by the drama that plays out in the Talk pages more than the encyclopedic content)
4. Beginning editors (start editing on a more regular basis; might move beyond grammar fixes to contributing significant content)
5. Advanced editors ("Wikipedians") (contribute full articles, or contribute to article's path towards Good Article or Fea-

Figure 3. The Huggle application automates the multiple steps involved in reverting an edit through its toolbar



tured Article status; might join or even start a WikiProject to organize around a particular topic)

6. Administrators (experienced users who have power to perform bureaucratic and administrative tasks – with this power, however, often comes a reduced involvement in editing and contributing new articles, due to time load)
7. “Gods” (founding members, Arbitration Committee, or Wikimedia employees, who have final say in controversial topics; involved mostly as oversight)

Though some of these roles are formalized (Administrators and “God”-level users), for the middle chunk of Wikipedia participants, behavior is self-determined rather than prescribed. Unlike systems where experience and contributions provide increased privileges, a newcomer to Wikipedia could be contributing content at the advanced editor level in a matter of a week or less, or could wait years before stepping up responsibility.

Task management by participant type

With these categories of people in mind, we probed further on the middle group of participation levels: beginning editors, advanced editors, and administrators. The following sections discuss how each of these groups finds, manages, and completes actions.

Beginning Editor

This first section addresses the action management of beginning editors. These editors are already aware of the editing capabilities of Wikipedia, and thus organize their actions around this activity. Task selection for beginning editors tends to consist of a “pull”, rather than a “push” mechanism.

Several of these participants join the “Recent Changes Patrol”, a group that monitors the latest modifications to Wikipedia to immediately revert vandalism [9]. Since identification of vandalism often involves little beyond simple recognition (most of the times not requiring any subject knowledge), it is a relatively easy activity for newcomers (and more expert participants) to perform. Community members have built software tools to support this activity; a popular one is *Huggle* (Figure 3), which presents the latest changes in a stream in one interface pane, and a large “Rollback” button for editors to immediately undo any vandalism. There are similar tools, all of which require editors to have met the most basic of participation requirements (having been on the site a few weeks, and made a few edits).

For editors who work in Recent Changes Patrol, managing tasks is a relatively easy issue – there is a never-ending stream of edits to check out and potentially revert, and editors who participate in this activity often report getting “hooked” on this activity, feeling it is rather game-like:

“We like to consider it a game of whack-a-mole.” –
Wikipedia editor

Recent Changes Patrol works well as an activity for editors because of its low barrier of entry, lack of self-doubt (there is no “expertise” required in identifying if “Look, ma, I can edit Wikipedia!” is a valid contribution), and constant availability of activity (and, therefore, actionable items). The activity here is similar to activity in video games, where actions are supplied by an external mechanism, and the participants’ responsibility is to respond to it in a time-dependent manner. It is a necessary and valuable role, but potentially detrimental towards the goal of having more people participate in content creation, as there is no clear path of participation from monitoring vandalism to adding information to articles. Further, research indicates that vandalism-related changes are only a relatively small fraction of the work on Wikipedia (around 6% as of 2006) [10].

For those beginning editors who don’t regularly participate in Recent Changes Patrol, the task of choosing what to edit is less clearly defined. The practice of watching pages through the Watchlist has been previously documented [9], [13], but beginning editors are less likely to use it for finding tasks to do, with one informant stating, “I really have no use for it”. Instead, the following strategies are present:

- Use the ‘Random Page’ feature to find a random page that needs attention.

Scattershot at best, this task-finding technique serves mostly as a cure for boredom. Given the large number of “stubs” on Wikipedia, it is likely that the Random Page feature will land the editor on a short page, of which she probably has no additional information to add at the moment. This behavior is most similar to the micro-task markets such as Amazon’s Mechanical Turk [4], where small, often repetitive tasks are performed by members of a crowd.

- Browse the Wikipedia within a constrained topic area, using category pages.

Similar to template tags, articles are categorized on Wikipedia through the addition of Category tags. Categories are non-exclusive; an article on “Mark Twain” might be in the “Authors” and the “Americans” category.

