Awareness for consensus-seeking:
A mixed-initiative approach to summarizing discussion in Wikipedia

Travis Kriplean

I've fleshed out a few of ideas about how to better support the consensus process in Wikipedia. Obviously only a single part of this can be tackled for the cse574 project requirement. I'm not sure which part exactly. So, my apologies for the length, but this is the best way for me to make progress. Also, its very light on specifics of the mixed-initiative and AI techniques to be employed.

Introduction

Wikipedia follows a consensus-based collaborative authoring ethos, where nearly everything is subject to continual critique and editing [3, 5, 2] (see Figure 1). Strict majority voting is considered "evil"[1]; by explicitly discouraging voting in favor of discussion, Wikipedians hope to raise the strength of reason over problematic quantitative voting mechanisms.

The Mediawiki software, however, largely lacks technical mechanisms to support communal management of the vast amount of discussion generated during consensus seeking. This is a central problem for consensus-based communities, and is part of the larger problem of capturing design rationale. Organizing discussion to make past arguments and decisions salient is important for supporting the feedback loop at the heart of Wikipedia's model of consensus, where prior consensus is supposed to be taken into account as contributors wrestle with the same or similar issues at a later time (see figure 1). As Wikipedia ages, this lack of attention to historical discourse becomes increasingly problematic, resulting in illegitimate power plays and a great deal of frustration to both new and veteran contributors alike [3]. This is especially the case on popular articles where thousands of pages of discussion have accumulated over the years.

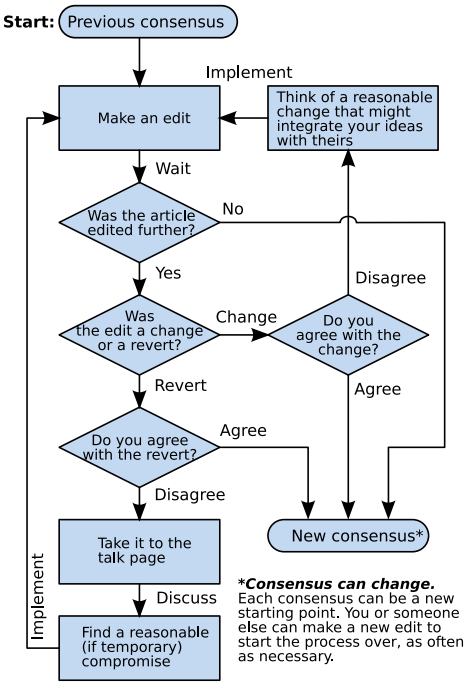

Figure 1: The idealized consensus process underlying Wikipedia's production (from Wikipedia:consensus).

We propose to design a mixed-initiative system to help Wikipedians create summary pages of discussions and changes to the articles. We target the *willing archivers*, participants who have taken upon themselves the task of using current mediawiki functionality to organize past discussions.

Summarizations of the discursive history underlying an article's production will give new contributors the ability to place the current issues being discussed with respect their historical context. Veteran contributors to an article will be able to easily refer new contributors to where a particular issue was already discussed. Participants seeking precedent on how to proceed on a content dispute can more readily consult other articles for similar situations.

Consensus in practice

> "It is very easy to create the appearance of a changing consensus simply by asking again and hoping that a different and more sympathetic group of people will discuss the issue. This, however, is a poor example of changing consensus, and is antithetical to the way that Wikipedia works. Wikipedia's decisions are not based on the number of people who showed up and voted a particular way on a particular day. It is based on a system of good reasons. Attempts to

