

Tools

The essays that follow explore the need for tool development to better support developers, organizers and moderators. These position papers were synthesized out of input from staff and community experts as well as prior independent research.



Published
January 2019

Authors

Niharika Kohli
Ryan Kaldari
Joe Matazzoni

Contributors

Amir Aharoni
Asaf Bartov
Adam Baso
Trevor Bolliger
Maria Cruz
Bryan Davis
David A.M. Goldberg
Aaron Halfaker
James Hare
Danny Horn
Ramsey Isler
Jon Katz
Josh Minor
Toby Negrin
Margeigh Novotny
Nirzar Pangarkar
Joseph Seddon
Alex Stinson
Nick Wilson
Leon Ziemba
Leila Zia



In early 2006, a large-scale vandalism attack wiped out thousands of articles on the English Wikipedia. The editors were defenseless and the site was vulnerable. In response, four Wikipedians came together to write TawkerBot, the first anti-vandalism bot [1] for Wikipedia. This bot proved to be a life-saver for the site, and today more than 300 bots work round the clock on English Wikipedia to ensure the smooth functioning of the site. Tools like Twinkle [2] (tool library that helps editors perform wiki maintenance), Huggle[3] (a diff browser for rapidly reverting vandalism), HotCat[4] (allows a user to easily add and remove categories to pages), AutoWikiBrowser [5] (semi-automated MediaWiki editor), etc. drive many of the tasks power editors do on English Wikipedia every day.[6]

At the same time, smaller language wikis like Hindi Wikipedia[7] have problems coping with vandalism and keeping up with content moderation needs. Unlike English Wikipedia, they don't have the corps of volunteer developers able to write tools to defend and curate the site's content. It's much harder for those communities to grow their content or their editor base, because active contributors are stuck doing manual drudge work that bigger wikis automated years ago.[8][9]

Tools for Developers

Wikis need code contributors as much as they need content contributors. Templates, gadgets, and bots act as superpowers in making editors more efficient at their tasks. Experienced editors use these tools to create and maintain backlogs, track incoming edits and their quality, perform mass actions on pages, ward off vandalism and more. However this superpower is limited to wikis which have contributors who are able to write such code for the site. Bringing these important resources to all wikis is fundamental to bridging the equity gap across all language wikis, and is going to be a key factor in helping us realize the goal of amassing the sum of all knowledge.

Tools for Organizers

The Foundation's 2018-19 annual plan recognizes organizers as "fundamental implementers" and a "core asset" of the free-knowledge movement. Unfortunately tools that support organizers' efforts are frequently ad-hoc, poorly documented and available only on certain wikis. Access problems can be particularly acute in smaller communities, where the technical skills required to set up and run bots, scripts and other technologies are often scarce.

Organizers' needs fall into four main areas. "Community-building" tools are required to help organizers inform, engage and manage the work of their communities. "Outreach and promotion" tools will help organizers advertise their activities and recruit new members. "Event-management" tools are necessary to more efficiently carry out tasks like event signup and conference scheduling. Emotional labor, the fourth category of need, demands better support infrastructure to nourish and sustain the passionate organizers on the ground, over time. Finally, two overarching meta-problems are key areas of interest among organizers. One is the need for better guidance about best practices and the tools that do exist. As one organizer put it, "There are a lot of tools we don't know about or know what they can do for us. We need someone to help us understand what we are missing, and what to do and how to do it." The second is the need for a mechanism that can replace or augment categories, so that organizers will be able to classify content effectively and more efficiently tap volunteers' subject interests—a primary motivator, especially of new editors.

Tools for Moderators

A critical, but often overlooked, aspect of the workflows that make our projects successful are the tools and processes used to review and moderate our content. For the most part, the Wikimedia Foundation has taken a hands-off approach to content moderation tools and allows the community develop their own solutions (with a few exceptions such as Recent Changes filtering). As one would expect, the systems

built by the community utilize the building blocks at their disposal: templates, gadgets, user scripts, and wikitext. These tools and processes suffer from several significant problems, many of which have already been mentioned above: lack of portability, limited capabilities, lack of automated testing and code review, lack of localization support, etc.

Another major problem with these tools and processes, especially those created for content moderation, is their high learning curve. For example, on English Wikipedia there is a system for submitting, reviewing, and publishing draft articles called Articles for Creation (AfC). In order to participate as a reviewer in AfC, you have to install a special user script, be granted a special right though a unique vetting process, and use several obscure templates and categories. The complexity of this process limits the number of people who are able and willing to participate, which in turn leads to a less diverse pool of reviewers. This lack of diversity may contribute to problems of systemic bias in our content.[10] [11] The small number of reviewers also makes the review process slow, often taking a week or longer to review a submitted draft. This likely contributes to the overall inefficiency, resulting in decreased newcomer productivity. [12] Unless we make moderation tools work for less technical users, it is unlikely the pool of moderators will grow or diversify.

Similar examples can be found throughout the moderation processes for our projects, including workflows for article assessment, deletion, and problem triaging; workflows for reviewing edits; workflows for reviewing and organizing multimedia contributions; workflows for proofreading Wikisource transcriptions; and more. While the Wikimedia Foundation has historically focused on building software for reading and editing content, the other critical pieces of the wiki ecosystem have been neglected, leading to volunteers feeling overwhelmed and unsupported. In a 2015 survey of experienced editors across 10 projects, only 35% said that the Foundation was mostly or completely meeting the community's technical needs around identifying and surfacing content problems.

[13] Unfortunately, the Foundation's hands-off approach has resulted in a lack of credibility in this area. To build our credibility, we should first focus on the areas where there is a clear need for better tools, such as fighting vandalism and sock-puppetry. We should also investigate how editors transition into becoming moderators so that we can better facilitate that transition. Once we've proven our capacity to understand their motivations and work with moderators to build effective tools, we will then have the mutual trust needed to tackle more difficult workflows such as article deletion and conflict mediation.