- Use one of the Wikipedia tools that suggests a task to do.

For this last category, there are a small number of tools available that seek to connect editors and articles. Wikipedia itself provides a simple page that shows pages that have been tagged as needing “cleanup”, or that need to be expanded. This tool orders tasks by age of tags. A separate tool, written by a Wikipedia participant, presents a random article and

uses general heuristics to suggest potential improvements to the article – such as adding an image, or expanding the text, or adding more intra-Wiki links [21]. This tool is separate from Wikipedia, and requires the editor to load a separate interface to receive a task suggestion. The most successful of these tools has come from research on providing an “intelligent task routing” interface for WP editors, implemented as an agent named “SuggestBot”. This routing increased participants likelihood of editing article by a factor of four, versus suggesting random articles that need attention [13].

As these editors often have not yet joined a WikiProject or carved out an area of interest, the feeling of “there’s nothing for me to do, I’m not an expert in anything” was common among interviewees. Efforts like the SuggestBot are a good step forward, but a more ground-up rethinking of how to encourage greater participation is required. Since these ideas are related to some advanced editor details, they will be described below.

Advanced Editor

This section continues the analysis by moving to advanced editors, where the role of an editor shifts slightly: no longer content to perform simply minor edits, an advanced editor is more likely to contribute significantly longer content. This often manifests itself as the editor joins a WikiProject and turns to this as the main source of tasks and activity. An advanced editor will invest the time in bringing an article from “Good Article” to “Featured Article” status, often as part of such a WikiProject. Previous work has shown how advanced editors (“Wikipedians”) carve out their own “voice” in the project; the collected observations demonstrate a similar shift from “pull” to “push” task management.

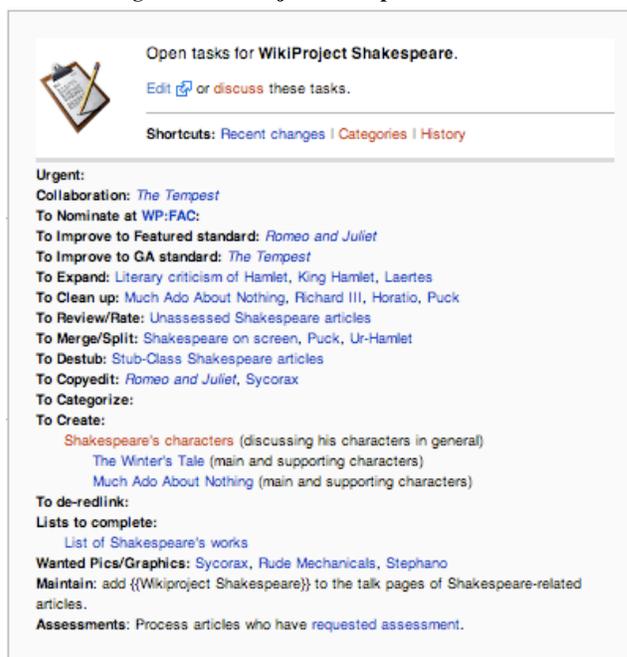
These editors often center their task management around one of these WikiProjects. These projects are an essential organizational pattern within Wikipedia. Just as communities that grow beyond a certain size need subdivisions to maintain fluid communication and participation, WikiProjects provide a sense of purpose, a smaller, more manageable set of actions, and the feeling of group membership. As of September 2008, there were 657 active WikiProjects on the English Wikipedia.

Two central activities are evident within WikiProjects: maintenance and expansion of content, and driving chosen pieces of content towards a very polished state. The former involves cleaning up messy articles, expanding ‘stubs’, and categorizing pages that should belong in the WikiProject. This is superficially similar to the actions of a beginning editor, though now within a smaller scope of articles and within a smaller community.

Instead of the broad, impersonal task lists in the Wikipedia general to-do lists, WikiProjects provide a more scoped view of what’s necessary within the project, as can be seen in Figure 4.

