

Mentoring in Wikipedia: A Clash of Cultures

David R. Musicant
Dept. of Computer Science
Carleton College
Northfield, MN 55057
dmusican@carleton.edu

James A. Johnson
Dept. of Anthropology
University of Minnesota
Minneapolis, MN 55455
joh04712@umn.edu

Yuqing Ren
Carlson School of
Management
University of Minnesota
Minneapolis, MN 55455
chingren@umn.edu

John Riedl
Dept. of Computer Science
University of Minnesota
Minneapolis, MN 55455
riedl@cs.umn.edu

ABSTRACT

The continuous success of Wikipedia depends upon its capability to recruit and engage new editors, especially those with new knowledge and perspectives. Yet Wikipedia over the years has become a complicated bureaucracy that may be difficult for newcomers to navigate. Mentoring is a practice that has been widely used in offline organizations to help new members adjust to their roles. In this paper, we draw insights from the offline mentoring literature to analyze mentoring practices in Wikipedia and how they influence editor behaviors. Our quantitative analysis of the Adopt-a-user program shows mixed success of the program. Communication between adopters and adoptees is correlated with the amount of article editing done by adoptees shortly after adoption. Our qualitative analysis of the communication between adopters and adoptees suggests that several key functions of mentoring are missing or not fulfilled consistently. Most adopters focus on establishing their legitimacy rather than acting proactively to guide, protect, and support the long-term growth of adoptees. We conclude with recommendations of how Wikipedia mentoring programs can evolve to take advantage of offline best practices.

Categories and Subject Descriptors

H.1.2 [Models and Principles]: User/Machine Systems

General Terms

Human Factors

Keywords

Wikipedia, newcomer retention, mentoring

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 republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WikiSym '11 October 3–5, 2011, Mountain View, CA, USA.
Copyright 2011 ACM 978-1-4503-0909-7/11/10 ...\$10.00.

1. INTRODUCTION

Recruiting and retaining new volunteers in Wikipedia is a challenge. Until recently, research that examined the growth of Wikipedia showed it be growing at an exponential pace [1]. Currently, however, the rate at which new articles and new users (referred to as “editors”) come to Wikipedia seems to have slowed dramatically [10]. According to Wikipedia’s own statistics page¹, the number of Wikipedians (editors who have done at least 10 edits over their Wikipedia lifetimes) continues to increase by about 1% a month. However, this rate is clearly decreasing; in 2005, for example, the rate was approximately 10% a month. Furthermore, the number of editors who contribute more than 5 or more than 10 edits each month is decreasing with time. It appears that an important inflection point has been reached.

Perhaps even more troubling is the difficulty that new editors have in engaging with the community. Before 2005, nearly 40% of new editors in English Wikipedia were active one year after their first edit; after 2007, only 12-15% of new editors were active one year later.² These same trends were seen in varying degrees across a number of Wikipedias in other languages. Wikipedia “newbies” are having considerably more trouble integrating with and becoming a part of the community than they did in years past. This means that groups of people that are currently underrepresented in Wikipedia [2] will find it even harder to be heard going forward, as it is becoming so difficult for people to socialize into the existing system.

Why does this matter? The reality, of course, is that there is still considerable knowledge missing from Wikipedia. Many Wikipedia articles are stubs, which are articles too short to provide more than rudimentary information.³ New knowledge continues to be created which needs documentation in Wikipedia. Existing editors may have already filled in the information they know about, so new editors may be needed to fill in the gaps. The underrepresentation of certain populations means there are many articles that belong in Wikipedia that are likely covered less well than they could be [2]. Therefore, in this current stage of Wikipedia’s growth, it is more important than ever to implement tech-

¹<http://stats.wikimedia.org/EN/TablesWikipediaEN.htm>

²<http://strategy.wikimedia.org/w/index.php?oldid=80191>

³<http://en.wikipedia.org/wiki/Wikipedia:Stub>

niques for retaining and inspiring new editors if Wikipedia is to continue to grow and thrive. The recent Editor Trends Study⁴ shows that new editors are coming to Wikipedia, so there is a currently a pool of people from which to recruit.

There are a variety of programs within Wikipedia that have been created, in whole or in part, to address the challenge of newcomer retention.⁵ Despite the existence of these programs, however, Wikipedia is still struggling. One of the key ways that “real-world” communities develop their volunteers is through mentors. Only a few of the newcomer programs in Wikipedia utilize an approach resembling mentoring (Adopt-a-user⁶ and Mentorenprogramm⁷ are two examples) with mixed levels of success. The Wikimedia Foundation has expressed concern with newcomer retention issues, and has very recently started another program with mentoring-like aspects.⁸ In this paper, we examine what social science research literature tells us about best practices for mentoring in traditional real-world workplaces. We point out that many of these best practices and the culture of Wikipedia are in direct conflict, and examine in detail the nature of this conflict. We then analyze in depth one particular mentoring program, namely the English Wikipedia “Adopt-a-user” program⁹. Finally, we look towards the future and discuss how a Wikipedia mentoring program might be able to synthesize real-world best practices and lessons learned from current mentoring implementations.

2. WHAT MENTORING IS, AND HOW IT OUGHT TO WORK

Mentoring has been described as “an intense interpersonal exchange between a senior experienced colleague (mentor) and a less experienced junior colleague (protégé) in which the mentor provides support, direction, and feedback regarding career plans and personal development.” [9] In offline volunteer and professional communities, assigning mentors to members has proven effective at helping them move through new roles and challenges [3]. New members find that mentors help them acclimate to the social and organizational norms of the community. Experienced members find that mentors may help them grow as leaders of the community. Furthermore, effective mentoring can result in benefits for the protégé such as faster promotion rates, higher compensation, accelerated career mobility, higher career satisfaction and self-esteem, and reduced stress and role conflict [9]. It seems reasonable that some of these benefits should transfer to Wikipedia as well. In this section, we examine what qualities are often distinguished as the key elements of successful mentoring.

The literature identifies two major functions of mentoring: *career functions* and *psychosocial functions*. Career functions are those “that enhance advancement in an organization” [6], and can be broken down into the following five sub-areas:

- *Sponsorship*. A mentor can actively nominate a protégé for desirable career moves [6], and can encourage the pro-

tégé to take on tasks that would specifically help to prepare for promotion [8].

