An analysis of Wikipedia's special tags and their implications for the nuanced spectrum between human edits and bot edits

Ashik AHAMED, Max WANG,
Amity Ramona MENTIS-CORT, Jeanna MATTHEWS

Clarkson University, Potsdam, New York, USA

ahameda@clarkson.edu, mawang@clarkson.edu,
mentism@clarkson.edu, jnm@clarkson.edu

Abstract: *Wikipedia articles are a key training source for Natural Language Processing (NLP) research. The platform's metadata records edit histories, including whether edits are made by humans or bots. A lesser-studied part of this metadata is the set of "special tags," which can mark edits made with semi-automated tools like Twinkle or fully automated bots such as Cewbot. This spectrum spans from light automation, such as spell checking, to heavy automation, such as AI-generated text later edited by humans. This paper examines how Wikipedia's special tags are used and explores their potential to trace the continuum between human and bot edits, an increasingly relevant issue with the rise of generative AI tools. We employ a structured SQL-based analysis through Quarry, complemented by manual inspection of revision histories to capture patterns beyond metadata. Using this hybrid approach, we move beyond surface-level records to uncover gaps between official data and editing behavior. The study demonstrates the value of combining SQL-based analysis with human verification to enhance the historical accuracy of Wikipedia's tag data. The resulting dataset supports machine-learning training, knowledge-graph construction, and analyses of cooperative online behavior*.

Keywords: Wikipedia, Special tags, Human–bot collaboration, Data science.

1. Introduction

Wikipedia is one of the most active and complex collaborative knowledge systems in the world, allowing real-time editing by a diverse pool of contributors: from casual contributors to professional experts to content custodians and even bots (Zheng et al., 2019). Unlike many corpora of text, Wikipedia maintains rich metadata about the historical editing process. This metadata can provide a unique window into the collaboration between humans and automated tools over time. With the rising prevalence of generative AI tools, it is increasingly important to ask when we should be including AI-generated content in the training set for new AI tools (Ferschke, Zesch & Gurevych, 2011; Halfaker et al., 2013). Automated tools are generally built by learning from data collected from humans, and replacing that human signal in training data with AI-generated content may not always be appropriate. For example, if you generate numerous articles for Arabic Wikipedia by translating from English with the purpose of making those articles more

https://doi.org/10.58503/icvl-v21y202601

accessible to Arabic speakers, that could be an appropriate use. However, prior studies on language-specific Wikipedia editions, such as Arabic Wikipedia, have shown that data derived from local corpora may not always represent the perspectives of native speakers (Alshahrani, Wali & Matthews, 2022). Using AI-generated content as the training data for new AI-generated tools can reach a point of diminishing returns, where the organic, effective signals are lost in an echo chamber of automated reinforcement of the dominant patterns.

We want to find effective models of collaboration between humans and automated tools. There are plenty of examples of collaborative models that give bad outcomes such as asking humans to jump in at the last moment when an automated system is faced with a difficult moral choice such as in the trolley problem (Elish, 2019). Similarly, humans rubber-stamping all automated suggestions is unlikely to give effective outcomes as humans in that position have little incentive to overrule the automation (Matthews, 2020). What we want is to build collaborative models that give good outcomes, and Wikipedia metadata including the Special Tags have the potential to give us a unique window to explore that. In this study, we analyze the Special:Tags ecosystem to recover and characterize patterns of tag adoption and use. We captured a snapshot of the Special:Tags page and then queried Wikimedia replica databases to extract usernames, timestamps, and diffs for tagged edits to augment this page with additional metadata. We also validated the first use for each tag and categorized tags by degree of automation (manual, semi-automated, automated). In this paper, we describe our methods, present results (including an illustrative example using the *Twinkle* tag), and discuss implications for human–automation collaboration on Wikipedia.