In keeping with the Wiki model, any editor can add or remove tasks to these lists.

Figure 4. WikiProject Shakespeare’s Task List



This seems, at first glance, to be ideal: a display of to-dos within an interest group, matching encyclopedia needs with editor interests. However, two main problems are present: lack of task specificity, and failure to match Wikipedians’ own action patterns. If one takes a “Getting Things Done” approach to this task list, one sees that the tasks, though scoped in interest, are still for the most part far too broad to be actionable. “Much Ado About Nothing” might require cleanup, but a passer-by would not know what exactly might need to be cleaned up merely from the presence of the item on the task list.

Beyond the specificity issue, most of our informants simply reported not using the WikiProject to-do lists. Rather than keeping to these lists, task choices from advanced editors were either self-directed content creation efforts for a specific article (informants reported keeping their own lists of “articles they’d like to write someday”), or fortuitous discovery of actions:

I generally don’t plan in advance what I am doing. Editing Wiki has become a fascinating hobby in which I feel I am doing something useful but enjoying myself at the same time. – Wikipedia editor and WikiProject contributor

This timely discovery of actions to do can also come from off-Wikipedia factors; for example, one Wikipedian reported that he researches and adds to articles relating to places where he will be visiting in the near future, or topics that have recently come up:

If I’m going somewhere - like Deptford or Dovedale or Amsterdam - then I will work on that article for a bit as

a way of gathering information about a place that I am about to visit. Or if I've seen a film, like I Am Legend, I may get involved in that article. – Wikipedia editor and WikiProject participant

The other previously mentioned WikiProject process is also worthy of note – bringing articles to a very polished state. This can take one of two forms: a scheduled collaboration, or a featured article effort. A scheduled collaboration is a timed “challenge” undertaken by a WikiProject to improve a particular page that falls within its coverage. Our informants reported that these collaborations are often unsuccessful, even though they are given prominent placement on a WikiProject's page.

From the gathered data, it seems these collaborations are unsuccessful for at least two reasons: first, Wikipedia participants do not react well to being “told what to do”, even when such funneling is done within a scope such as a WikiProject; second, the “promise” at the conclusion of the collaboration is often vague, seeking only to “make an article better”.

This is not to say that scheduled collaborations are never successful. On larger WikiProjects, the large number of participants means that, even with low percentage of participation on a scheduled collaboration, large changes can still be effected during these collaborations. The success of a WikiProject appears to depend on three main factors, according to our observations and informants: the number of editors interested in the topic (smaller interest groups are often subsumed into larger Projects as Task Forces), the group's ability to standardize conventions without inciting large arguments, and the availability of a wealth of off-Wiki articles of encyclopedic value that can be adapted into the Project's articles.

Often, a featured article effort is a more successful way of moving an article forwards. According to one of our informants, who was a WikiProject founder:

A push towards FA is generally a much more effective way to bring together editors with relevant knowledge than a scheduled collaboration. – Wikipedia editor and WikiProject founder

The highest form of recognition for an article on Wikipedia is for it to reach Featured Article status – Featured Articles are shown on the front landing page for an entire day, thus receiving attention and bringing positive “buzz” to the WikiProject(s) responsible for the page. Viegas et al. studied the deliberation process behind choosing nominated articles for inclusion in Featured Articles [27], but a separate, equally interesting process is the work done on an article to get it to that stage. Contrary to the image of Wikipedia as the joint effort of a large crowd of users, Featured Article efforts are often undertaken by a core group of 5 or less people.

For advanced editors, equally important as the process of performing actions as part of a WikiProject is the process of establishing which actions need to be performed by others. Editors can mark pages as being actionable by marking them

with a set of tags, such as `{{wikify}}` (indicates that a page needs intra-wiki links, or `{{notability}}` (indicates that it is unclear whether the topic of the article merits inclusion in an encyclopedia.)