- *Exposure-and-Visibility*. A mentor can assign responsibilities or encourage the protégé to take on responsibilities that enhance the protégé’s visibility to people higher up in the organization [6]. Specifically, the mentor can introduce the protégé to new colleagues, or directly provide assignments designed to help increase contact with others [8].
- *Coaching*. A mentor can suggest “...specific strategies for accomplishing work objectives, for achieving recognition, and for achieving career aspirations.” [6] A mentor can share the history of his/her career, can encourage the protégé to prepare for advancement, and can suggest specific strategies for achieving specific goals [8].
- *Protection*. A mentor can help to shield a protégé from damaging contact with more senior workers, and can intervene in situations where the protégé does not yet have the skills to resolve the situation appropriately [6]. Likewise, the mentor can work to preemptively reduce unnecessary risks that could threaten promotion for the protégé [8].
- *Challenging Assignments*. A mentor can assign or encourage the protégé to take on challenging work, supported with appropriate training and feedback. This helps to develop the skills of the protégé, and also to help the protégé build a sense of accomplishment [6, 8].

Psychosocial functions, alternatively, are “those aspects of a relationship that enhance an individual’s sense of competence, identity, and effectiveness in a professional role” [6]. These can be broken down into:

- *Role modeling*. A mentor serves as an example to the protégé regarding attitudes, values, and behavior. The mentor may consciously attempt to provide such examples, or may merely do so unintentionally by doing the job at hand [6]. Shared attitudes and values between the mentor and protégé likely enhance the relationship [8].
- *Acceptance-and-Confirmation*. Both the mentor and the protégé can enhance each other’s sense of self via positive regard and compliments [6]. The mentor can also ask the protégé for help with a problem that the mentor has [8], which helps the protégé feel more competent and effective.
- *Counseling*. The protégé can discuss with the mentor personal matters of concern [6], such as questions of competence and relationships with colleagues that should be held confidential [8].
- *Friendship*. Social interaction and exchange about work and outside experiences can help enhance work by protégés, particularly during the early and middle stages of a protégé’s career [6]. Going to lunch together and interacting outside of work are two possible ways of developing friendship [8].

3. HOW MENTORING INTERSECTS WITH WIKIPEDIA

Most of the traditional mentoring functions listed in the previous section have obvious Wikipedia interpretations. A few of them, however, do not map quite so well. Wikipedia does not have a clear concept of promotion. There are roles

⁴<http://strategy.wikimedia.org/w/index.php?oldid=79956>

⁵http://en.wikipedia.org/wiki/WP:Community_portal

⁶<http://en.wikipedia.org/wiki/Adopt-a-User>

⁷<http://de.wikipedia.org/wiki/WP:Mentorenprogramm>

⁸http://en.wikipedia.org/wiki/WP:Wiki_Guides

⁹<http://en.wikipedia.org/wiki/Adopt-a-User>

such as administrators, bureaucrats, and a few others, but very few editors attain the roles. Without a more clear-cut eventual “expected promotion,” much of *sponsorship* loses its relevance. While a mentor certainly could recommend tasks to a protégé that would help the protégé learn to do a better job, this becomes very similar to *challenging assignments*. *Exposure-and-visibility* ends up being highly related to these two as well. In Wikipedia, helping a protégé identify tasks to work on that would increase that protégé’s capabilities would nearly automatically help the protégé learn more about Wikipedia and help other Wikipedians become familiar with the protégé’s work.

Apart from the above issue of mapping traditional mentoring functions to the organization and goals of Wikipedia, these mentoring functions also face challenges and opportunities related to Wikipedia’s culture and technology. Some characteristics of Wikipedia allow the above mentoring functions to easily occur; there are other characteristics that are hostile to some of these practices.

Challenges from working in wikis

MediaWiki talk pages are organized poorly for communication. It is unfortunate from this perspective that the MediaWiki software is optimized for wiki work, and not for interpersonal communication. Most communication between editors occurs on public talk pages, which are simply more wiki pages within Wikipedia. Indenting and threading are not automatically managed, as they are in most web forums. (The Liquid Threads¹⁰ extension adds threading support for talk pages, but has not caught on in general Wikipedia use.)

It is unclear where communication between two editors should occur. Wikipedia culture does not specify where the correct place is for two editors to communicate. All Wikipedia editors have their own personalized talk pages. When two editors communicate, there is little consistency as to whether editors post entirely on the talk page of one of the two users, or whether they cross-post responses to each other’s pages.

Challenges from working online

Wikipedia does not cleanly “push” communications. Editors must return to Wikipedia to know that they have received a communication from another editor. It is possible to set up RSS feeds of changes to talk pages, but this mechanism is quite clumsy, and it is unlikely that new editors would make the effort. The more fundamental issue here is that most new editors have a low level of presence and awareness in Wikipedia when compared to the real world or perhaps other online environments (such as Facebook). To help counter this problem, most other online environments offer a simple and reliable push to email, which increases awareness.

Communication is public. Wikipedia culture seems to encourage all communication to occur on talk pages, as opposed to in other non-Wikipedia forms of communication. This cultural norm is due, in part, to the fact that many editors are anonymous and thus not able to be contacted in any other way. As all content posted to Wikipedia pages is both publicly viewable and archived, there is no expectation of privacy in communication between editors.

¹⁰<http://www.mediawiki.org/wiki/Extension:LiquidThreads>

Challenges from Wikipedia culture

Socializing is discouraged. In both policy and culture, Wikipedia is opposed to too much socializing among its editors. This is perhaps best exemplified via the Wikipedia policy page “What Wikipedia is not,”¹¹ particularly the section titled “Wikipedia is not a blog, webspace provider, social network, or memorial site.” Specifically, it says that “Wikipedians have their own user pages, but they may be used only to present information relevant to working on the encyclopedia... The focus of user pages should not be social networking, or amusement, but rather providing a foundation for effective collaboration.” Another policy, which merely says “Do not ask for another’s personal details,”¹² would similarly seem to discourage socialization.