The MediaWiki maintains a collection of special tags, defined within Wikipedia itself on their Special:Tags page, which offers a centralized, up-to-date summary of all defined tags used on Wikipedia. While the exact number can vary over time, at the time of this work, there were roughly 300 special tags listed. Basic metadata is displayed for each tag, such as the number of edits tagged at the time of viewing, the source (core MediaWiki, extensions, or external tools), and the tag's activation status. This number of edits is dynamic and constantly shifting as users all over the world make new edits using the tags. While the *Special:Tags* page provides a useful overview of existing tags, it is important to understand how these tags are actually created and recorded. In MediaWiki, tags originate from three main sources. First, some tags are added automatically by the MediaWiki software itself to capture details of the editing interface, such as whether an edit was made using the VisualEditor or source editor. These system-generated tags are typically consistent across different language editions of Wikipedia, as they are applied by the core MediaWiki software rather than individual communities (Geiger & Ribes, 2010; Steiner, 2014). Second, external tools and bots, such as Twinkle or AutoWikiBrowser, can define and apply tags when performing automated or semi-automated edits. Prior research shows that such tools often predate their formal tagging, resulting in older edits that lack corresponding metadata (Geiger, 2014;

Halfaker & Riedl, 2012). Some tags are inferred algorithmically rather than explicitly applied; for example, revert-related tags are assigned based on heuristic detection of reverted edits (Halfaker & Riedl, 2012). Finally, some tags are applied manually by editors. Together, these mechanisms make tags a powerful but imperfect source of information, as their completeness and reliability vary depending on their origin and time of introduction.

Special tags in MediaWiki record *how* an edit was made and serve multiple purposes. Some indicate routine maintenance—for example, HotCat for quick category fixes and other minor cleanup (e.g., fixing typos/formatting or updating/archiving dead links). Reverts are marked with mw-rollback or mw-undo. Other tags capture tool-assisted work, such as *Twinkle*, which supports semi-automated tasks like issuing warnings and reverting vandalism. Large-scale, repetitive updates are often carried out by bots (e.g., Cewbot); those edits may carry task-specific tags when applicable or simply be recorded under the bot flag. Prior work has used these tags to study sociotechnical coordination and governance on Wikipedia (Geiger & Ribes, 2010; Halfaker et al., 2013). Collectively, tags help researchers examine platform governance, content integrity, and the division of labor between humans, tools, and bots in peer-production environments. Although the current special tags page shows a snapshot of current tag activity, it does not include deeper historical metadata. Important details like the time a tag was first applied and the user (i.e., the first username) are left out. Furthermore, it doesn't differentiate between edits made by bots and human users, nor does it classify tags according to their degree of automation (manual, semi-automated, or fully automated). Additionally, it's unclear whether tags were used for system-level operations like file uploads or user account creation, or for content edits. In addition to the special tags page itself, each row includes a total number of tagged changes, which links to an interface to query recent changes with that particular tag. While useful, there are some limitations in this interface that make it difficult to use for thorough, long-term analysis. For example, when filtering edits by tag and user type (e.g., bots vs. humans), the platform only displays results from the last 30 days and restricts output to 500 entries. These limitations significantly hinder large-scale or historical studies of the tool usage and tag behaviour on Wikipedia.

To overcome these limitations, we employ a hybrid research methodology that combines automated data extraction with manual verification. First, we use Quarry, an open SQL platform that interfaces with important tables like change_tag, revision, actor, and page, to query Wikimedia's public replica databases. This enables us to pinpoint the exact location of the tag's application, whether it was applied to a particular article edit (like changing the content of a Wikipedia page) or to a system-level action caught in the site's logging system (like deleting pages, creating accounts, or uploading files). We can also identify the first recorded use of each tag and identify the user in question (either a human or a bot). Second, to complement this structured analysis, we manually reviewed the revision histories of a select group of users whose usernames were identified using the change_tag data.

We visited Wikipedia user contribution pages to examine editing activity associated with specific tools. Tools such as Twinkle were in active use long before their edits appeared in the change_tag table. This manual examination helped identify earlier instances of tool usage that predate the official introduction of corresponding tags into the tracking system. This paper contributes: (1) a reproducible pipeline (SQL + manual verification) to infer first-use and usage patterns across 301 tags; (2) a curated bot whitelist/blacklist that replaces naive '%bot%' heuristics; (3) a public dataset reporting counts by human vs. bot and by automation class (manual, semi-automated, automated) and (4) a discussion of how special tags reveal how humans and bots collaborate along a nuanced spectrum from light automation, such as spell checking, to heavy automation, such as AI-generated text later edited by humans.