Once again, the community has written a set of end-user programmed tools that facilitates this process. These tags serve two functions: readers of a page will see these tags, and might be enticed to participate by the call-to-action embedded in them; and the tagged pages will end up in special “smart search” listings of pages that need attention. This process is a significant departure from the process of delegation that occurs in a business, since no single person (or even group of people) has been charged with fixing these problems. The fact that much of this work remains not done could be a product of this lack of accountability – a common problem in peer-based production systems [7].

The rapid manner in which template tags can be assigned is dual-edged; while it means that advanced editors can quickly filter through articles to tag them for cleanup, it also means that the tags rarely have the granularity or guidance to be useful beyond a large sign-post that the article requires attention.

Relating back to the beginning editors, WikiProjects could serve as an effective bridge between them and advanced editors. As of now, there is no way of indicating what each editor is currently working on, or to claim responsibility for a set of actions. WikiProjects could serve as a meeting-point or clearinghouse for action items (defined by advanced editors), which would then be “taken up” by rising editors, giving the latter a sense of responsibility and shared ownership of the project.

In summary, our investigation found that advanced editors center their Wikipedia activity around WikiProjects, which serve as the functional units of organization on Wikipedia. These WikiProjects contain both day-to-day tasks marked by template tags, and wider tasks and collaboration – scheduled collaborations on particular articles, and Featured Article efforts.

Administrator

Finally, this section addresses the Administrators. These power users are responsible for the day-to-day maintenance of the Wikipedia community, and thus have a steady stream of work to do. Says one administrator:

Some admins are specialist, ie. they use their tools in certain areas. Some, like myself, are all-rounders. – Wikipedia Administrator

For “specialist” administrators, the main activity is to watch over an area of Wikipedia—at the scope of a WikiProject, for example—watching for higher-level questions that need resolving, such as edit controversies. For these administrators, the Watchlist is a valuable tool for finding actions to perform. Also, WikiProject Talk pages provide for issue escalation, as well as messages written on an administrator's talk page (the primary form of between-user communication on Wikipedia

is by editing another editor's talk page). For these specialist administrators, the task management is mostly ad hoc—tasks emerge as issues emerge.

For generalist administrators, there is never a lack of things to do. At any given moment on Wikipedia, pages are being marked for deletion, editors are being nominated for “adminship”, and the front page might need updating—all of these are actions that a generalist administrator might perform. Task management for this category of editors is, in the words of one admin, a mix of “templates and experience”. Templates provide automatic listings for categories of articles that need attention (candidate articles for deletion, for example), giving an administrator a launch-point for actionable items. However, one administrator cannot perform the actions related to all template categories at once. Since some administrator tasks have a permanent backlog due to a heavy flow of activity, experience informs editors which category will probably need the most attention, and thus guides them towards which set of action items to tackle first.

Since administrators have no easy way of knowing what other administrators are doing, a shared view into the administrative activity on Wikipedia at any time would seem to be a valuable tool, including the current backlog for several categories, information which is now presented as simple number count of action items. Similarly, a mixed-initiative interaction scenario might be possible, by suggesting the next administrative act to be performed, based on current backlog and past admin behavior.

Another open question with regards to administrators is how to balance their administrative responsibilities with their desire to still contribute content to Wikipedia. Interviewees reported that, upon switching to adminship, they sometimes found they did not have enough time to do the editing and contributing that they had done to get to adminship in the first place. Explicitly looping the administrator-level participants in at times when the voice of experience would be needed – and not just in negative situations, but in positive ones like the effort to move an article to Featured status – might be an effective way of counteracting this imbalance.

Discussion

Based on this account of task management for different levels of participation on Wikipedia, we obtained five principal insights in response.

Bottom-up Structure

Unlike many traditional volunteer organizations, where the product of the volunteering effort is not consumed or enjoyed by the volunteers themselves, Wikipedia editors are also Wikipedia readers. This means that editors are also familiar with what makes a good article, and can thus tag pages that need attention. The “template tags” referred to in the advanced editor description are the method through which pages are flagged for attention. The first insight that emerges is as follows: tasks get assigned – to the community – from individual editors through a bottom-up structure using template tags.