Clearly, there's a lot of work for us to do in this area as we have only scratched the surface thus far. If we want to increase the capacity of our communities to efficiently and effectively moderate content, it is time for the Foundation to begin investing seriously in this area.

Examples

- HotCat
- Huggle
- Twinkle
- AutoWikiBrowser
- Programs and Events Dashboard
- Wikimedia Cloud Services
- CentralNotice
- GeoNotice

Areas of Impact

- Template
- Gadgets
- Bots
- Editing and Administration APIs
- Discussion systems
- Messaging systems
- Contributor Analytics
- Developer Advocacy and Outreach
- Translation and Localization Infrastructure
- API and Tool Documentation

01

Tools for Developers

Empowering our volunteer developers to write better code that can work across wikis is going to be a key factor in helping us gather the sum of all knowledge. Wikis need code contributors as much as they need content contributors. Templates, gadgets, and bots act as superpowers in making editors more efficient at their tasks. Experienced editors use these tools to create and maintain backlogs, keep track of quality of incoming edits, perform mass actions on pages, ward off vandalism and more. However this superpower is limited to wikis which have contributors able to write code for the site. This creates disparity in the resources available to

wikis. Bringing these important resources to all wikis is fundamental to bridging the equity gap across all language wikis.

Pan-wiki Tools Platform

Empowering our volunteer developers to write better code that can work across wikis is going to be a key factor in helping us gather the sum of all knowledge. Wikis need code contributors as much as they need content contributors. Templates, gadgets, and bots act as superpowers in making editors more efficient at their tasks. Experienced editors use these tools to create and maintain backlogs, keep track of quality of incoming edits, perform mass actions on pages, ward off vandalism and more. However this superpower is limited to wikis which have contributors able to write code for the site. This creates disparity in the resources available to wikis. Bringing these important resources to all wikis is fundamental to bridging the equity gap across all language wikis.

We must evolve our platform so that support tools can work on all our wikis seamlessly. Right now a lot of developer code lives on specific wikis (gadgets, Lua modules, templates) where it really isn't possible to do any type of testing, code reviews or debugging; nor is there any straightforward way to add localization or RTL support. This often leads to issues like security vulnerabilities, [1] conflicts with MediaWiki deployed extensions, [2] and bugs due to lack of maintenance. Also, in its current state, having code hosted on the wikis (in a per-project fashion) makes it hard to get in the mindset of having the code work across wikis. It's easy to get sucked into customization and forget to think about things like RTL rendering or localization. This evolution depends on core services being available to developers across communities which they can use for building tools. Examples of these services include APIs can be used to do copyright violation detection, vandalism detection, and image recognition, provide access to better statistics, and so on. Part of growing these services involves better partnerships with

companies like Google, Turnitin and others already providing such functionality. Finally, for this to succeed, Wikimedia staff will need to collaborate with our volunteer developer communities to come up with documentation and best practices for creating new tools. Tools that facilitate communication among engineers and volunteer developer communities will be key to achieving this goal.

An example of these might be tutorials and guidelines like these documenting how gadgets can make use of OOU [3] to standardize our interfaces and make them more accessible for everyone.

Bots : Potential and Risks

Bots automatically perform repetitive tasks, but like any type of automation the advantages come with risks. The way that bots are wielded by editors can shape the character of a wiki's community. An editor's ability to rapidly apply changes that impact contributions that may have taken hours of individual work is a form of power that can alter contributors' attitudes about their willingness to participate. Bots are also complex pieces of programming, typically developed by a single individual trying to solve a specific problem in a specific context. There is nothing inherently wrong with this, as volunteer development is a core component of our culture. However, there is little standardization in the development of bots, or guidelines for creating them to serve the wider community of projects. We must begin to evolve our thinking about the future of Bots in the context of the discussions on structured data outlined in the Ubiquity paper, and in our Developer, Moderator and Organizer Tool positions.

Gadgets & Gadget Usage

Unlike bots which may make use of external resources or specialized programming (e.g. neural net programming for ClueBot),[4]

gadgets are built for and within the browser ecosystem in JavaScript or CSS. The number of gadgets enabled by default for English Wikipedia, sixteen, [5] far exceeds the number for other wikis such as German [6], Hebrew, [7] Tamil, [8] and Italian [9], while eight other gadgets are manually enabled by over 30,000 active users. We have seen some migration of gadgets between projects, based on volunteer efforts by those who are familiar with JavaScript and CSS. Examples include HotCat, [10] Navigation popups, [11] UTC:LiveClock,[12] and WikEd. [13]

Gadgets are comparable to apps and extensions on other platforms, and we are faced with the same challenges of testing, reliability, developer standards, and review. Gadgets can easily break things for users due to bad programming practices, poor testing, or conflicts with other MediaWiki extensions, [14] [15] [16] and there is no easy way to identify them as the source of the problem. Most importantly, gadgets do not go through any sort of staff review process before being deployed which means they are a potential vector for abuse.

In addition to the issues of segregated usage, many gadgets are prone to browser performance issues due to their large JavaScript code consuming bandwidth and processor resources [14].

Conclusions & Implications

Though gadgets are heavily used on most Wikipedias their application is uneven due to three factors: overspecialization, concentration (in terms of development and use), and lack of instrumentation.

Larger projects with larger language communities (e.g. English) dominate the development of gadgets. And while correlation between programming expertise and English usage does not dictate their development, it does reflect the concentration of support based on privilege seen in other contexts.

The functionality of gadgets tends to be overspecialized and adapted to the wiki for which they were developed, thereby limiting internationalization, adaptation between left-to-right/right-to-left writing systems, and the development of more general and migratable solutions.

Finally, we lack instrumentation to track gadget requests, development, distribution and usage. Without the means to measure and analyze the gadget space we will continue to suffer the inefficiencies of redundant development efforts, overspecialization without awareness of the costs or consequences of these issues.