We pull the list of defined tags from Special:Tags and analyze them with SQL on Wikimedia's public databases (Quarry), joining change_tag/ change_tag_def for tag IDs and names, revision/actor for edit details and usernames, and page for titles, while using logging for non-edit actions. For each tag, we find the earliest confirmed use, the associated user, a canonical diff URL (using rev_id and rev_parent_id), and total usage counts, splitting human vs. bot activity with a curated bot allow/deny list rather than naive "%bot%" name filters (Johnson et al., 2016). Finally, we classify tags by automation level: manual (human-only), semi-automated (human triggers a tool like *Twinkle*/*Huggle*/*AWB*), and automated (bot or server-driven without per-edit confirmation).

We applied this method to a snapshot of Special:Tags captured on May 9, 2025. This snapshot contained 301 distinct tags. Queries were run using June 12–16, 2025. For each tag, we computed total edits and broke them down by human vs. bot; we also identified the earliest validated use and recorded the user and diff link for that event. In Section 2, we describe how tags were classified and present the results of that classification. Section 3 provides a deeper dive into *Twinkle* and other common tags, illustrating first-use validation and tag evolution. Section 4 concludes the paper with reflections on future work.

2. Tag classification

Following the identification of the initial users and timestamps for each of the 301 tags, we conducted a detailed first-use analysis. Each tag was classified based on whether its first recorded user was a human or a bot. After this step, we examined which users were responsible for introducing the largest number of tags overall, both among humans and bots. In addition to identifying first users, we also computed the total number of edits made by humans and bots for each tag using structured SQL queries. This helped us to compare the extent of participation between human editors and automated agents across different tag types. Finally, we performed a tag description analysis to understand the purpose and function of each tag. Using the descriptions provided on the *Special:Tags* page, we categorized every tag into one of three automation classes: *Manual*, *Semi-Automated*, or *Automated*.

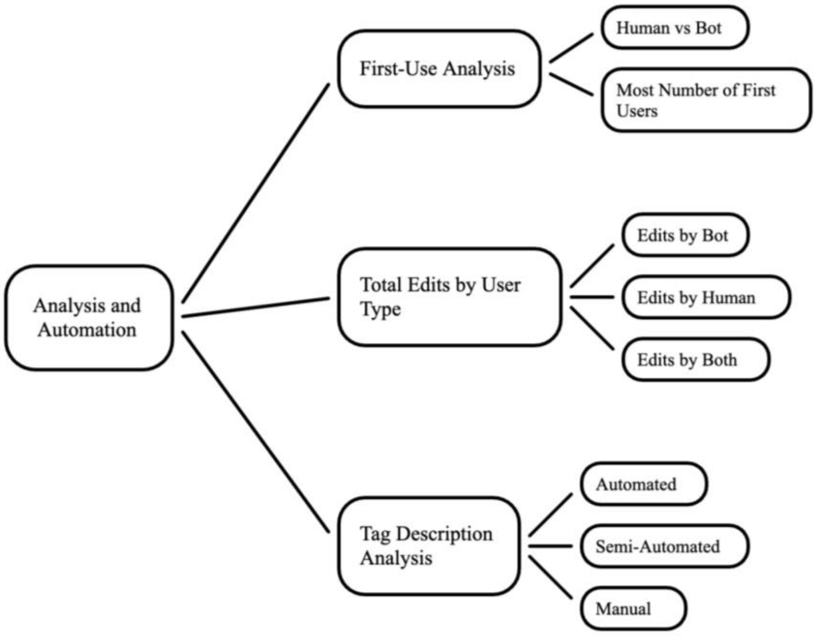

Figure 1. Overview of Analysis and Automation of User Interactions: Workflow from First-Use Analysis to Automation Classification

We also distinguish between human editors and bot editors. Distinguishing between human and bot users is helpful for analyzing the patterns of tag usage, especially for tools like *Twinkle* that are mainly used by humans, but bots can also use. Initially, we used a simple SQL filtering method by applying the condition LIKE '%bot%' to the actor_name field. Although this approach yielded results rapidly, it generated numerous false positives because it identified any username that contained the substring "bot"—even if Wikipedia did not officially classify the account as a bot. We performed a focused comparison using the *Twinkle* tag by determining the first five distinct users who triggered the *Twinkle* tag in the change_tag table in order to measure the amount of error caused by the %bot% filter. To confirm their classification, we next looked to see if these users were listed on Wikipedia's official bot list. Five usernames were identified by this comparison: Robotnick2, Simbotin, RobotGoggles, Ahechbot (borderline), and NinjaRobotPirate. Ahechbot was the only one that was formally registered as a bot. The other four were either unverified accounts or human editors. On the other hand, we found that only two official bots are being used by *Twinkle*, and they are Helsabt and Ahechtbot. In assessing tool usage like *Twinkle*, this result shows that four out of five usernames (80%) detected by the %bot% filter were false positives, underscoring the need for more precise classification techniques. In our study, we improved bot detection by using a curated whitelist instead of substring matching.