Policy and culture encourage communications to be brief. “Be concise” is listed as a best practice for talk page communication: “If your post is longer than 100 words, consider shortening it.”¹³ Precisely whether or not this is a challenge is somewhat unclear. Some other online forms of communication have this sort of restriction; Twitter would be an obvious extreme example.

Opportunities afforded by Wikipedia

Editors are encouraged to be civil to each other. A considerable number of policies and guidelines make this point in different ways, including the Wikipedia pillar on Civility¹⁴, as well as various guidelines such as “*Please do not bite the newcomers.*”¹⁵ and “*Assume good faith.*”¹⁶ Few other online communities likely have such direct and verbose descriptions of the value of civil behavior to the community.

Editing history is easily retrievable. For every editor in Wikipedia, there is a “User contributions” link that can show chronologically all editing work done by that editor. Many online environments have a variety of ways of logging and displaying (or keeping private) the activity of its users. Wikipedia is fairly unique in that any user can quickly see a complete log of all edits done by another user.

The final two characteristics above facilitate mentoring; the remaining characteristics would seem to oppose the successful implementation of at least some of the mentoring functions listed in Section 2. In order to get a better sense of how these concepts have played out in practice, we move to looking at how mentoring in Wikipedia has been actually implemented.

4. MENTORING IN WIKIPEDIA, IN PRACTICE

English Wikipedia has a considerable set of resources for assisting new editors, though few of them focus in their design on long-term relationships or connections between editors. The Help Desk¹⁷ and #wikipedia-en-help IRC channel¹⁸ are likely the two most well-known examples. These

¹¹<http://en.wikipedia.org/wiki/WP:NOT>

¹²<http://en.wikipedia.org/w/index.php?oldid=420792073>

¹³<http://en.wikipedia.org/w/index.php?oldid=420792073>

¹⁴<http://en.wikipedia.org/wiki/WP:CIV>

¹⁵<http://en.wikipedia.org/wiki/WP:DONTBITE>

¹⁶<http://en.wikipedia.org/wiki/WP:AGF>

¹⁷<http://en.wikipedia.org/wiki/WP:HD>

¹⁸<http://en.wikipedia.org/wiki/WP:IRC>

programs and others like them do not resemble mentoring programs, however, and so we do not assess their popularity or effectiveness in this paper. There is evidence showing that WikiProjects have been an effective mechanism to direct editors' attention to challenging tasks, though these projects do not pair mentors with protégés [5]. An alternative set of "mentoring" efforts in Wikipedia exist regarding the rehabilitation of previously banned editors.^{19,20} It appears that the title "mentor" may be misapplied for these programs, as the role of the mentor here was really more of a parole officer.²¹ Both of these programs essentially died due to inactivity in less than a year.

It appears that essentially none of the newcomer programs in English Wikipedia (including those listed above and others) resemble a true mentoring program, with one exception: Adopt-a-user²². Adopt-a-user has a specific focus on setting up an "adoption" relationship (i.e., a mentoring one) between an experienced Wikipedia editor (the "adopter") and one who is looking for assistance (the "adoptee"). Therefore, we present a detailed analysis of the Adopt-a-user program to assess its effectiveness from two perspectives: how well does it embody the functions of a successful mentoring program, and what effect does the program have on those editors who become adoptees?

We now proceed to discuss Adopt-a-user in more detail.

5. ADOPT-A-USER STRUCTURE

The Adopt-a-user program was created in September, 2006. Since then, over 1000 Wikipedia editors have been adopted by other Wikipedia editors. We review here the mechanics of the program.

Adopt-a-user has two specific kinds of participants: the adopter (i.e., the mentor) and the adoptee (i.e., the protégé). For simplicity of language, we will refer to an editor who wishes adoption as an *adoption seeker*, or simply a *seeker*. Once an editor decides to seek adoption, the program offers two main paths. One path requires the seeker to add an `{{adoptme}}` template to his/her user page. This template displays an image on that seeker's user page, indicating the desire for adoption; additionally, it dynamically adds the seeker's username to a webpage containing a list of editors seeking adoption. The seeker then waits for an adopter to browse that list, find the seeker's name, and contact the seeker directly via the seeker's talk page.

An alternative route is for the seeker to directly contact a potential adopter and ask to be adopted via making a direct request on the adopter's talk page. There are two different ways in which adopters advertise that they are willing to take on adoptees. The first is by placing an `{{adopting}}` template on their user page. This displays an image on their user page indicating that they are adopting, but also dynamically adds the adopter's username to a page containing a list of adopting editors. The second way an adopter can advertise is to directly edit the "Adopters" page²³, which vaguely resembles a listing of "classified ads"; adopters specify their names, their interests, additional information about themselves, and whether or not they are currently adopting.

These varying approaches are not mutually exclusive; a seeker can advertise via the "adoptme" approach while simultaneously looking for an adopter on either or both of the pages listing adopters. Likewise, adoptees who connect with their adopters through a non-traditional approach (such as working together on an article) may skip this connection phase of the adopting process entirely.

Once a seeker and a potential adopter have communicated, the adopter may choose to offer adoption to the seeker. This offer is typically done via conversation on the seeker's and adopter's talk pages. In principle, the adopter is supposed to leave an `{{adoptoffer}}` template on the seeker's user page, which "officially" indicates that the offer has been made. In practice, however, the `{{adoptoffer}}` template is not always used. (Some statistics on this can be found in the next section.)

Once the adoptee decides to accept the adoption offer, the adoptee and adopter each are supposed to place `{{adoptee}}` and `{{adopter}}` templates on their user pages, which each display images indicating the relationship that has been established.

Once the adoption relationship has been established, it exists as long as the two parties wish. The most common way that the relationship ends is simply for the adopter and adoptee to stop talking. Typically, the templates seem to stay up on the editors' pages beyond the time the relationship has ceased; eventually, the templates are removed as part of some sort of cleanup operation. A less frequent way of ending the relationship is for the adoptee and adopter to agree that the adoptee has advanced to the point of no longer needing assistance, at which point the adoptee "graduates" from the program. An `{{AdopteeGrad}}` template exists for this purpose, but it is rarely used.