Wishlist & Technical Requests

In addressing the annual Wishlist, [18] the Community Tech team has recognized that many of the frequently requested tools, bots and gadgets already have existing solutions created by the community. Some of these tools have alternate versions adapted to specific languages (such as the TemplateWizard [19] developed for Hebrew [20] and German [21] wikis, while others, like the Blame tool, have been created multiple times [22][23]. Tools like Crosswatch [24] and a pageview stats tool were on the wishlist in 2015 but had already been developed by volunteers. Similarly, user Legoktm developed the GlobalPreferences extension in 2013, well before it was nominated for the wishlist [25]. Looking forward, all of these examples demonstrate the need to improve the tracking of requests and implementations across our various communities.

Conclusions & Implications

Clearly we must continue to rely upon a robust and dedicated community of volunteer developers to deliver the means to interact with Wikipedia's complex features. Much of their work has become indispensable to their specific communities and on more global scales as well, and in this sense they have the potential to have a

multiplying effect on the efforts of the Foundation. However, in the context of tools, we once again confront the challenges that come from our uniquely heterogeneous character: knowledge and tools concentrated in certain communities do not cross over to those that might make use of them. Often there are communication gaps between different wikis, which mostly act as individual ecosystems of users and tools. We lack a place to surface the tools used by various wikis, which would not only address issues of redundancy in both development and requests, but encourage the type of cross-community tool portability that is envisioned here. This is a challenge that mixes technical and cultural issues, some—like multilingualism and different community densities across different projects—are inherent to Wikipedia, while others—like widely divergent practices in documentation and standardization—are associated with any volunteer development community. By prioritizing the development of a shared platform for tool identification, user requests, and pre-deployment assessment we can begin to address both of these issues simultaneously.

Notes

- 1 “Persian Wikimedia cryptocurrency mining incident” <https://lists.wikimedia.org/pipermail/wikitech-l/2018-March/089636.html>
- 2 Using Hotcat after saving with visual editor tries to edit old version <https://phabricator.wikimedia.org/T170896>
- 3 OOUI is a UI component library created by WME, aimed at providing a consistent UI experience that works well for all languages. <https://www.mediawiki.org/wiki/OOUI>
- 4 ClueBot https://en.wikipedia.org/wiki/User:ClueBot_NG
- 5 Default gadget usage on English
- 6 Wikipedia [Special:GadgetUsage](https://en.wikipedia.org/wiki/Special:GadgetUsage)
- 7 <https://de.wikipedia.org/wiki/Spezial:GadgetUsage>
- 8 <https://he.wikipedia.org/wiki/תבנית:GadgetUsage>
- 9 <https://ta.wikipedia.org/wiki/சிறப்பு:GadgetUsage>
- 10 <https://it.wikipedia.org/wiki/Speciale:GadgetUsage>
- 11 Hotcat <https://en.wikipedia.org/wiki/MediaWiki:Gadget-HotCat>
- 12 Navigation Popups https://en.wikipedia.org/wiki/Wikipedia:Tools/Navigation_popups
- 13 UTC LiveClock <https://en.wikipedia.org/wiki/MediaWiki:Gadget-UTCLiveClock>
- 14 WikEd <https://en.wikipedia.org/wiki/MediaWiki:Gadget-wikEd>
- 15 <https://phabricator.wikimedia.org/T178348>
- 16 <https://phabricator.wikimedia.org/T170896>
- 17 <https://phabricator.wikimedia.org/T22134>
- 18 <https://phabricator.wikimedia.org/T142461>
- 19 Community Wishlist https://en.wikipedia.org/wiki/Wikipedia:Bot_requests
- 20 Template Wizard <https://www.mediawiki.org/wiki/Help:Extension:TemplateWizard>
- 21 Hebrew Template Wizard https://he.wikipedia.org/wiki/תבנית_קישור:Gadget-TemplateParamWizard.js
- 22 German Template Wizard <https://de.wikipedia.org/wiki/Wikipedia:Technik/Skin/Gadgets/Vorlagenmeister>
- 23 <https://github.com/wikiwho/WhoColor>
<http://wikipedia.ramselehof.de/wikiblame.php>
- 24 Crosswatch <https://tools.wmflabs.org/crosswatch/welcome>
- 25 Global Preferences <https://www.mediawiki.org/w/index>

02

Tools for Organizers

The Foundation's 2018-19 annual plan recognizes organizers as "fundamental implementers" and a "core asset" [1] of the free-knowledge movement. But tools that support organizers' efforts are frequently ad-hoc, poorly documented and not universally available—particularly to smaller communities. As the movement puts an increasing emphasis on knowledge equity, the need to understand and support movement organizers is more vital than ever. This white paper is an early effort to analyze and document organizers' main areas of need. This examination

will be followed and deepened soon by the annual plan-mandated Movement Organizer Study. [1]

Who are organizers and what do they contribute?

For the purposes of this analysis, a “movement organizer” is anyone who purposely seeks to motivate, attract and/or direct volunteer attention. Organizers’ efforts generally fall into the following categories: content creation (e.g. content drives, editathons, photo expedition and writing competitions,) knowledge dissemination (e.g. training, conferences and campaigns,) process improvement (e.g. standardizing sources, developing procedures, style guides and templates,) and outreach (e.g. lobbying and partnership development, especially in relation to GLAMs, governments, NGOs, etc.) and emotional labor.

Some organizers work independently, possibly having been trained by more formal groups. But most operate within the framework of various structures that support their efforts – they range from legally incorporated nonprofit entities (Chapters, of which there are about 40 that work to improve laws and negotiate content donations for example,) to Independent Organizations operating on the edges of recognized groups within a given movement. Unincorporated User Groups, with regional or thematic focuses (e.g., Community User Group of Greece, Wiki Medicine,) of which there are about 100 provide the other structures of support.