We compiled the full set of active and inactive bots from Wikipedia's "Category: All Wikipedia bots" (2,000+ accounts), then ran SQL on the Wikimedia Quarry replicas. For each first editor, we labeled the edit bot if actor.actor_name was in that whitelist (IN (...)) and human if it was not (NOT IN (...)). This method sharply reduced false positives and yielded more reliable user-type attribution.

Based on our findings, we divide the 301 tags into different usage groups. Figure 2 shows the results. The majority of tags—194—were initially used only by human editors, and this emphasizes how important manual contribution is to tag adoption. Only two tags, like "Cewbot" and "MediaWiki message delivery" were initially utilized exclusively by bots; these tags are both well-known for carrying out completely automated maintenance tasks. We found 83 tags that demonstrated mixed usage, which means that both human users and bots used them initially at different times, in addition to these exclusive-use categories. This overlap demonstrates a degree of convergence in the ways that bots and humans interact with Wikipedia's tagging infrastructure, and reflects the common use of specific tools by both user types.

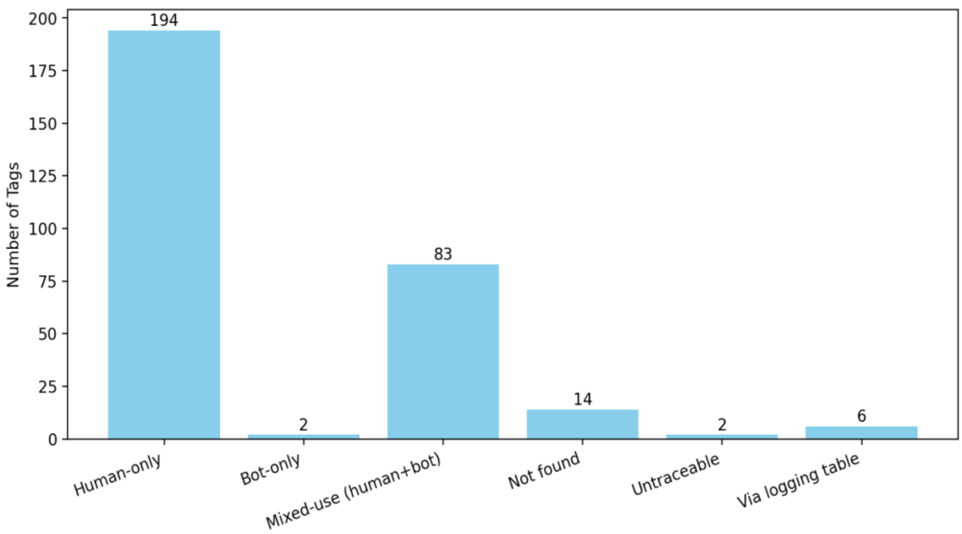

Figure 2. Distribution of Wikipedia Tags by Usage Category
This visualization highlights the dominant role of human contributors in initiating tags, along with other tag usage patterns across the Wikipedia ecosystem.

Additionally, we identified 14 tags that were never observed in either revision or logging data and categorized them as *Not Found*, indicating that, although they were defined in the system, they were never used for editing. Two additional tags (OAuth CID: 951 and OAuth CID: 1013) were classified as *Untraceable*, as it was not possible to reliably identify who used them or when. Although OAuth CID: 1013 appeared twice in the change_tag table and was associated with page revisions rather than system actions, attempts to retrieve the corresponding revision IDs (e.g., 823744470 and 823744540) returned the message

"This revision does not exist," indicating that the edits were likely deleted or suppressed. The same issue occurred for OAuth CID: 951. As a result, despite being tied to genuine edits, neither tag could be attributed to a human or bot editor. Finally, we identified six tags that appeared exclusively in the logging table and were used for system-level or administrative operations rather than content edits. These included two tags for creating user accounts via email (newusers/byemail), one tag for page creation (create/create), and three tags for file uploads (upload/upload). These tags capture behind-the-scenes operations that support Wikipedia's infrastructure but do not directly alter article content.