Adopt-a-user, incidentally, is a classic example of how a Wikipedia page or program can gain a life of its own beyond its creator. The initial creator of the program was the Wikipedia editor Flameviper, who was 12 years old at the time. Five months later, Flameviper was blocked from Wikipedia for "sock puppetry."²⁴ Flameviper has since returned to Wikipedia with the username Ziggy Sawdust, and seems to have won back support of the community²⁵. However, he is no longer involved with Adopt-a-user.

6. ADOPT-A-USER PARTICIPATION STATISTICS

Before we move on to examine the effectiveness of Adopt-a-user as a mentoring program, we first review participation levels in the program. Nearly all of our data for the Adopt-a-user program is based on the most recent full-text dump of English Wikipedia, which contains data through January 2010.

We note that we have made a number of simplifying assumptions for purposes of our analysis. `{{adoptme}}` templates can appear, disappear, and reappear repeatedly on a single editor's user page. We simply use the date of the first appearance as the date that the editor requested adoption. A relatively small number of editors have multiple adopters, possibly at disjoint periods in time or possibly with some overlap. For simplicity, we have eliminated such editors from our analysis. When an editor is adopted, there should

¹⁹<http://en.wikipedia.org/wiki/WP:REHAB>

²⁰<http://en.wikipedia.org/wiki/WP:MENTCOM>

²¹<http://en.wikipedia.org/w/index.php?oldid=294488170>

²²<http://en.wikipedia.org/wiki/WP:Adopt-a-user>

²³<http://en.wikipedia.org/w/index.php?oldid=421602681>

²⁴<http://en.wikipedia.org/wiki/WP:SOCK>

²⁵<http://en.wikipedia.org/w/index.php?oldid=374505883>




Figure 1: Participation in Adopt-a-user.

be symmetric `{{adoptee}}` and `{{adopter}}` templates on the adoptee and adopter user pages, respectively. In reality, this is not always the case. Therefore, we take the adoptee’s perspective, and consider someone to have been adopted if there is an `{{adoptee}}` template on that editor’s user page.

The Venn diagram in Figure 1 shows total numbers of editors from the start of the program through January 2010, as they work through the adoption process. The “Adoptme” circle on the upper left represent all editors that at some point in history had an `{{adoptme}}` template on their user pages. Similarly, the “Adoptoffers” circle represents all editors that at some point had an `{{adoptoffer}}` template. Note that the “Adoptoffers” circle does not represent all editors that received an offer for adoption. For a number of them, offers come through conversation on talk pages, and are not actually coded via the `{{adoptoffer}}` template. Finally, the “Adoptees” circle represents those editors with an `{{adoptee}}` template after eliminating those who had more than one adopter.

Figure 1 shows that a total of 2042 editors requested adoption via the `{{adoptme}}` mechanism, yet only 749 of them (417 + 332) actually received an offer. However, it actually turns out that during the period of measurement, most editors acting in good faith who wanted to be adopted were made offers or were adopted outright. We randomly sampled 14 editors who were not made offers and discovered that there was a good explanation for most of them. Five of them were vandals, sock puppets, or other sorts of troublemakers; four actually were adopted but they coded the template incorrectly; and three more did receive offers via conversation, but not via the `{{adoptoffer}}` template.

Figure 1 also shows that 1095 editors received adoption offers via the `{{adoptoffer}}` template. It is fairly clear that the `{{adoptoffer}}` mechanism is typically only used when an `{{adoptme}}` has been posted; while many editors seek adoption through contacting a potential adopter directly, only 51 `{{adoptoffer}}` templates were placed (41 + 10) when an `{{adoptme}}` had not been placed previously. One number to note in particular is the 712 seekers who posted an `{{adoptme}}` template and received an `{{adoptoffer}}`, yet never “closed the loop” and got adopted. A random sample of 13 showed that many of them simply continued to edit Wikipedia and ignored the offers. Of those 13, one was a sock puppet, and eight more simply continued to edit Wikipedia and ignored the offers. Four of them vanished from Wikipedia before the offers came, and so perhaps for these editors, a quicker offer might have made a difference.




Figure 2: Count of “adoptmes” and “adoptees” by month.

Figure 2 shows the frequencies of seekers and adoptees over time. The number of people participating in the program has been decreasing, overall.

The stated intent of the Adopt-a-user program is to help out new editors.²⁶ The median number of days an editor has been in Wikipedia, when adopted, is 41 days. However, 23% of all editors who request adoption via the `{{adoptme}}` template have been editing in Wikipedia for at least six months. Similarly (data not shown here), 24.5% of adoptees have been editing in Wikipedia for at least six months when they are actually adopted.

For those editors who ask to be adopted via an `{{adoptme}}` template, how long does it take to get adopted? The median wait time is four days. Over half are adopted within the first week, and most are within the first month. The distribution does have a long tail, and waits did take many months for a smaller number of editors. Moreover, the delay for editors to be adopted grew considerably during the summer and fall of 2010.²⁷

Now that we have reviewed the scope of the Adopt-a-user program and the participation levels within, we consider the matter of the experience that program participants received.

7. ADOPT-A-USER QUALITATIVE ANALYSIS

The first research question we aim to address is the following:

RQ1: *To what extent does the Adopt-a-user program fulfill the key functions identified in the mentoring literature?*

In order to answer this question, we undertook a manual analysis of communications between adopters and adoptees in order to better understand how frequently each of these mentoring functions occurred with Adopt-a-user. Specifically, we randomly sampled 25 adoption pairs (adoptee and adopter). One coder then read all communications between each of these pairs, and coded each communication as one of the prominent mentoring functions.

²⁶<http://en.wikipedia.org/wiki/WP:AAU>

²⁷<http://en.wikipedia.org/w/index.php?oldid=418970618>

After some initial consideration of the mentoring functions in Section 2, we realized that not all of the mentoring functions could be cleanly coded. As explained in Section 3, *sponsorship*, *exposure-and-visibility*, and *challenging assignments* are quite difficult to distinguish within Wikipedia. We therefore chose to code these three mentoring functions together as one, which we refer to as *exposure-and-visibility*. Though we re-use the phrase from the mentoring literature, we choose to use the word “exposure” to have two different meanings: to expose the adoptee to other Wikipedians and vice-versa (as is intended in the mentoring literature), but we also use “exposure” to mean exposing the adoptee to different and enhanced tasks within Wikipedia.

