

content curation

Content curation is the aspect of wiki activity related to editing, refining, and cleaning up content that has been generated. The Wikimedia movement's ambitious aspiration to make the sum of all knowledge available to everyone in the world means that the movement has a tremendous amount of work to do with respect to making judgments about what information belongs, and how to organize, phrase, and cite it. Most of the hundreds of languages in the world have Wikipedias with less than 10% the number of articles that English Wikipedia has, and even the largest Wikipedias have serious gaps in terms of the depth of their articles, and the subject matter covered by their articles. As all that content gets added, the curation workload will increase beyond what humans are capable of doing.

Augmentation is a potential pathway to curating the massive amount of information needed in the Wikimedia projects. By applying algorithms and artificial intelligence in the right ways, human editors can be aided in making the most important judgments about the content in the wikis, allowing the content to be well-organized and reliable. This kind of human-machine partnership is not new in the wikis. Tools like [Twinkle](#) and [Huggle](#) have been helping to automate the tasks of reviewing recent changes and patrolling for vandalism since 2007 and 2008. ClueBot has been independently reverting vandalism since 2011. And in more recent history, [ORES machine learning models](#) have begun to surface the edits and pages most in need of attention.

As humans and machines work together to curate content, we can think about that interaction on a spectrum of how much work the human editor does and how much work the machine does. In some scenarios, the machine may just direct human attention to important curation needs. In other scenarios, the human may review tasks completed (e.g.

edits reverted) by an algorithm. This paper explores some specific examples of content curation activities that can exist in the future, drawing from all along the spectrum of the human-machine partnership.

Because bias and unfairness already exist in the contents of the Wikimedia projects, algorithms have the potential to magnify and exacerbate those problems. The Wikimedia movement should confront this with the same principles that have led to our success in the past: transparency and the ability for anyone to contribute.

Sections

[Sections](#)

[Augmentation strategy summary](#)

[Definition of content curation](#)

[Aspiration](#)

[Augmented content curation](#)

[Content curation strategy](#)

[Notes](#)

[Sources](#)

Augmentation strategy summary

This is a summary of the [overall strategy](#) for augmentation, which this document applies to the specific aspect of content generation.

In order to meet our movement's goal of making all the world's information available to everyone, we have more work to do than human editors can do alone. We need help in the form of augmentation, which is when humans and algorithms work together. Though augmentation in the wikis is not new, it will be a growing part of the future of the wikis. To ensure that the contributions made by algorithms are productive, unbiased, and fair, we will need to stick to our movement's principles of openness, transparency, and the ability for anyone to contribute. We should build closed-loop infrastructure and interfaces that allow anyone to contribute new algorithms, and for even non-technical editors to participate in training and tuning those algorithms. These principles would apply to all types of augmentation, whether it is in the aspect of content generation, content curation, or governing interactions between people.

Definition of content curation

Content curation is the aspect of wiki activity related to editing, refining, and cleaning up content that has been generated. This is in contrast with content generation, which is about adding new facts, writing, translations, or images to the wikis. It is also separate from governance,

which is about the interactions and communications between wiki editors.

Aspiration

The Wikimedia movement wants the sum of all knowledge to be available to everyone in the world. That ambitious goal means that the movement has a tremendous amount of work to do. Most of the hundreds of languages in the world have Wikipedias with less than 10% the number of articles that English Wikipedia has, and even the largest Wikipedias have serious gaps in terms of the depth of their articles, and the subject matter covered by their articles. Assembling all that knowledge is about more than just compiling it -- it means curating it: making judgments about what information belongs, and how to organize, phrase, and cite it. This makes the sum of all knowledge more accessible, and also makes it more trustworthy. There is going to be too much curation work in the future for humans to do it unassisted.

Augmented content curation

Augmentation is a potential pathway to curating the massive amount of information needed in the Wikimedia projects. By applying algorithms and artificial intelligence in the right ways, human editors can be aided in making the most important judgments about the content in the wikis, allowing the content to be well-organized and reliable. This kind of human-machine partnership is not new in the wikis. Tools like [Twinkle](#) and [Huggle](#) have been helping to automate the tasks of

reviewing recent changes and patrolling for vandalism since 2007 and 2008. [ClueBot](#) has been independently reverting vandalism since 2011. And in more recent history, [ORES machine learning models](#) have begun to surface the edits and pages most in need of attention.