Focusing on tags with observable first-use events in revision data, we analyze when tags were first introduced and whether their initial use was by human editors or bots. This temporal analysis allows us to examine patterns in tag adoption over time and to compare early human and bot involvement in the tagging ecosystem. After excluding 22 tags that were categorized as *Not Found*, *Untraceable*, or logging-only, we retained 279 tags for further analysis. As shown in Figure 2, 194 of these tags were first used exclusively by human editors, 2 were first used exclusively by bots, and 83 were used by both humans and bots at different points in time.

We retrieved the total number of edits associated with each of the 279 tags by running our custom SQL queries on the Quarry platform individually for each tag. These queries allowed us to calculate not only the overall number of edits per tag, but also to separate those edits by user type. These things help us to distinguish between contributions made by human editors and those made by bot accounts. This breakdown provided insight into how different types of users (human and bot) interact with specific tags, and we also get an idea about whether certain tags are more commonly associated with automated tools or manual workflows.

Figure 3 shows the use of a single tag, *mw-reverted*, over time. The y-axis represents the cumulative number of edits, separated into edits made by human editors, bot editors, and all edits combined (human + bot + uncategorized). Edits are classified as uncategorized when the associated accounts cannot be reliably attributed, such as in cases of renamed, deleted, or suppressed accounts. The first recorded use of the tag by a human editor occurred in 2001, while the first recorded use by a bot occurred later, in 2004. From 2001 through 2018, usage remained minimal across all categories. Beginning in 2019, however, tag usage increased sharply, with human contributions growing substantially faster than bot contributions.

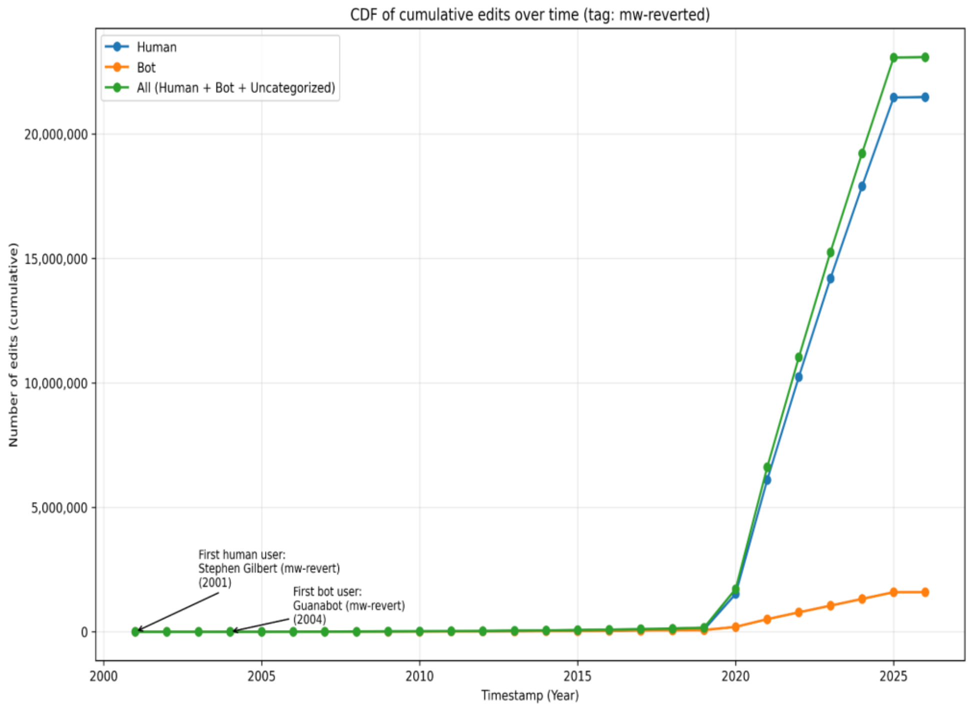

Figure 3. CDF of Cumulative Edits Over Time for the mw-reverted Tag, Showing Human-Only, Bot-Only, and Combined (Human + Bot + Uncategorized) Contributions from 2001 to 2026