Many groups coalesce around shared interests in organizing or improving efforts around subject areas or types of tasks (e.g. copyediting). These “WikiProjects,” “Portals” or just “Projects” appear across many wikis, with about 300 concentrated on en.wiki. Finally there are Campaigns: unaffiliated working groups, usually in part supported by organizers at existing affiliates or user groups. These are usually time limited efforts run with support from the more permanent Affiliates and User groups,, such as Art+Feminism, Wikipedia Asian Month, or Wiki Loves Monuments.

Wiki “functionaries” such as stewards, admins, and bureaucrats, whose activities generally don’t focus on “attracting and directing volunteer attention,” are not considered organizers.

Problems and Needs

As one organizer put it, “There are a lot of tools we don’t know about or know what they can do for us. We need someone to help us understand what we are missing, and what to do and how to do it.”

As one staffer put it,

“We need to give people the scaffolding of how to be successful.”

We are some distance from addressing these two points because Wikipedia organizer tools remain hard to discover, largely restricted to use by people with technical skills, and limited by documentation that is non-standardized and/or inconsistently maintained. Organizer tools broadly fall into four groups: community-building, outreach and promotion, event-management, and conference tools.

Common Needs

These have many current issues in common. There is no step-by-step process that organizers can follow to create a new project or campaign and make sure that it will be successful. A lack of end-user documentation (as opposed to technical) often makes the tools organizers do locate unusable. Most organizers do not have technical backgrounds, so the bots, scripts, Wikidata tools and other technologies that benefit some groups enormously are not available to all. This problem can be particularly acute in smaller communities, where it is less easy to find people with the required skills. Organizers want an easy way to announce an event and find out who is going to come, for which they generally must turn to off site tools. Setting out work for a community to accomplish is another common organizer need, yet the wikis lack even the most basic functions of task-management software.

Community-Building Tools

“Community-building” refers to a nexus of functions that organizers require to inform, engage and motivate their communities — to

build “a movement not an event”, as one organizer put it. Performing these functions currently requires a patchwork of tools, a high degree of technical sophistication and considerable manual effort.

The key community-building function is conversation: between group members and between organizers and the group. Organizers need easy ways to make announcements and invite large numbers of people to participate in a discussion, while group members need better ways to share and discuss information among themselves.

Outreach and Promotion Tools

Community-Building tools address a need for more effective conversation within a group, but organizers also need to reach editors, readers and others who are not yet in their circle of contacts. There is currently no simple or effective way for organizers to promote their groups or events to wiki readers or editors who are likely to be interested (e.g. based on geography or demonstrated subject interest). A founder of one of the largest and most active user groups recently named their inability to reach out beyond a circle of existing, very active users as one of the biggest issues with current tools. Finding new members is slow and often accomplished through personal networks.

Tools like CentralNotice [2] and related tools like Sitenotice exist and do reach readers, but they lack important features such as targeting. Whereas Geonotice, [3] for example, can target by location but reaches only registered users and only on their Watchlists. Such tools are also subject to many restrictions and layers of approval. [4]

Event-Management Tools

Organizers of edit-a-thons, training sessions, photo walks and similar in-person events have urgent needs related to event management. The Event Metrics [5] project, currently in development, will aid event organizers with better data about their contributions, but event management itself is out of scope for that project. An opportunity exists for synergy between Event

Metrics and future event-management tools, since both make use of similar input data about the event and its participants. Event organizers need specific tools that support in person participant signup and sign-in, “day of” Wiki account creation, and conference tools.

Generally, event management is complicated by privacy issues and technical restrictions designed to prevent exploitation of the wiki platform. In the first case, the management of event participants requires providing an email address, or the use of third-party management tools such as Eventbrite. In the second, limitations on the number of wiki accounts that can be created from a single IP complicate the process by which participants engage the event. There are workarounds for some needs, but they are not universal or well known.

Conference Tools

At a higher level of organization, conferences are large annual or semi-annual events with their own needs. Each conference must rebuild or remix the infrastructure for conference calendaring, signup, payment, scheduling etc for each event. This is true even for recurring annual events. More direct links to on-wiki activities, such as the pages related to sessions, including proposal submission, scoring of submissions, scheduling, and presenting scheduled programs that happen in isolation on wiki.

Recommendations

The Foundation is currently undertaking a Movement Organizer Study which will provide more data on how organizers do their work, and their needs and challenges. This broad survey of organizer workflows be useful for community interventions, product teams and others wanting to better support organizers’ functional and emotional needs. Addressing these opportunities will be non-trivial, but the benefits will be felt more broadly than by organizers alone. “Subject interest” and “Group conversations”, in particular, are fundamental tools whose potential applications are widespread.

Improve The Category System

Research and experience indicate that an abstract interest in the movement, per se, is not what motivates volunteers at the beginning of their wiki journey: they come to Wikipedia to share knowledge and passion for some subject. The Wikiproject directory [6] demonstrates the wide range of interests organizers and volunteers pursue formally, from folklore to pharmacology, football to firearms. However, such projects are hampered by the fundamental weaknesses of the wiki category system [7].

Fundamental improvements to categorization will produce benefits across all communities, but serve organizers particularly well. Better categorization would introduce a foundation for better traversal and interconnection of subject areas, and support better recognition and promotion of interests for participants and organizers, respectively. Organizers need automated ways to classify (and therefore assemble and search for) articles and tasks by subject, and to reach out to potential participants based on demonstrated subject interest.

Facilitate Group Conversations and Task Management

The consultation about “fixing talk pages” [9] in our 2018-2019 annual plan will initiate a conversation about the shortcomings of this core wiki communication platform. However changes to the talk page system may not be the only way to address organizers’ communication needs.