As humans and machines work together to curate content, we can think about that interaction on a spectrum of how much work the human editor does and how much work the machine does. In some scenarios,

the machine may just direct human attention to important curation needs. In other scenarios, the human may review tasks completed (e.g. edits reverted) by an algorithm. Below are some specific examples of content curation activities that can exist in the future, drawing from all along the spectrum of the human-machine partnership.

Activity	Algorithm role	Human role
Identifying vandalism	Flag edits that are likely vandalism	Review the flagged edits
Identifying unsourced content	Flag parts of articles that make claims that do not appear to be sourced	Review the flagged content and correct or delete
Checking sources	Check sources for the facts cited from them	Review any flagged citations found by the algorithm
Identifying copyright violations	Check edits for whether their contents appear in other sources	Review flagged edits and correct or revert
Grouping tasks	Assemble related individual curation tasks into a prioritized queue	Use the queue to work more efficiently
Routing tasks	Route individual curation tasks to the editors who are most likely to be interested or capable	Receive and take action on routed tasks
Improving article composition	Make automatic improvements to the tone, style, grammar, and spelling of written content	Review automated improvements

Content curation strategy

The human-machine partnership scenarios described above are easier said than done, and the technical effort to build such algorithms is only part of the challenge. The more important challenges are how to build a technical framework and establish design principles to ensure that as algorithms play a growing role in content curation, they are a force for unbiased and fair curation.

Content curation as done by humans necessarily reflects the biases that the humans have. For instance, one human editor's preference for writing with a certain style might accidentally exclude edits done by members of other demographic backgrounds. Because bias and unfairness already exist in the contents and practices of the Wikimedia projects, algorithms that learn from human work have the potential to magnify and exacerbate those problems. The Wikimedia movement should confront this with the same principles that have led to our success in the past: transparency and the ability for anyone to contribute.

Concretely, for algorithms that participate in content curation, these things should be true to ensure transparency and the ability for anyone to contribute:

- > Algorithms should be able to be created and deployed by anyone. If content curation algorithms are only contributed by a select group, the way

they curate content will reflect the biases of that select group.

Example: if the creators of algorithms for identifying vandalism are all from the English-speaking world, the algorithms might do a poor job at identifying vandalism in other languages.

- > It should be clear what work is being done by algorithms and where those algorithms come from.

Example: if an edit is reverted by an algorithm, the user who initially made the edit should know that it was reverted by an algorithm. This will increase transparency, and potentially encourage human editors to improve the result.

- > We should always build “closed loop systems” with humans in the loop. This means that any content curation done by an algorithm should involve a human to edit, improve, and audit the work.

Example: in the Recent Changes feed, ORES models suggest edits that need attention, but do not automatically take action.

- > Our “closed loop systems” should allow corrections made to machine work to be fed back into the system to improve the algorithm going forward.

Example: the [Recent Changes feed](#) currently lacks a way for humans to flag

ORES suggestions that are incorrect so that ORES can be improved.

- > Shepherding, tuning, and training algorithms should be an important wiki role that non-technical editors can take on. Any editor should be able to wield, adjust, and provide data for improving augmentation. This work should “count” as wiki work, as actual edits, and editors should find their way to this augmentation niche.

Example: if the Recent Changes feed were to include a way to flag ORES judgments as incorrect, flagging those judgments should count as an edit.

As described in the overall augmentation theme strategy, the Wikimedia Foundation should do two concrete things to make the above possible:

1. Build an infrastructure platform for many people to contribute augmentation tools, coupled with Wikidata (or something like it) to serve as a repository of facts.
2. Provide interfaces that make it possible for non-technical editors to adjust and contribute to those tools.

Notes

[1] TBD

Sources

M. Miller and J. Klein : [Research and Insights](#), Other contributors¹: D. Garry, A. Halfaker, D. Horn, J. Katz, J. Minor, T. Negrin, M. Novotny, N. Pangarkar

Twinkle <https://en.wikipedia.org/wiki/Wikipedia:Twinkle>

Huggle <https://en.wikipedia.org/wiki/Wikipedia:Huggle>

Cluebot https://en.wikipedia.org/wiki/User:ClueBot_NG

ORES Machine Learning Models <https://www.mediawiki.org/wiki/ORES>

Recent Changes Feed <https://en.wikipedia.org/wiki/Special:RecentChanges>

¹ If your name was left off the list by mistake please contact JMinor or MNovotny