We also discovered (after an initial round of coding) that the most instances of the mentoring function *role modeling* were found in the initial adoption offer and conversation. An adopter will often reference his/her own experiences during the initial communication, and so this sort of role modeling was really more a form of self-promotion or a way of demonstrating expertise. Such forms of role modeling are worth noting, but not quite in the spirit of what the mentoring literature intended. Therefore, we split the mentoring function of role modeling into two separate measurements: “Role modeling/introduction,” which clearly takes place in the context of the adopter and the adoptee meeting each other, and “Role modeling/task,” which takes place at a later stage when the adopter and adoptee are working together on a particular task.

Table 1 shows the results of our coding; several patterns emerge that are worth noting.

1. *The most frequent interaction occurred around role modeling in the introduction, though role modeling beyond the introduction was rare.* Introductory role modeling is different from what workplace mentoring literature describes. This makes sense, however; it is actually quite hard for someone to introduce oneself *without* in some way referring to one’s own experiences. For example, one adopter states: “I recently entered the program, and if you want I can adopt you. I’ve been here for three months, have 4000+ edits.” We note that there are other potential forms of role modeling beyond direct communication: a motivated adoptee could make use of the “User contributions” tool on his/her adopter, which would allow the adoptee to see all of the work undertaken by the adopter. However, an adoptee would not see much of this detail in Wikipedia without making an intentional effort to do so, and it is perhaps unlikely that an adoptee would make this effort without being directed to do so. Similarly, an adopter performs other forms of role modeling that we do not directly measure; the adoptee sees communications between the adopter and other individuals on the adopter’s talk pages, for example. It does seem clear, however, that adopters are at least not actively providing role modeling much beyond the initial communication.

2. *Acceptance-and-confirmation and coaching are the next most common mentoring functions seen.* It seems reasonable that acceptance-and-confirmation would occur somewhat. There are some limitations on what can be said due to the lack of privacy in communication, and it not always clear if a communication will be received by the other party; nonetheless, it seems that both the adopter and adoptee can communicate with each other in a way that shows appreciation for each other’s efforts. Coaching could be easily achieved in the question-and-answer model that we see many

ID	Career			Psychosocial				
	E/V	CH	P	RM/I	RM/T	A/C	CN	F
1	1	5		2	2	1		
2		12	3	6	2	11	1	7
3								
4		4	3	1				2
5		3		1				1
6				1		2		
7				1				
8				1				
9	5	2		2		3		
10		3		1		4		
11				1		1		
12	1	13	1	3		22	11	4
13		3		3	1	2	2	7
14								
15	9	7	1	1		10		2
16				1				
17		4		1		5	4	
18		1		1		1	1	
19						2	2	
20		6	1	2		3		1
21	1	7	1	2		6	10	4
22		15	2	1		4	3	4
23								
24				1				1
25								
≠ 0	5	14	7	20	3	15	8	10

Table 1: Coding of 25 random adopter/adoptee pairs. ID=pair ID; E/V=Exposure-and-Visibility; CH=Coaching; P=Protection; RM/I=Role modeling/introduction; RM/T=Role modeling/task; A/C=Acceptance/Confirmation; CN=Counseling; F=Friendship. An empty row means that no communications were found satisfying any of the mentoring functions listed. The bottom row contains the number of nonzero items in each column.

adoptees/adopters assuming. An adoptee asks a question, and the adopter coaches the adoptee in some way by answering the question.

3. *Counseling and friendship are more problematic, occurring in less than half of the adoption pairings that we sampled.* There are some good reasons for this: both of these would seem to be more challenging to do in Wikipedia than in a face-to-face environment. The difficulties in organizing communications (both on a talk page and across talk pages), the emphasis placed on making conversations brief, the unreliability of receipt of a message, and the lack of privacy all are counter-productive for achieving success at these two mentoring functions. Friendship, in particular, may be further hindered by the Wikipedia policies against socializing. While it is clear that friendship *does* occur, the cultures and policies at Wikipedia make this somewhat challenging. This culture against socialization was perhaps most notably observed during the deletion process for Esperanza, a now-defunct Wikipedia destination intended partially as a help resource for editors and partially as a place for social communication.²⁸ The entire set of pages was nominated and approved for deletion, in part because a number of editors perceived it as a social networking destination that was distracting editors from editing the encyclopedia.²⁹

4. *The mentoring functions of exposure-and-visibility and protection seem to suffer the most* (in addition to role model-

²⁸<http://en.wikipedia.org/wiki/wp:Esperanza>

²⁹<http://en.wikipedia.org/w/index.php?oldid=348450461>

ing/tasks, as described above). While these mentoring functions do occur in some cases, Adopt-a-user is failing at bringing them to fruition. It also seems as though these activities are quite doable in Wikipedia. A good adopter should be able to help an adoptee find interesting and challenging assignments to work on, such as new articles to write or technically challenging reorganizations. An adopter could certainly introduce an adoptee to other editors via talk pages, and could encourage the adoptee to work in areas likely to increase the adoptee's visibility in areas of interest. Protection could be achieved by the adopter advising the adoptee what sorts of pages are politically safe for a new editor to work on, or by the adopter participating in difficult discussions with or about the protégé. For example, if the adoptee is the subject of an investigation or proposed block, the mentor can participate in those discussions to explain the adoptee's activities. Likewise, the mentor can advise the adoptee as to what sorts of activities are likely to be dangerous or potential policy violations. (We note that it is possible that a few of these protection communications may have occurred elsewhere than on adopter/adoptee talk pages, and so we may have missed them in our coding; nonetheless, we would expect to see significantly more discussion of them on talk pages if this function was being carried out more frequently.)