Tagging pages, at first blush, might seem akin to assigning a task. However, there are two fundamental differences: the dates on the template tags are backwards-looking, rather than forwards-looking (“this page has been tagged since August 2007”, rather than, “this page should be cleaned up by September 2008”), and there is no single person or group that has been assigned the task of rectifying whatever problem merited the template tag. As we saw from our observations, this leads to tags that grow stale and are never completed by a participant.

Template tags do not function well as next actions

While a free-form task tagging system is attractive to participants due to the low cost of adding a new tag to a page, it does not function well as a method for organizing and motivating work on Wikipedia. This suggests that existing systems might work well to match users and tagged pages, but that this in itself is not enough to encourage work and accommodate editors' preferences. This is analogous to Whitaker & Sidner's findings that e-mail does not function effectively as a task management platform [29].

Particularly, this is due to the lack of actionability for most template tags. An advanced editor browsing around pages within her area of expertise has probably seen several times that a particular article needs cleanup, but has no way of knowing whether what the steps towards achieving this cleanup in this article might be, and how much of a time commitment this will be.

Lack of triage

As important as the step of collecting and potentially distributing tasks is the action of reviewing the assigned tasks to make sure they have not grown stale, irrelevant, or otherwise useless. While traditional task management in an organization contains this triage stage as an essential step of the process, the bottom-up, self-directed structure of Wikipedia complicates this stage. As template tags accumulate, it is hard for editors to tell data from noise. Since it currently falls to individual editors to make decisions as to whether a page's task tags are old and invalid, the task of triage is extremely burdensome. Thus, a system which opened this process to the community at large and broke down the overall challenge into smaller chunks performed by many people (for example, voting on which aspect of a page most needs cleanup) would both fit the overall Wikipedia *ethos* and increase the likelihood that tasks would be triaged.

Disconnect between Individual and Site Tasks

Many of our respondents reported having a personal to-do list or list of articles that they would like to expand, given the time. Since this list usually remains either in the participants' heads, on their own Wikipedia page (not linked to any site-wide display), or elsewhere, the community cannot benefit from the knowledge of what a topic expert might do, given enough time on that page.

Further, the disconnect between the site's task tracking and a user's own task tracking means that there is no current way

to understand what planned work will happen on Wikipedia. Allowing editors to selectively publish their planned (and current) tasks would create the notion of public commitment and accountability. The literature on offline volunteer organizations often highlights the role of scoped projects for volunteers, coupled with a shared, common display of currently underway tasks [22].

Lack of support for contextual discovery

Finally, our research highlighted the current lack of support for contextual discovery of relevant tasks. If an editor currently desires to edit articles in sequence, their source of related tasks is the WikiProject – though browsing through tasks on projects is currently clunky, and involves paging through several pages of category lists. Further, tasks for a project are not sortable by difficulty / time commitment, or any other facets.

WikiTasks

Inspired and informed by our qualitative research, we have set out to design a tool that accommodates and extends the task practice of Wikipedia, with an eye towards its applicability in the overall space of crowd-powered social software systems. Conceptually, this system is an application of *Getting Things Done*-related ideas and inspiration from open-source bug tracking systems such as Bugzilla, combined with our research and analysis of current task practice of Wikipedia (see Figure 5). For this tool, we focus on the task practice of advanced editors, who are uninterested in simple tasks such as participating in Recent Changes patrol, but instead want to move through the site and contribute as needed.

Design goals

We set out to design WikiTasks with a set of design goals, based on our research into Wikipedia task management.

1. Support seamless transitions between public and private tasks
2. On-page task management for articles for management and fortuitous discovery of tasks
3. Contextual display of tasks for related articles, to support a flow through a set of articles.