[1] http://meta.wikimedia.org/wiki/Polls_are_evil

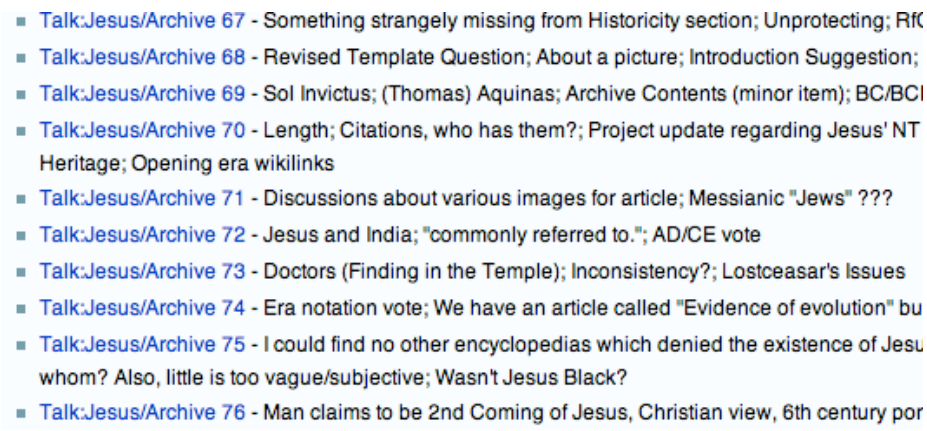

Figure 2: Archive annotation on Jesus article.

> change consensus must be based on a clear engagement with the reasons behind the current consensus." – *from Wikipedia:consensus*[2]

As the quote above from Wikipedia's consensus policy indicates, consensus-seeking in practice deviates from the ideal model for a number of reasons. Bad-faith participants can simply continue to raise the same issues until the current group of contributors to an article either grow too weary to argue against a change, or happen to be "sympathetic". Likewise, a group of editors can stonewall other contributors from making changes by dismissing new arguments and simply stating that prior consensus has been established. These *power plays* are illegitimate with respect to Wikipedia's ideals and body of policies, yet are quite prevalent [3].

The model also routinely breaks down among good-faith participants. Much of the time contributors simply cannot gain the proper awareness of what has been discussed in the past because of poor threading mechanisms, unorganized content, and an overwhelming amount of text. Technical mechanisms for awareness are low-level, consisting mainly of watchlists (subscribe to an article's edit history) and the ability to hyperlink to "diffs" between two given revisions of the article or talk page. Continuing the quote from the consensus policy, Wikipedians are supposed to collectively employ these and other features in their pursuit of consensus:

> "... and so in the new discussion section, provide a summary and links to any previous discussions about the issue on the articles talk page, or talk page archives, to help editors new to the issue read the reasons behind the consensus so that they can make an informed decision about changing the consensus."

On popular pages, however, it is impractical to expect contributors to parse the old discussions for relevant discussions. Instead, old discussions are usually lost, the same or similar arguments are repeated over and over, and illegitimate power plays with respect to the ideal consensus process are enacted. In short, Mediawiki does not provide adequate tool support for the important articulation work [4] of summarizing past discussions for input into future discussions. Even following an unfolding discussion can be daunting. This is a significant problem for maintaining quality articles and attracting/retaining community members.

Summarization tools

Despite the lack of tool support, we often observe a number of local strategies that Wikipedians have developed to combat proliferating discussion and recurrent topics. Common to each is the practice of creating discussion *archives*. These archives are created using Mediawiki's "move page" feature[3] and essentially moves non-current discussion threads (where currency is determined by the archiver) to a new page.[4]

[2] Accessed on 1/22/08.

[3] This feature moves a page and its edit history to a new location.

[4] It is also quite common for archives to be created by simple copy and paste, whereupon the edit history is not moved to the archive.

While archiving can help raise awareness of which discussions are active, they do not help contributors locate relevant past discussions. Moreover, on many of the popular pages, the archives themselves are overwhelming. For example, at the *Jesus* article, Wikipedians have created an index of the over 90 archives (probably around 50-100 pages of printed text each), painstakingly labeling each of the archives with a rough summary of their contents (see figure 2). The attempt may confuse more than help, but it reflects a truth about the Wikipedian community: there is almost always someone willing to work hard to compensate for the lack of technological support. These *willing archivers* are performing a task that is central to Wikipedia's production process and we can do better in providing the proper technological machinery.