In summary, it appears that Adopt-a-user is managing to carry out some of the mentoring functions proposed by the literature, but it is systematically failing to carry out others. The good news is that exposure-and-visibility, protection, and role modeling/tasks may be “low-hanging fruit.” The key factor that distinguishes these mentoring functions from the others is proactivity. While these might on some occasions occur in response to a question from an adoptee, they much more demand a proactive effort on the part of the adopter. In our reading of the communications between adopters and adoptees, few (but not none) of the adopters take a leadership role. Rather, they take a responsive one, dealing with questions from adoptees as they come in. A change in program structure that encouraged adopters to think about some of these proactive mentoring functions could help the implementation of these significantly. We say more about this in Section 10. In the next section, we look more directly at editing behavior by adoptees.

8. ADOPT-A-USER QUANTITATIVE ANALYSIS

The second key research question we hope to address is:

RQ2: *How does being an adoptee in the Adopt-a-user program affect editing behavior (if it affects it at all)?*

Mentoring functions are carried out primarily through communication. Therefore, we examine the effect of communication on adoptee's edits. Specifically, we look at:

- *Amount of communication between the adopter and the adoptee.* We measure this by counting the number of edits made by an adopter on one of the adoptee's talk pages, and the number of edits made by the adoptee on one of the adopter's talk pages. One of the challenges here is that in Wikipedia, a conversation often happens entirely on one user's page. We therefore also attempted to count the number of communications that an adoptee made on one of his/her own talk pages that was intended to for the adopter. To do this, we automated a procedure that scanned every edit made by an

adoptee to his/her own talk pages, and looked upward on the page to see if the adopter had made a comment within five paragraphs above that. If so, we assumed that the adoptee comment was intended to be read by the adopter, and counted it as such. We did the same in reverse for adopter communications.

- *Amount of communication the adoptee does, in general.* This was measured by all edits that the adoptee made to anyone's talk page in Wikipedia, for a particular period of time.

The experience level of the adopter might considerably affect the success of the relationship; similarly, the experience level of the adoptee at the time of adoption could have a dramatic effect as well. We therefore also look at how editing behavior matches with:

- *Tenure of the adoptee* (how long the adoptee has been active in Wikipedia). We calculated this as the difference in days between when the adoptee was adopted, and when the adoptee made his/her first edit.
- *Tenure of the adopter* (how long the adopter has been active in Wikipedia). We calculated this as the difference in days between when the adoptee was adopted, and when the adopter made his/her first edit. Note that this means that an adopter who has multiple adoptees has a different tenure for each adoptee. This is perfectly reasonable, as we are examining the effects of adoption on each adoptee.

In order to measure the contributions of each of these variables, we focus on edit counts over two week intervals. Experimentation showed us that intervals much shorter than two weeks tended to be too volatile, and detecting patterns was challenging. Intervals considerably longer than two weeks (such as one month) tended to be too long; it required too long a delay from effect to measurement to be able to measure it appropriately. We also observed that communications between adopters and adoptees tended to diminish fairly rapidly after adoption, and so finding significant results for time periods beyond a few weeks after adoption was challenging. For each adoptee, then, we concern ourselves with three such intervals:

- “Period -1”: the two week period before an editor was adopted; the last day of this period is the day before adoption.
- “Period 0”: the two week period immediately following adoption: the first day of this period is the day that the editor was adopted.
- “Period 1”: the two week period immediately following Period 0.

We then measure the connections between editing and the variables above by performing linear regressions over these three time periods. Period -1 is used to establish an editing baseline, period 0 is the one where the “communication interventions” occur, and period 1 is the period where we measure the results. The philosophy is similar to one that might be used for a medical study: period -1 would contain the patient's baseline medical data, period 0 would be the treatment period, and period 1 would be the period which one

Variables	Mean	Std Dev	Min	25% Percentile	50% Percentile	75% Percentile	Max
1. edits ₁	10.67	37.03	0	0	0	3	426
2. edits ₋₁	11.63	31.06	0	0	1	10	384
3. comm ₀	3.852	5.301	0	1	2	5	67
4. user_talk_edits ₀	7.613	15.13	0	1	3	8	203
5. adoptee_tenure	321.2	397.0	14	37	136	472	2283
6. adopter_tenure	1097	563.0	65	652	1026	1441	2762

Table 2: Descriptive statistics for data to be used in regression analysis.

Variables	edits ₁	edits ₋₁	comm ₀	user_talk_edits ₀	adoptee_tenure
1. edits ₁					
2. edits ₋₁	.39 ***				
3. comm ₀	.25 ***	.11***			
4. user_talk_edits ₀	.42 ***	.31***	.61***		
5. adoptee_tenure	-.04	-.05	-.04	-.04	
6. adopter_tenure	-.02	-.01	-.10**	-.06 ⁺	.05

Table 3: Correlation analysis.

We use the following notation to represent p-values: *** $p < .001$, ** $p < .01$, * $p < .05$, + $p < .1$

would use to measure the immediate outcome of the treatment. Our dependent variable, then, is editing in period 1; all other variables are used as independent variables. Note that for this dependent variable, we specifically focus on the total number of edits to articles (Wikipedia “namespace 0”) as opposed to talk pages or other namespaces in Wikipedia. This is because our goal is to determine if adoption actually has an effect on the encyclopedia itself. Also, including communication in our dependent variable would overlap with some of our independent variables listed above. Since we use the two weeks before adoption as a critical part of this analysis, we restrict our data to those adoptees whose tenure is at least 14 days. This results in a set of 686 adoptees that we use for the models below.

In the following results, we use a subscript to indicate the period under consideration. We verify that linear regression assumptions are satisfied and eliminate a small number of irregular outliers. When appropriate, we use logarithmic transforms; the notation $\log(x)$ should be read as an abbreviation for $\log_2(1 + x)$. Edit counts and communication variables were logarithmically transformed; tenure measurements were not. Tables 2 and 3 present the descriptive statistics and the correlations of the variables used in the analysis before log transformation for the subset of adoptees under consideration.

Table 4 then shows the results of running four successive linear regression models. Model 1 shows the success of predicting an adoptee’s edit behavior based on past behavior alone. With this model, we see that edit behavior can be predicted with some reliability from past editing. Approximately one-third of the variance in the data is explained by this model, and the regression coefficients are statistically significant.