Tracking, annotating and debating changes are fundamental “gestures” for all engaged wiki editors and contributors, but organizers bring an additional set of requirements such as task assignment and management that could potentially transform communication for everyone. Organizers strive to keep groups engaged and informed, and their efforts cross the border between the Wikipedia universe and the rest of the world in synergistic ways.

Organizers and group members need tools to work with what may become a new core wiki concept: the task. Different from the implicit notion of the edit which drives the development of an article, a task would have

a fixed identity, a lifespan, the means of being classified (e.g. by type, subject area, or degree of difficulty,) ownership privileges, and associated discussion threads. Tasks might not necessarily be circumscribed by or destined for single wikis (e.g. getting decentralized support on a communications plan, solicit support for event organization, or other “jobs” that don’t aren’t limited to single, on-wiki activity,) thereby introducing a new mode of cross-community connection.

Email is the only message-delivery mechanism outside talk pages that our system currently supports. But the world we live in now requires that organizers broadcast and stay in touch with members on multiple social-media platforms at the same time (see also Experience). Feeding these multiple platforms manually is labor-intensive. An important area will be determining how we should best should incorporate social-media channels into our notifications and communications system.

Improve Outreach and Promotion Tools

Given the enormous traffic the wikis command, failing to tap the communication potential of our platforms would be an enormous opportunity missed. Better outreach and promotion tools would enable organizers to more effectively reach desired audiences. Overcoming the problems associated with using the wikis for mission-focused promotion will require efforts in both the technical and social arenas.

Meeting organizers’ promotional needs may require the reconsideration of some longstanding ideas and prohibitions, e.g. experimenting with limited, noncommercial, movement-focused advertising to wiki readers. Given the wikis’ enormous traffic, even narrowly focused banners might prove effective, but this would require new oversight to make sure the level of promotion is not excessive. There may be ways to respect editors’ privacy while still targeting them with messages based on geography or demonstrated subject interest (for example editors might be requested to voluntarily submit such information).

As highlighted in earlier sections, subject interest

is a key motivator of wiki activities. Providing simpler and better ways for organizers to perform outreach based on subject will pay dividends in increased efficiency and effectiveness across a wide range of activities. Rather than focus on revising and reforming the MediaWiki Category system itself is unlikely to be the fix to this problem. Future solutions may come from experiments currently ongoing with structured data, AI projects like ORES Draft topic model, [10] or from a system based on link analysis, like the one that powers Recommendation API. [11]

Improve Tool Discovery, Documentation, Ease of Use

Wiki tools are hard to discover, use and install, and they are poorly documented. It is time to dedicate more resources that are responsible for documentation and organizing and standardizing community writing about best practices and model workflows. To address discoverability, the existing Tools Directory [12] and the Toolhub [13] project demonstrate a concrete effort. The Directory provides a list of tools by title, author and categorization tags, while Toolhub documents the underlying structure [14] for cataloging them. However, there is much room for improvement in terms of how these tools are explained and promoted to less-experienced users.

Immediately Fix “Day-of” Account Creation

There is something we can do address a core need in the very near term. The limits and controls around account creation on shared IPs are the most immediate block on effective events. This ticket [15] includes a good discussion of solutions. Current solutions share many of the common problems identified across organizers: it is not available on all wikis, not well known or understood and thinly documented.

Notes

- 1 https://meta.wikimedia.org/wiki/Wikimedia_Foundation_Annual_Plan/2018-2019/Audiences#Outcome_1:_Progressive_Onboarding
- 2 Central Notice <https://meta.wikimedia.org/wiki/Special:CentralNotice>
- 3 GeoNotice <https://en.wikipedia.org/wiki/Wikipedia:Geonotice>
- 4 Usage Guidelines for Central Notice https://meta.wikimedia.org/wiki/CentralNotice/Usage_guidelines#Goals
- 5 Event Metrics [16] https://meta.wikimedia.org/wiki/Community_Tech/Tools_for_program_and_event_organizers
- 6 Wikiproject Directory https://en.wikipedia.org/wiki/Wikipedia:WikiProject_Directory
- 7 A few of these weaknesses: Categories are monolingual, ad-hoc and completely nonstandard, so developing universal tools based on them is challenging. Items in sub-categories don't inherit from parents, which is why searching broad categories typically yields few article results, contrary to user expectations. This non-inheritance also leads to hyper-specificity and bloat, to the point of absurdity. [8] Technically speaking, the category system is not a “taxonomy”—a logical structure where all child categories are more specific classifications of the parents and wholly contained by them (e.g., Musical Instruments > String Instruments). It is, instead, a “category network,” where children have some relationship with parents, but the relationship is unpredictable and therefore less useful (Musical Instruments > Orchestras). Categories can even be circular, with one category being both a parent and a child of itself.
- 8 Twitter feed dedicated to surfacing absurd categories related to Cats on Commons <https://>

twitter.com/CommonsCat

9 Talk pages Consultation [17] https://meta.wikimedia.org/wiki/Wikimedia_Foundation_Annual_Plan/2018-2019/Audiences#Outcome_2:_Communication

10 ORES Draft topic model https://meta.wikimedia.org/wiki/Research:Automatic_new_article_topics_suggestion

11 Recommendation API https://meta.wikimedia.org/wiki/Recommendation_API

12 Tools Directory <https://tools.wmflabs.org/hay/directory/>

13 Toolhub <https://meta.wikimedia.org/wiki/Toolhub>

14 Toolhub data model https://meta.wikimedia.org/wiki/Toolhub/Data_model

15 <https://phabricator.wikimedia.org/T202759>

16 Eight organizer interviews: As part of the research into the ongoing Event Metrics project, I conducted about 12 hours of interviews with seven event organizers suggested by program staff. (I also received one written interview.) These interviews, for which I have extensive notes, cover the full workflow of event creation, management and reporting. I did not ask subjects for the right to publish but could request if desirable. https://meta.wikimedia.org/wiki/Community_Tech/Event_Metrics

17 Extensive talk page discussion: Also as part of Event Metrics, J. Matazzoni engaged extensively with organizer on the project talk page. (The discussion is organized by subject. E.g., here on the problems associated with Account Creation.) https://meta.wikimedia.org/wiki/Talk:Community_Tech/Tools_for_program_and_event_organizers, https://meta.wikimedia.org/wiki/Talk:Community_Tech/Tools_for_program_and_event_organizers#Step_3:_Wiki_account_creation

03

Tools for Moderators

The tools and processes used for reviewing and moderating content are critical to the growth of the projects, but have historically been overlooked. Communities have, for the most part, been left to develop their own moderation workflows, tools and solutions – which means that community-built systems tend only to utilize the building blocks at their disposal: templates, gadgets, user scripts, and wikitext.[1] This approach results in solutions with limited capabilities, that tend not to be portable, lack automated testing and code review, and don't support localization.