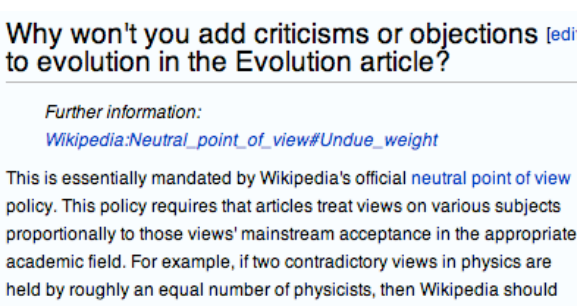

Figure 3: FAQ on the Evolution article.

We propose to build a mixed-initiative tool designed to help archivers create summary pages for current and past discussions for a given article.[5] On the article *evolution*, contributors have built what we take as our ideal summarization.

There are two main parts to the summary page: (1) A FAQ that contains a question stating the issue, a set of links to relevant policy, and a description of the prior consensus on the issue with a reasoned argument for the decision (see figure 3) and (2) a collection of topics, each with a set of links to particular discussions in the archives that pertain to that topic (see figure 4).[6]

Figure 4: Discussion posts clustered by topic.

Mixed-initiative approach

There are a number of tasks that a machine might perform to help an archiver build a summary page for a given article. Here we outline three steps that might be taken in a mixed-initiative approach. Note that the user interaction model is definitely not well thought out now.

Step 1: Generating topic groups. The goal of this step is to identify a number of conversation threads that can be treated as a *topic group* in some salient way, such as conversations indicating prior consensus about some issue. This is analogous to the manual discussion clustering exhibited in Figure 4. The first step is to employ a clustering algorithm which defines similarity based on a number of features, such as inverse term frequency (discussions sharing many rare words are more likely similar), types of policies cited, etc.[7] We will start with a basic model, but it may be worthwhile to learn the set of features that best captures related posts. The user would be involved in indicating whether the groupings are logical, which posts are relevant, and creating the title for the grouping (and perhaps a description as in Figure 3).

Step 2: Adding contextual information to the topic grouping. A number of properties of the discussions in a topic group (as generated in step 1) may be useful to contributors seeking consensus, and should be included in the summarization:

[5]This tool could be extended to work across thematically related pages, identifying, say, similar posts on the *evolution* and *creationism* articles. Working at a Wikiproject level may prove to be an even more useful approach than at an article level.

[6]The builders of the summary articulate the same motivation as us:

> "This page was created in response to certain topics being brought up again and again on Talk:Evolution, wasting many editors' time and energy by forcing them to respond repeatedly to the same issues. The FAQ serves to address these common concerns, criticisms, and arguments by answering the various misconceptions behind them"

[7]It may prove useful to have participants identify prototypical conversations they are interested in archiving in order to seed the clustering.

- Revealing who has been involved in conversations about a topic group can allow participatns to contact those who have been previously engaged in the topic. Because of turnover in the active contributors, this is affordance is critical for elevating the status of prior consensus.[8]
- It is useful to identify the types of policies that have been brought to bear to interpret the topic group, as done by the contributors on the *evolution* article. A couple techniques could be employed to automatically extract relevant policies. First, policy citations (links to policy pages) and keywords (clear-text references to policy) can be a first tier indication of relevant policy. Second, by building models of policy citation co-occurance throughout Wikipedia, policies that have a high probability of co-occurance with the policies referenced in the topic group may be included in the possible implicated policies. Third, again at the global level, a model could be built of the commonly co-occuring text surrounding policy citations to figure out which policies might be referenced "between the lines". Snowball-like techniques may be applicable here. Of course, the archiver would have the final say on the salient policies.
- Some discussions are very strongly coupled to specific features of the article being discussed and are therefore less *historically salient* (e.g. the wording for a specific sentence). It may be useful to build a model of how closely coupled the topic group is to the specific layout of a given version of an article. Presumably a summary page should privilege loosely coupled topic groups (recurrent high level concerns) and topic groups directly related to current discussions.