Further, we were informed by prior work [5], and fulfill the design guidelines in this work as follows:

The WikiTask system

The WikiTasks system adds a task sidebar to Wikipedia articles, Wikipedia user pages, and WikiProject pages. These were selected as entrypoints due to the different facets into the Wikipedia that each presents: page-level, user-level, and topic-level. At any point, if the editor is browsing using the WikiTasks proxy, the task list will be automatically inserted into these pages. For article pages, the visible tasks are those relating to the current article (that are currently unassigned to an editor, or assigned to an editor who has chosen to make their task list public), and tasks relating to articles within that article’s “neighborhood” (similar articles belonging to

the same WikiProject). For user pages, tasks assigned to that editor are displayed (to all visitors, if the task list is public, or to just the editor). Finally, for WikiProjects, tasks under that project’s jurisdiction are listed, which extends the current ability of projects to have a “to-do list”.

The WikiTasks system breaks down tasks along two dimensions: tasks can be *personal* or *site-wide*, and can have varying levels of time investment / difficulty. The distinction between personal and site-wide preserves the current ability for task creation to be performed by anyone in a bottom-up manner, but adds a private dimension and the ability to transition tasks between both of these spheres through a toggle in the user interface. Further, editors can elect to publish their task list, providing shared visibility into tasks and a sort of public commitment factor. This distinction is inspired by open-source bug tracking system like Bugzilla, which can default to assigning a bug to “anyone”, with the understanding that a particular contributor will eventually accept the bug and commit to fixing it.

The second distinction – levels of time / difficulty – emerged from our conversations with editors, who distinguished between small tasks that were mostly related to consistency, presentation, and grammar, and larger tasks that required research and a longer time commitment. Providing this additional facet allows editors to enter a flow-state through performing a set of smaller tasks, or diving into a larger task.

Our prototype is completely opt-in, so most anonymous or first-time editors of Wikipedia would probably not be exposed to WikiTasks. Since this might lead to a disconnect between the state of the task list and of the actual article, WikiTasks includes the notion of task staleness – if a task list attached to a particular article has not been edited for a substantial period of time, a colored bar at the bottom of the list indicates the staleness of the list.

Finally, WikiTasks provides a bookmarklet that can be installed in an editor’s browser that allows for “clipping” text from the Web to their own personal research area.

Implementation

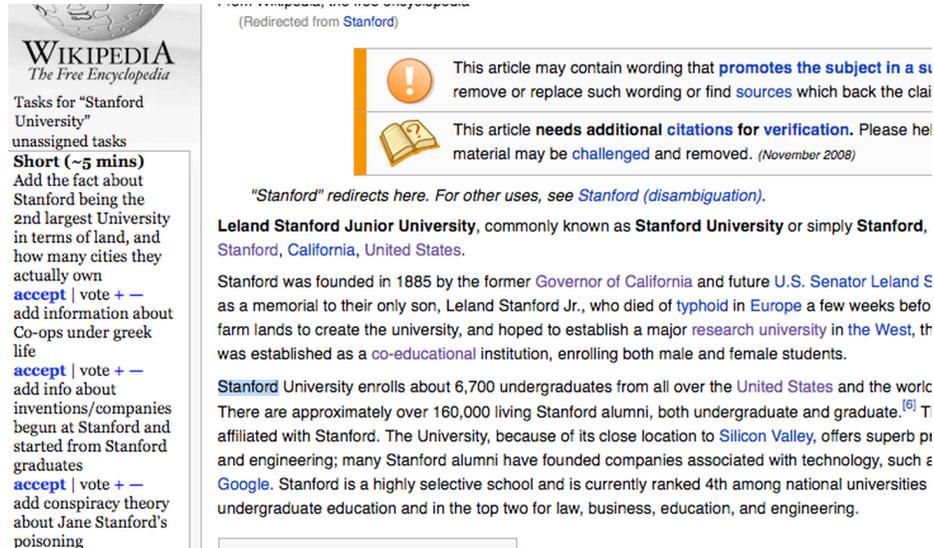
The WikiTasks system is composed of three parts: a rewriting proxy that adds dynamic content to Wikipedia pages (this approach was suggested to us by the authors of prior work in this domain [25]), a Django-based Web service that tracks and stores tasks, and a tool that runs on the Wikimedia-provided Wikipedia toolserver¹ that fetches related and relevant pages.