Many of these tools and processes also have steep learning curves and require considerable technical acumen to use. [2] The more complex the process is, the more it limits the number of people who are able or willing to participate. This, in turn, leads to a less diverse pool of reviewers and contributes to problems of systemic bias in

our content.[3][4] And, fewer moderators makes the review process slow, impacting newcomer productivity.[5] In order to grow and diversify the pool of moderators, moderation tools must be made to work for less technical users, and the Foundation must drive investment that enables community developers to create more portable, localizable tools, and in resources to support better practices in tool development.[6]

The Example of English Articles for Creation

“New article creation is a battlezone where socializing newcomers seems to take a back seat to ensuring quality control.”

Since Wikipedia’s dramatic rise in popularity it has faced an endless stream of spammers, vandals, PR firms, and “clueless” newbies. In response, around 2009, English Wikipedia began directing new users to create new articles in the Draft namespace and seek review before publication. This new process, Articles for Creation (AfC), was meant to ensure that new articles measure up to Wikipedia standards for notability and verifiability, [2] but it can be a frustrating and discouraging process for newcomers. This is indicated by a low percentage of drafts that are reviewed (29.4% of AfC drafts are never submitted for review. 11% of reviews take longer than a week) and a falloff in new article longevity. Since AfC was introduced (2007), the proportion of articles of each new page creator that survive at least 30 days has declined. One reason for this is that draft articles are hidden from potential collaborators. It also suggests that the review process is often slow and confusing.

Implications

Overall, we should create better tools for the community to review new articles with, increase the percentage of drafts that are reviewed and decrease the number of reviews that take a week or longer to happen. By improving these metrics, we will improve newcomer productivity.

What Experts Say

Wikipedia has a very strong need for better tools to deal with vandalism and sock-puppeting. Though the means to block such malicious actors exist, they are easy to evade, especially for dedicated vandals who know how to use proxies. [7] Current proxy blocking tools aren’t very effective, especially in combating zombie proxies. [8]

Community-built tools are important for

moderators, but have lots of problems including lack of consistent or ongoing developer support, being tailored to one wiki and with generally steep learning curves for installation and use.

Discussion-based workflows (see Articles for Deletion, [9] Bot Request for Approval, [10] Article Peer Review, [11] for example, are often overly complicated and would benefit from dedicated interfaces to streamline their processes..

Generally, this points to the need for clearer guidelines for collaborating with the community on tool development to make sure we aren’t building things in isolation. Further, they must be built in close collaboration with the community in order to be successful. One key requirement here is consistent communication: project managers should provide weekly updates for the community, post screenshots, ask questions, and respond to feedback.

Implications

The top priority for moderators’ tools should be counter-vandalism tools, particularly automated proxy-blocking tools.

Later, we can focus on the tools used to facilitate discussion-based workflows like Articles for Deletion.

Decision Quality [12]

Because disagreements and conflicts on Wikipedia are frequent and inevitable, effective decision-making and conflict resolution processes are essential to community health.

The Articles for Deletion (AfD) process is one of the core moderation processes on English Wikipedia. AfD works through a consensus process rather than voting, and thus far we have learned that many factors affect the quality of AfD decisions, for example the likelihood that they will be reversed later. On a broader level it has been demonstrated that larger groups of participants make better decisions, but that the degree of participant experience makes a difference. Lacking understanding of Wikipedia norms and processes tends to yield worse decisions. Bias also plays a decisive role in decision quality: biased admins make worse decisions than impartial ones, and diverse

groups tend to reduce bias. At the same time, orchestrated bias does affect decision quality, especially when users with established positions are specifically recruited.

Implications

More tools should be built to facilitate more inclusive participation in moderation processes such as Articles for Deletion since this will likely improve the quality of the decisions. These tools should provide onboarding support and education for newcomers so that they can participate constructively.

Notes

1 Wikitext https://meta.wikimedia.org/wiki/Help:HTML_in_wikitext

2 Articles for Creation (AfC), just one example from English Wikipedia, is a system for submitting, reviewing, and publishing draft articles. In order to participate as a reviewer in AfC, you have to install a special user script, be granted a special right through a unique vetting process, and use several obscure templates and categories. https://en.wikipedia.org/wiki/Wikipedia:WikiProject_Articles_for_creation

3 Lam, Shyong K., et al. (2011). "WP:Clubhouse? An Exploration of Wikipedia's Gender Imbalance", WikiSym '11.

4 Purtil, Corinne; Schlanger, Zoë (October 2, 2018). "Wikipedia rejected an entry on a Nobel Prize winner because she wasn't famous enough", Quartz.

5 Schneider, Jodi, et al. (2014). "Accept, decline, postpone: How newcomer productivity is reduced in English Wikipedia by pre-publication review", OpenSym '14.