Step 3: Discussion dissection and visualization. Participants may want to be able to drill down in-depth into what was happening in a specific discussion, perhaps even extracting arguments. A number of potentially useful features require deep inspection of discussions:

- It would be very useful to be able to identify or isolate arguments in a discussion, or at least model whether two posts (and/or contributors) are in opposition to one another. A number of features might be used: reverts of the author's edits on the article page, discussion and indentation patterns on the discussions pages, escalating policy citations, etc.
- Ultimately, we want to provide tools for contributors to construct visualizations of topic groups, using techniques such as gIBIS [1]. These visualizations attempt to make salient the actual issues being discussed, the positions being considered, the arguments for and against, and the logical relationships between them. Identification of opposing arguments, as described in the previous bullet point, would be a first step in this direction.

Possible datasources

Data for training and evaluation may come from a number of sources. Our current intuition is that two data sources will be particularly useful:

- *Archiving practices.* Some of the pages, such as those discussed earlier, have manually organized past discourse. This data may be mobilized as a way to evaluate different techniques and perhaps even as training data (for example, see Figure 4).
- *Dispute resolution processes.* There are a number of dispute resolution mechanisms Wikipedia has established to allow disputing participants to request third party mediation. We can mine these in order to find interesting editing activity and evaluate our techniques in their ability to help summarize the dispute.

Other use cases

Although we have targeted the "willing archiver" with this tool, we ultimately want to support a wider range of use cases. Some of the important ones:

Dispute resolution. Help a group of people trying to work through a dispute gain a different view of the discussion landscape through an argument visualization. This feature may be appropriated into the dispute resolution processes that Wikipedia has in place, such as Requests for Comment, Mediation, and Arbitration. Before obtaining third party intervention, dissenting participants may

[8] It would be nice to do an empirical analysis of the turnover in the main contributors to popular articles.

be strongly encouraged to at least agree on a visual representation of the argument at hand. This would greatly reduce the difficulty of providing outside assistance to content disputes. The work proposed here is a necessary precursor to such support.

On demand related discussion identification. A contributor may want to be able to quickly identify the past conversations that are most closely related to an ongoing discussion. This use case will be an intermediate result of the proposed research, although it will receive minimal UI attention.

Lightweight visualization. Unfolding discussions in Wikipedia are difficult to follow. What's the issue? Who's involved? What are the arguments? On one article, we found this post illustrating the difficulty of following a large discussion and moving toward consensus: "Thanks for clarifying your views on the article. Unfortunately most of the editors who work on this article do not share the same view (i hope i am speaking for most of the editors)." Lightweight rapid construction of visualizations for current ongoing discussions may help. This is more focused on immediate discourse, not so much on increasing the saliency of discourse history.

Timeline

02/13 – Milestone 1

- Investigate archiving practices on other popular pages. Note similarities and differences, and, if possible, try to figure out the reasons for their emergence.
- Try to identify a broader class of training data similar to that on the evolution page.
- Extract conversation threads and text from talk pages.
- Experiment with simple methods of clustering conversation threads.
- Define the possible stages of the mixed-initiative interaction. Build mockups if possible.
- Identify the specific AI techniques that should be employed.

02/27 – Milestone 2

03/12 – Presentation

03/21 – Writeup

References

[1] J. Conklin and M. L. Begeman. gIBIS: A Hypertext Tool for Exploratory Policy Discussion. In *CSCW '88: Proceedings of the 1988 ACM conference on Computer-supported cooperative work*, pages 140–152, New York, NY, USA, 1988. ACM.

[2] A. Forte and A. Bruckman. Scaling consensus: Increasing decentralization in wikipedia governance. *To appear in HICCS*, 0, 2008.

[3] T. Kriplean, I. Beschastnikh, D. W. McDonald, and S. Golder. Community, Consensus, Coercion, Control: CS*W or How Policy Mediates Mass Participation. In *Proceedings of GROUP 2007*, pages 167–176, 2007.

[4] K. Schmidt and L. Bannon. Taking cscw seriously. *Computer Supported Cooperative Work (CSCW)*, 1(1):7–40, March 1992.

[5] F. B. Viégas, M. Wattenberg, J. Kriss, and F. van Ham. Talk before you type: Coordination in wikipedia. *In Proceedings of HICSS*, pages 78–78, 2007.