In our prototype system, when an editor visits a proxied page², the system inserts a small piece of Javascript into the page. This Javascript uses the XMLHttpRequest object to fetch dynamic content for that page from our Web service. This Web service, written using the Python-based Django framework, queries its own data store about site- and user-specific tasks that are relevant to the current page (and the current page’s context), and queries our tool on the Wikimedia tool-

¹<http://toolserver.org>

²<http://WikiTasks-proxy-url/wikipedia/wiki/Dogs>

Figure 5. WikiTasks on an article page



server. This tool returns which pages are related and “neighbor” to the current page or WikiProject, and the Django server queries its data store for relevant tasks for these additional pages, too. Then, this is returned to the user and inserted into a page’s Infobox.

There are a few trade-offs to this implementation. Proxy-ed visitors to the Wikipedia are forbidden by Wikimedia from making edits to the encyclopedia directly from the proxy (to avoid anonymous, untraceable edits), so WikiTasks redirects Edit requests to the Wikipedia page itself. Also, a large-scale deployment of this prototype’s ideas would benefit from integration with the MediaWiki software that powers Wikipedia, rather than serving pages through a proxy request.

Future Work and Conclusion

WikiTasks is an instantiation of broader ideas that are applicable to other social Web sites concerned with collaboration and production. Namely, the idea of contextual display of task lists according to the section of the site being browsed (showing both tasks relating to the current section, as well as related sections), the bottom-up creation of tasks that can be initially entered into the system but later taken up by individual contributors, and visibility into how long it has been since a participant has had a look at the list of tasks.

Though our task tool is grounded in previous research on task list design, it would also benefit from a wide deployment on Wikipedia, including eventual integration with the MediaWiki platform so that it could be used on any MediaWiki-based wiki. Further, as we explored with the community triage component, the social aspect of Wikipedia can open new possibilities for task management that might be explored in future work.

Our insights, requirements analysis, and design work have implications that are Wikipedia-specific, but have generalizable applications as well. Specifically, we believe the bottom-

up structure of our task management system, the tight coupling of tasks and the content they relate to, and the attempt to seamlessly intermingle personal and community-wide tasks have applications to Wiki-style collaboration more broadly, social software on the Web that focuses on user-generated content, and with other large-scale collaborations, such as open-source software collaborations.

This paper has presented an in-depth qualitative look at task practice on the English Wikipedia, with an expanded list of roles on Wikipedia and a description of how the core segment of roles (beginning editors, advanced editors, and administrators) create, select, and perform tasks. We found that, despite the ease of creating tasks in a bottom-up manner on Wikipedia, actual task selection and performance is done on an individual level that often ignores the template tasks. Further, the existing structure does not support the discovery of actions as editors move throughout the Wikipedia. From these observations and insights, we developed WikiTasks, a tool that integrates with Wikipedia and allows for site-wide tasks to be created in a bottom-up manner, but also personal tasks to be accepted from this larger pool of community tasks.

Acknowledgements

We thank Sep Kamvar for invaluable project advice and the Wikipedians who graciously gave us their time and ideas. This work was supported by seed funding from the Stanford Computer Science Department.

REFERENCES

1. Adler, B., and Alfaro, L. A content-driven reputation system for the wikipedia. *WWW '07: Proceedings of the 16th international conference on World Wide Web* (May 2007).
2. Allen, D. *Getting Things Done: The Art of Stress-Free Productivity*. Penguin Books, 2003.