6 Code review procedures, documentation standards, etc.

7 Open Proxies <https://en.wikipedia.org/wiki/>

Wikipedia:Open_proxies

8 A zombie proxy is a proxy server being run on a computer that has been compromised by hackers.

9 Articles for Deletion https://en.wikipedia.org/wiki/Wikipedia:WikiProject_Articles_for_creation

10 Bot Request for Approval https://en.wikipedia.org/wiki/Wikipedia:Bots/Requests_for_approval

11 Article Peer Review https://en.wikipedia.org/wiki/Wikipedia:Peer_review/guidelines

12 Lam, Karim and Riedl (2010) "The Effects of Group Composition on Decision Quality in a Social Production Community" <https://dl.acm.org/citation.cfm?id=1880083>

Platform agnosticism

“Stroll through Sanlitun, a bustling neighborhood in Beijing filled with tourists, karaoke bars, and luxury shops, and you’ll see plenty of people using the latest smartphones from Apple, Samsung, or Xiaomi. Look closely, however, and you might notice some of them ignoring the touch screens on these devices in favor of something much more efficient and intuitive: their voice.” [1]

The Chinese language, as many other languages, was not built for typing tiny letters on a small screen. But that’s okay because technology, as it usually does given large-enough demand, is making its way around such initial difficulties. In this particular case, the answer might be voice search, alongside AI and new messaging paradigms.

This is just one example of how growing populations are coming to the internet with new needs, new languages, and new modes of expression. What’s certain is that their arrival will change the fabric of and forms in which knowledge is created, shared, and used. As internet access and usage rises in growing economies, the internet will become a more diverse place and platforms will be required to adapt to the needs and motivations of these new users.

If Wikimedia projects want to be participate in this growth and to “break down the social, political, and technical barriers preventing people from accessing and contributing to free knowledge”, we must ensure adaptability to any platform or mode of usage.

Yet predicting trends can be tricky and the risks that have prevented us from being at the forefront of technical innovation so far still apply. Unlike Google, who have the resources to try to do everything-everywhere-all-the-time. We don’t have the luxury or expertise to take large risks, especially if they’re not initiated from

the the ground-up (i.e. from our communities).

For us, ubiquity means adaptation - skipping the guessing game of what will be big in the future, investing in the needs of our current and potential communities to make sure our content is available for use in any future trend and for presentation on any device.

To support the goal of ubiquity we must focus on re-structuring our content so that it can easily be repurposed, remixed and repackaged by us, our communities, or other platforms. For us, structured Wikipedia content could significantly content porosity between our projects over time - facilitating use cases like a reader’s smooth transition from the Wikipedia article on Istanbul, to the Wikivoyage guide, and onward to related media about the city from Commons. Structured content would also support the establishment of global and customizable modular templates for articles, portals and projects. Standardized formats for the subcomponents of these experiences (such as sections, ideas or themes) via well-documented Wikipedia APIs—ie. the means of retrieving knowledge in whole or part—would support both non-Wikipedia platforms and future Wikimedia uses. A more structured content API platform would also make it easier for our diverse communities to generate the tools they feel they need - while maintaining consistent and reliable standards, and that work smoothly across the entire Wikimedia platform (see also Tools: For Developers).

Finally, as we consider this issue of ubiquity at the intersection of Wikipedia and its consumers, structured content would relieve us from the requirement to anticipate, monitor or otherwise be directly aware of how all populations in all emerging economies are developing their own unique relationships to the internet.

Syndication

“Wikipedia content appears to play a substantially more important role in the Internet ecosystem than anticipated, with other websites having critical dependencies on Wikipedia content.”

“Google becomes a worse search engine for many queries when it cannot surface Wikipedia content (e.g. click-through rates on results pages drop significantly) and the importance of Wikipedia content is likely greater than many improvements to search algorithms. Our results also highlight Google’s critical role in providing readership to Wikipedia. However, we also found evidence that this mutually beneficial relationship is in jeopardy: changes Google has made to its search results that involve directly surfacing Wikipedia content are significantly reducing traffic to Wikipedia.” [2]

So far, Wikipedia’s relationship with Google has been fairly symbiotic. We provide a trusted source they can show at the top of the page; they provide an increase in pageviews and, in turn, an increase in donations, in new editors, and in the continued creation of quality content they can then show to users. Everybody wins and information is distributed freely.

Yet exposing more information outside of the site, such as in Google’s knowledge panels, has decreased pageviews to Wikipedia. It is unfortunate that this is an issue. While we still met our goal of providing the information a reader sought without the direct traffic to our sites, we face not only a decrease in funds, but eventually a decrease in quality. Potential editors never see the site, let alone have a way to contribute, and current editors have less motivation to continue writing. Over time, we’re in trouble.

But, so is Google. The study quoted above clearly shows that Google is a worse search engine in a world without Wikipedia. Wikipedia’s importance is so large that the “mere presence of Wikipedia links may have an effect

approximately 80 times larger than the difference between a good ranker algorithm and bad one (for many queries)”. Similar patterns have been found for other online websites such as Reddit and StackOverflow, where Wikipedia content is widely shared.

Thus we find ourselves in an odd paradox where our current level of ubiquity is also a potential threat. One option would be to take a purely defensive stance and work towards preventing any information from usage outside of the site. Needless to say that that goes directly against the free-culture underpinnings of our movement, as well as our licensing. The other option would be to take syndication for granted - to imagine our content spread throughout the fabric of the internet, and shift our content creation and revenue model to such a future.

We need to open or deepen conversations with our partners, to provide them with insights into our side of the relationship. Being able to present them with our perspectives, such as those outlined in the previous sections will make it easier for them to respond to this more nuanced recognition of our interdependency. Larger institutions in particular must be made aware of the financial, legal, trust and cognitive dimensions of a relationship where they are getting a tremendous amount of value for no cost while putting the sustainability of that resource at risk. For example, our partners need to be more aware of information attribution issues: sampling our content without attribution that links to its full context not only negatively affects Wikipedia pageviews but potentially diminishes its functional value (i.e. outside the context of the community that can vouch for/dispute its veracity).