Model 2 adds to Model 1 the effect of adoptee/adopter communication. For this model, we combine all four kinds of adoptee/adopter communication (written by adoptee on adopter’s talk pages, written by adopter on adoptee’s talk pages, etc.) into a single variable. (We combine these four variables together because they are all highly correlated with each other; communication tends to reciprocate. For example, if an adoptee leaves a message on an adopter’s page, it is highly likely that the adopter will respond.) This model

again shows statistical significance among the regression coefficients, and a small increase in adjusted R^2 . In other words, we see a weak correlation between adopter/adoptee communications and edits by adoptees to articles. There are a number of possible interpretations of this result. The most obvious one to fans of the Adopt-a-user program would be that adopter/adoptee communication increases adoptee editing to articles. Another possibility, however, may be that this result merely illustrates the phenomenon that “editors who increase/decrease their Wikipedia activity (for whatever external reason) do so consistently over all kinds of Wikipedia pages; article edits as well as talk page edits are both similarly affected by external causes for page edit activity changes.”

To truly learn which of the above interpretations is likely to be correct, one needs a proper comparison group of editors who did not participate in Adopt-a-user. Choosing a historical control group for adoptees is actually quite challenging. There is a strong self-selection bias among those people who choose to participate in the Adopt-a-user program. People who ask to be adopted are not typical editors: they often have a particular problem they wish to solve that indicates a level of engagement not necessarily descriptive of the “typical” Wikipedia editor. Choosing a random historical set of editors who share these same characteristics but were not adopted is difficult. Another possible approach would be to use editors who wished adoption but were not offered it; as explained in section 6, however, most editors asking in good faith for adoption were made offers. To legitimately try to tease out these various interpretations would require a controlled study where editors are randomly selected to participate in Adopt-a-user or not. Such an intervention is outside the scope of this paper; nonetheless, we acknowledge it as a limitation of this work.

Model 3 revisits Model 2, but uses all edits made by the adoptee to user talk pages. In other words, instead of measuring only adoptee/adopter communication, we measure all communications made by the adoptee intended to be seen by other editors. We choose to use only edits to user talk pages, as opposed to article talk pages as well; this is because one might argue that article talk pages are clearly associated with article edits. We see in Model 3 that consider-

	Model 1		Model 2		Model 3		Model 4	
	Coef	P	Coef	P	Coef	P	Coef	P
Intercept	.343	**	.236	+	-.093		-.206	
$\log(\text{edits}_{-1})$.598	***	.576	***	.451	***	.448	***
$\log(\text{comm}_0)$.152	**				
$\log(\text{user_talk_edits}_0)$.320	***	.318	***
adoptee_tenure							.000319	
adopter_tenure							.000136	
R^2	.3409		.3497		.4095		.4108	
Adjusted R^2	.3399		.3478		.4077		.4073	

Table 4: Linear regression results. In all cases, the dependent variable is $\log(\text{edits}_1)$.

ably more variance in the data is explained than when using adopter/adoptee communications alone. This suggests that the driving factor in Model 2 is not particularly the direct effect of the Adopt-a-user program, but merely an indication that communication by adoptees correlates with article editing activity. Whether communication actually causes more editing would have to be determined via a controlled study. The correlation alone shows us that there is a strong connection between editing behavior and communication behavior. However, this connection may instead illustrate the principle that users who edit more generally communicate more, with both their adopters and other editors.

Finally, Model 4 is an update to Model 3 that includes adopter and adoptee tenure as well. We see that neither the tenure of the adopter nor that of the adoptee is statistically significant.

To summarize, we find two key conclusions from the above analysis. The first is that there are significant correlations between editor communication and article editing activity, but we need more evidence to make any actual claims about causality. The second conclusion is that communications specifically between adopters and adoptees do not seem to offer advantages over other forms of editor communication.

9. MENTORENPROGRAMM

Wikipedias in other languages also have a variety of mentoring programs. One example of these is the German Mentorenprogramm³⁰ (which translates simply as “Mentor program”). Mentorenprogramm is technically similar to Adopt-a-user, but has some important differences in its higher level structure. Our work in this paper is generally restricted to English Wikipedia due to limitations of time and space, but we spend a brief detour looking at this program in particular because of its recent publicity at Wikimania [4], and because we were able to perform some rapid and comparable quantitative analysis.

Mentorenprogramm uses a very similar approach to Adopt-a-user regarding requesting adoption and getting adopted. The key distinctions between the two programs are not in the technical details, however, but in the dramatically stronger oversight and monitoring that Mentorenprogramm implements. In Adopt-a-user, there are some guidelines on who should be a mentor, essentially based on amount of experience.³¹ Mentorenprogramm has a similar set of guidelines, but requires that a mentor submit a blurb about him/herself to a voting page and be voted in by other mentors via a two-

thirds majority.³² Mentorenprogramm also requires that each mentor has a co-mentor, to help handle issues that arise while a mentor may be unavailable. One more major way that Mentorenprogramm differs from Adopt-a-user is that it has a “Neulingsbörse,”³³ which roughly translates as “Beginners Exchange.” This page lists all current mentees and their level of activity within German Wikipedia.

We performed the same quantitative analysis for Mentorenprogramm that we did for Adopt-a-user, as described in Section 8. Though we do not include the results here due to space considerations, the outcomes are nearly identical to those shown for Adopt-a-user: the same variables are significant with comparable P-values, and the coefficient magnitudes are comparable. Mentorenprogramm has a number of interesting characteristics described above that may have positive effects on the program. From the perspective of this quantitative analysis, however, our findings for Mentorenprogramm are the same as those that we find for Adopt-a-user.

10. LOOKING TO THE FUTURE

In the previous two sections, we have shown that Adopt-a-user is struggling to bring the full experience of mentoring to its editors. How could this program, or others like it, be modified to have effects which are larger in magnitude or longer lasting? We offer some suggestions based on our previous discussions and findings.

Better communication capability. Mentoring critically depends on strong communication between the mentor and the protégé. Optimally, features such as threaded discussions, simplicity in cross-page communication, and more reliable notification of message receipt could make a huge difference. These are considerable software changes to Wikipedia, however, and may not be directly achievable. There are, alternatively other changes that are quite possible. Other communication mechanisms that lie outside of Wikipedia such as email, instant messaging, Facebook, and other alternatives likely allow for better and more private discussion between the mentor and the protégé. Rather than ignored at best or discouraged at worst, these sorts of communications can be directly encouraged. The mechanism for such encouragement is easy: at the start of an mentoring relationship, the mentor/protégé can be sent an automatic “mentor welcome” message that encourages them to swap email addresses, Facebook names, and so on.