3. Almeida, R., Mozafari, B., and Cho, J. On the evolution of wikipedia. *Proc. of ICWSM* (Dec 2007).
4. Amazon. Mechanical turk [online]. Available from: <http://mturk.com>.
5. Bellotti, V., Dalal, B., Good, N., Flynn, P., and Bobrow, D. What a to-do: studies of task management towards the design of a personal task list manager. *Proceedings of the SIGCHI conference on Human factors in computing* (Dec 2004).
6. Bellotti, V., Ducheneaut, N., Howard, M., and Smith, I. Taking email to task: the design and evaluation of a task management centered email tool. *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems* (Apr 2003).
7. Benkler, Y. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press, 2006.
8. Beschastnikh, I., Kriplean, T., and McDonald, D. Wikipedian self-governance in action: Motivating the policy lens. *Association for the Advancement of Artificial Intelligence* (2008).
9. Bryant, S., Forte, A., and Bruckman, A. Becoming wikipedian: transformation of participation in a collaborative online encyclopedia. *GROUP '05: Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work* (Nov 2005).
10. Buriol, L., Castillo, C., Donato, D., and Leonardi, S. Temporal analysis of the wikigraph. *Proc. of Web Intelligence* (Dec 2006).
11. Community, W. Wikipedia [online]. Available from: <http://wikipedia.org>.
12. Cosley, D., Frankowski, D., Terveen, L., and Riedl, J. Using intelligent task routing and contribution review to help communities build artifacts of lasting value. *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems* (Apr 2006).
13. Cosley, D., Frankowski, D., Terveen, L., and Riedl, J. Suggestbot: using intelligent task routing to help people find work in wikipedia. *IUI '07: Proceedings of the 12th international conference on Intelligent user interfaces* (Jan 2007).
14. Forte, A., and Bruckman, A. Why do people write for wikipedia? incentives to contribute to open-content publishing.
15. Gabilovich, E., and Markovitch, S. Computing semantic relatedness using wikipedia-based explicit semantic analysis. *Proceedings of the 20th International Joint Conference on ...* (Dec 2007).
16. Kittur, A., Chi, E., Pendleton, B., Suh, B., and Mytkowicz, T. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. *CHI '07: Proceeding of the twenty-fifth annual SIGCHI conference on Human factors in computing systems* (2007).
17. Kleek, M., Bernstein, M., Karger, D., and mc schraefel. Gui — phooey!: the case for text input. *UIST '07: Proceedings of the 20th annual ACM symposium on User interface software and technology* (Oct 2007).
18. Kreifelts, T., Hinrichs, E., and Woetzel, G. Sharing to-do lists with a distributed task manager. *ECSCW'93: Proceedings of the third conference on European Conference on Computer-Supported Cooperative Work* (Sep 1993).
19. Kriplean, T., Beschastnikh, I., and McDonald, D. Articulations of wikiwork: uncovering valued work in wikipedia through barnstars. In *Proceedings of the ACM 2008 conference on Computer supported cooperative work* (2008), ACM New York, NY, USA, pp. 47–56.
20. Kuznetsov, S. Motivations of contributors to wikipedia. *ACM SIGCAS Computers and Society* (Jan 2006).
21. Magnus. Wiki todo [online]. Available from: <http://toolserver.org/~magnus/wikitodo.php>.
22. Pearce, J. Volunteers: The organizational behavior of unpaid workers. *books.google.com* (Jan 1993).
23. Riehle, D. How and why wikipedia works: an interview with angela beesley, elisabeth bauer, and kizu naoko. *WikiSym '06: Proceedings of the 2006 international symposium on Wikis* (Aug 2006).
24. Stalder, F., and Hirsh, J. Open source intelligence. *First Monday* 7, 6 (2002).
25. Suh, B., Chi, E., Kittur, A., and Pendleton, B. Lifting the veil: improving accountability and social transparency in wikipedia with wikidashboard. *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems* (Apr 2008).
26. Viégas, F., Wattenberg, M., and Dave, K. Studying cooperation and conflict between authors with history flow visualizations. *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems* (Apr 2004).
27. Viégas, F., Wattenberg, M., and McKeon, M. M. The hidden order of wikipedia. *HCI* (2007).
28. Wattenberg, M., Viegas, F., and Hollenbach, K. Visualizing activity on wikipedia with chromograms. *LECTURE NOTES IN COMPUTER SCIENCE* (Dec 2007).
29. Whittaker, S., and Sidner, C. Email overload: exploring personal information management of email. *CHI '96: Proceedings of the SIGCHI conference on Human factors in computing systems: common ground* (Apr 1996).