This recognition of the mutual downside is a potential opportunity for deepening our relationship with these high traffic drivers. Just as we now have a process for reviewing and adopting potential affiliate chapters, so too could we institute a model of official corporate affiliation with Wikipedia (e.g. “Google, an official partner of Wikipedia”), that makes that partner an official sponsor of the Open Knowledge Movement, according to some

mutually agreeable terms. A similar concern was voiced in recent research with regard to GLAM partners - that we have no way of bringing them into the fold in an official manner... “we can’t even provide them with a logo to use on their website”.[3] In this way we have the potential to amplify the “building a better world” missions of, for example, Apple and Google... and even to elevate the “don’t shoot the messenger” vibe of Reddit. Providing current and future partners with access to structured content, including contribution actions, via an API would support more symbiotic relationships, and open the door to creating workflows of contribution from other places where our content is used.

Content Relevance

“In the English Wikipedia, articles of strong insufficient quality alone receive close to half of the pageviews, and in the Russian Wikipedia, they receive more than half.” [2]

For our projects to be ubiquitous, we must provide relevant content to all of our users. Not all wikis are the same, nor do they grow in similar fashions and users of different projects have widely varied motivations for reading. For example, our research shows that readers in Western-language Wikipedias are more likely to focus on quick-fact information whereas the speakers of languages in growing economies are more likely to use Wikipedia for deeper learning and for work or school purposes. To be able to cater to the needs of individual wikis or groups of wikis, we must be able to distinguish their needs. Features that might work great on one target audience, might not work for another. Similarly, content that might be notable for a particular community, might not be for another. Focusing on targeting our work to match our unique audiences as well as providing them with the tools to build according to their needs will help us cover the entire range of content that our current and future readers will require. Only by analyzing the needs of readers, editors, and moderators can we address imbalances

in projects which constrain their growth. For example, knowing which Wiki projects may have quality content but low readership, or a high volume of low quality content, would targeted interventions toward more sustainable approaches to growth. Achieving that sustainability will mean assisting projects according to their specific needs. Being able to model the extent of a wiki’s content gaps along with nuances of its editor retention history would allow us to more effectively focus on the factors that limit that project’s ability to scale. We will need new tools to do so, based on a foundation of structured content and communication.

1. Wikitext https://meta.wikimedia.org/wiki/Help:HTML_in_wikitext
2. Articles for Creation (AfC), just one example from English Wikipedia, is a system for submitting, reviewing, and publishing draft articles. In order to participate as a reviewer in AfC, you have to install a special user script, be granted a special right though a unique vetting process, and use several obscure templates and categories. https://en.wikipedia.org/wiki/Wikipedia:WikiProject_Articles_for_creation
3. Lam, Shyong K., et al. (2011). “WP:Clubhouse? An Exploration of Wikipedia’s Gender Imbalance”, WikiSym ’11.
4. Purtill, Corinne; Schlanger, Zoë (October 2, 2018). “Wikipedia rejected an entry on a Nobel Prize winner because she wasn’t famous enough”, Quartz.
5. Schneider, Jodi, et al. (2014). “Accept, decline, postpone: How newcomer productivity is reduced in English Wikipedia by pre-publication review”, OpenSym ’14.
6. Code review procedures, documentation standards, etc.
7. Open Proxies https://en.wikipedia.org/wiki/Wikipedia:Open_proxies
8. A zombie proxy is a proxy server being run on a computer that has been compromised by hackers.

Ban, Perc and Levnajic (2017) “Robust clustering of languages across Wikipedia growth”

<http://rsos.royalsocietypublishing.org/content/4/10/171217>

“Can NC-licensed content be used in the Wikipedia project? – No. Wikipedia contents are being used commercially.” https://meta.wikimedia.org/wiki/Free_knowledge_based_on_Creative_Commons_licenses#8._Can_NC-licensed_content_be_used_in_the_Wikipedia_project?_%E2%80%93_No._Wikipedia_contents_are_being_used_commercially

Creative commons non-commercial license https://meta.wikimedia.org/wiki/Free_knowledge_based_on_Creative_Commons_licenses#8

De Zuniga, Weeks and Ardevol-Abreu (2016) “Effects of the News-Finds-Me Perception in Communication: Social Media Use Implications for News Seeking and Learning About Politics” <https://onlinelibrary.wiley.com/doi/full/10.1111/jcc4.12185>

Dixon (2018) “Why Decentralization Matters”, <https://medium.com/s/story/why-decentralization-matters-5e3f79f7638e>

Econsultancy (2018) “Why we need to stop repeating the “50% by 2020 voice search prediction” <https://econsultancy.com/why-we-need-to-stop-repeating-the-50-by-2020-voice-search-prediction/>

Guazzini, Vilone, Donati, Nardi and Levnajic (2015) “Modeling crowdsourcing as collective problem solving” <https://www.nature.com/articles/srep16557>

Hagedom et. al (2011) “Creative Commons licenses and the non-commercial condition: Implications for the re-use of biodiversity

information” <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3234435/>

Halfaker et. al. (2018) “The Rise and Decline of an Open Collaboration System: How Wikipedia’s reaction to popularity is causing its decline” https://www-users.cs.umn.edu/~halfaker/publications/The_Rise_and_Decline/halfaker13rise-preprint.pdf

Jacobs (1961) “The Death and Life of Great American Cities”, Random House

Jun, Lee and Jeong (2016) “Early onset of structural inequality in the formation of collaborative knowledge, Wikipedia” <https://arxiv.org/abs/1610.06006>

Jun, Lee and Jeong (2015) “Intellectual interchanges in the history of the massive online open-editing encyclopedia, Wikipedia” <https://arxiv.org/abs/1510.06092>

Kittur, Pendleton and Kraut, (2009) “Herding the Cats: The Influence of Groups In Coordinating Peer Production” <http://citeseerx.ist.psu.edu/viewdoc/>