3. Deeper dive: Twinkle and other common tags

We also decided to take a closer look at five specific tags so that we can provide a more grounded, detailed understanding of how tagging actually works in practice. These five tags correspond to some of the most frequently used tools on Wikipedia: *Twinkle*, *Huggle*, *AWB* (AutoWikiBrowser), *mw-replace*, and *STiki*. They are important for semi-automated editing on Wikipedia. Most of them need a human to approve each edit, but some, like *mw-replace*, can make changes automatically based on set rules. still require a human to initiate, configure, and supervise its operation. That puts it between manual tools and fully autonomous bots. By studying how these tools are used, we can better understand how Wikipedia content is created and controlled. Our goal was to check how well Wikipedia's tag system shows the real history of editing tools. We did this by comparing two things: the first time a tool is recorded in the change_tag table (using SQL from Quarry) and the actual first time it was used, which we found by manually checking edit summaries and user contributions. In the following section, each tag is discussed in detail.

In this section, to illustrate some nuances in using Wikipedia Special Tags, we present a concrete case study of the *Twinkle* tag. *Twinkle* is considered to be one of the most widely used semi-automated tools on Wikipedia. The *Twinkle* tag indicates use of the *Twinkle* tool, a JavaScript-based application developed in 2007 by AzaToth. The goal of this tool was to help editors with common tasks like sending user warnings, reverting vandalism, and tagging pages. *Twinkle* saves editors time and effort by simplifying these repetitive tasks so they can apply templates and issue notices with just a few clicks. Even though *Twinkle* has been widely used for many years, it wasn't officially tracked in Wikipedia's change tagging system until September 5, 2020. To get a clearer picture of how *Twinkle* was used before that, we combined SQL-based analysis with manual checks of edit histories. This mixed approach helped us trace how the tool was actually used over time, spot gaps between the metadata and real edits, and distinguish between edits made by humans and bots. What we found was striking: *Twinkle* had been actively used for more than a decade before it ever showed up in Wikimedia's tagging system. This points to a major delay in how tool-assisted edits were being recorded—showing that metadata alone doesn't always tell the full story. Based on our analysis of the change_tag table, we found that *Twinkle* has been used in a total of 9,788,150 edits. Of these, the overwhelming majority number is 7,173,736 edits that were performed by human editors, reaffirming *Twinkle's* role as a tool designed for active contributors. In contrast, only 12 edits were attributed to official bot accounts. Other edits fall under "uncategorized". This large disparity underscores *Twinkle's* design focus: it is built to assist human workflows such as reverting vandalism, tagging pages, and issuing user warnings, not for automated editing.

We also found disparity between the tool's initial recorded use in the change_tag system. As there is a gap between its real first appearance in the Wikipedia edit history. The first time Bruce1ee used the *Twinkl*e tag was on September 5, 2020. We got these results by using our custom SQL query from the Quarry website. However, through manual investigation of edit summaries and user contribution pages, we discovered that *Twinkle* had been actively used as early as 1 February 2007 by developer, AzaToth. It was used to edit the article *Animism*. That's a gap of more than 13 years between the actual use and when it was officially recorded. A similar inconsistency was found in bot usage. The first recorded bot use of *Twinkle* was by Helsabt on 12 September 2021, as per the change_tag table. However, on 20 December 2014, it was first used, but this thing was not recorded on the change_tag table. These disparities show how the change_tag system may greatly underestimate the historical timeline of tool adoption and misses early tool activity.

Table 1. First actual use, first recorded use, and delay between them for selected wikipedia tags

TAG	**First recorded use (Date and username)**	**First actual use (Date and username)**	**Days Between**
Twinkle	2020-09-05 Brucelee	2007-02-01 AzaToth	4,976
AWB	2018-01-17 Reedy	2006-05-18 Reedy	4,262
Huggle	2015-09-10 Petrb	Equazcion 2008-03-11	2,739
STiki	West-andrew.g 2018-04-11	West-andrew.g 2010-02-18	2,974
mw-replace	2017-12-07 ltitanthompson	PRehse 2006-11-25	4,030

We found a similar gap for four other tools: *Huggle*, *AWB* (AutoWikiBrowser), *mw-replace*, and *STiki*. In all these cases, the tools were used years before they were tracked with a tag. This shows that metadata alone doesn't give the full picture. Many tools were in active use long before they were added to the tagging system. So, relying only on SQL tag data can miss important history. That's why we used a hybrid method—both manual checking and database queries—to build a completer and more accurate timeline of when each tool was actually used on Wikipedia.