Tolerance (or encouragement!) of social connections. The mentoring literature indicates that forging a friendship be-

³⁰<http://de.wikipedia.org/wiki/WP:Mentorenprogramm>

³¹<http://en.wikipedia.org/w/index.php?oldid=416476593>

³²<http://de.wikipedia.org/w/index.php?oldid=85641490>

³³<http://de.wikipedia.org/w/index.php?oldid=86979910>

tween the mentor and the protégé is important. Moreover, other research has shown that having strong social ties within a community helps keep people within that community [7]. Wikipedia’s culture, however, is generally opposed to active acknowledgement and fostering of social relationships. A large-scale shift in Wikipedia culture is difficult to achieve, but perhaps the effort should be started. Alternatively, in a similar manner as described above, the mentor and protégé should be encouraged to pursue social connections outside of Wikipedia. They can be Facebook friends or follow each other in Twitter; they can work on a blog together; they can simply email each other from time-to-time to see how each other are doing. If the community would allow it, a set of pages within the Adopt-a-user program could serve this purpose as well. Implementing an automated prompt encouraging mentors and protégés to do this is simple: see the previous point.

Facilitation of role modeling. Protégés learn from the examples set by their mentors. Protégés could be encouraged to look at the contributions of their mentors. What are their mentors doing, and how are they going about it? The protégés could be prompted to ask questions of their mentors regarding their work, and how they go about it. This could again be handled via an occasional regularly scheduled message (possibly by email or other “push” mechanism) encouraging protégés to pay attention to what their mentors are doing. One could also envision using machine learning techniques to identify which episodes from a mentor’s editing history would be particularly interesting to the protégé, and encourage the protégé to take a look.

Stronger leadership by mentors. In our readings of communications between adopters and adoptees, we saw that many adopters serve merely as “question answerers.” An adoptee asks a question, and the adopter answers it. This is useful, but it doesn’t at all address the career development aspects of mentoring. A Wikipedia mentor could proactively recommend types of work for the protégé to work on based on what the mentor knows of the protégé’s interests and strengths. A mentor could also encourage the protégé to work on specific tasks with the goal of exposing the protégé to new ideas, or enhance the protégé’s capabilities. Mentors could be encouraged to do this via occasional automatic reminders that specifically direct the mentor to think about the above issues. Finally, mentors could do a better job of this if they share areas of interest with their protégés. Matching algorithms could be used to help a protégé find a mentor with appropriate areas of interest and expertise. It is likely that most adopters in the Adopt-a-user program do not think of the sorts of activities in the last recommendation as part of their job description. It might be quite helpful to simply have a page that provides guidelines as to what mentors should do, and which would also further help encourage mentors to proactively engage in exposure-and-visibility and protection. Developing a software framework to encourage such behaviors might assist considerably in driving mentor motivation and proactivity.

Future Research

There is much more work in this area to be done. We have only analyzed in detail one online mentoring program (Adopt-a-user), and briefly looked at another (Mentorenprogramm). Other Wikipedias in other languages, as well as other online communities, have additional mentoring pro-

grams that would be worth examining. The study we present here has been based on historical data and conversations archived within Wikipedia; it would be interesting to conduct actual interviews with protégés and mentors that have participated in these programs. A controlled study where editors are selected to participate or not would tell us considerably more about the effectiveness of such programs. Finally, putting into place the suggestions described earlier and measuring the effects of them would tell us much about how well best practices for mentoring offline transfer to being effective in an environment such as Wikipedia.

11. ACKNOWLEDGMENTS AND DATA

We are grateful for the strong support and assistance by the various members of GroupLens Research. We also very much thank Tim Moritz Hector, who quickly sent us data and provided support regarding Mentorenprogramm. This work was supported by NSF grant IIS 08-08692.

We have made the user comments that we coded and our detailed analyses thereof freely available.³⁴

12. REFERENCES

- [1] R. B. Almeida, B. Mozafari, and J. Cho. On the evolution of wikipedia. International Conference on Weblogs and Social Media, Mar. 2007.
- [2] N. Cohen. Define gender gap? Look up Wikipedia’s contributor list. The New York Times, January 30 2011. <http://www.nytimes.com/2011/01/31/business/media/31link.html>.
- [3] D. L. DuBois, B. E. Holloway, J. C. Valentine, and H. Cooper. Effectiveness of mentoring programs for youth: A meta-analytic review. *American Journal of Community Psychology*, 30:157–197, 2002.
- [4] T. M. Hector. Structure of the German Mentoring Program and international comparison. Wikimania, 2010.
- [5] A. Kittur, B. A. Pendleton, and R. E. Kraut. Herding the cats: the influence of groups in coordinating peer production. In D. Riehle and A. Bruckman, editors, *Int. Sym. Wikis*. ACM, 2009.
- [6] K. E. Kram. *Mentoring at work: Developmental relationships in organizational life*. Scott Foresman, Glenview, IL, 1985.
- [7] J. M. McPherson, P. A. Popielarz, and S. Drobnic. Social networks and organizational dynamic. *American Sociological Review*, 1992.
- [8] R. A. Noe. An investigation of the determinants of successful assigned mentoring relationships. *Personnel Psychology*, 41, 1988.
- [9] J. E. A. Russell and D. M. Adams. The changing nature of mentoring in organizations: An introduction to the special issue on mentoring in organizations. *Journal of Vocational Behavior*, 51(1):1 – 14, 1997.
- [10] B. Suh, G. Convertino, E. H. Chi, and P. Pirolli. The singularity is not near: slowing growth of wikipedia. In *Proceedings of the 5th International Symposium on Wikis and Open Collaboration*, WikiSym ’09, pages 8:1–8:10, New York, NY, USA, 2009. ACM.

³⁴http://www.grouplens.org/system/files/adoption_data.zip