By drilling down into these five tools, we uncovered not just how these systems are used—but also where the metadata system falls short. Many tools were active long before their usage was officially tracked, and others are inconsistently tagged depending on how or by whom they were used. Table 1 summarizes the mismatch between actual first use and first use recorded in the change tag system, revealing clear evidence that the formal tracking system lags behind reality. This gap emphasizes the need for improved tagging infrastructure, especially as Wikipedia becomes a foundational source for datasets, LLMs, and real-time information systems.

Additionally, Wikimedia staff or administrator accounts, such as "TSevener (WMF)," made some of the earliest edits identified in our first-use data. These cases likely represent internal testing or early configuration rather than organic community adoption. For example, in the case of the tag *app-image-add-top*, the first use was by the account "TSevener (WMF)," which is an admin account. Recognizing such administrative activity provides a more nuanced interpretation of the data, distinguishing between genuine first use by the community and preliminary system testing. Future tagging studies could build on this finding by explicitly differentiating administrative and non-administrative activity to refine timelines of tool adoption.

4. Conclusion and future work

This study examined Wikipedia's *Special:Tags* system to better understand the nuanced continuum between human and automated contributions. Through SQL-based data extraction and manual verification, we found that tool usage on Wikipedia is more complex and historically underrepresented than metadata alone suggests. Tags such as *Twinkle, Huggle,* and *AWB* revealed multi-year gaps between actual tool use and their first recorded appearances in the change_tag table, underscoring the limits of relying solely on structured datasets. Our refined human–bot classification method also showed that common heuristics (like "%bot%") can produce up to 80% false positives, which points to the need for verified bot registries and improved labeling standards.

We need to develop richer and more transparent tagging frameworks in the future to effectively capture emerging forms of AI-assisted collaboration. We propose extending the tagging taxonomy to include indicators such as *llm-generated* (AI-authored text), *prompt-inserted* (AI-suggested edits), and *llm-verified* (human edits checked with AI). These additions would allow researchers and moderators to trace AI involvement with greater precision and accountability. As large language models become integral to editing and moderation, standardized tagging of human–AI interactions will be essential for preserving the reliability of Wikipedia's metadata and ensuring that future datasets reflect not just what was edited, but also how and by whom.

REFERENCES

Alshahrani, S., Wali, E. & Matthews, J. (2022) Learning from Arabic Corpora but Not Always from Arabic Speakers: A Case Study of the Arabic Wikipedia Editions. In: *Proceedings of the Seventh Arabic Natural Language Processing Workshop (WANLP), December 2022, Abu Dhabi, United Arab Emirates*. Stroudsburg, PA, Association for Computational Linguistics. pp. 361-371.

Elish, M. C. (2019) Moral crumple zones: Cautionary tales in human–robot interaction. *Engaging Science, Technology, and Society*. 5, 40–60. https://doi.org/10.17351/ests2019.260.

Ferschke, O., Zesch, T. & Gurevych, I. (2011) Wikipedia revision toolkit: Efficiently accessing Wikipedia's edit history. In: *Proceedings of the ACL-HLT 2011 System Demonstrations, June 19-24, 2011, Portland, Oregon*. Stroudsburg, PA, Association for Computational Linguistics. pp. 97–102.

Geiger, R. S. (2014) Bots, bespoke, code and the materiality of software platforms. *Information, Communication & Society*. 17(3), 342-356. https://doi.org/ 10.1080/1369118X.2013.873069.

Geiger, R. S. & Ribes, D. (2010) The work of sustaining order in Wikipedia: The banning of a vandal. In: *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work, February 6-10, 2010, Savannah, Georgia. New York, ACM*. pp. 117–126.

Halfaker, A. & Riedl, J. (2012) Bots and cyborgs: Wikipedia's immune system. *Computer.* 45(3), 79-82. https://doi.org/10.1109/MC.2012.82.

Halfaker, A., Geiger, R. S., Morgan, J. T. & Riedl, J. (2013) The rise and decline of an open collaboration system: How Wikipedia's reaction to popularity is causing its decline. *American Behavioral Scientist*. 57(5), 664–688. https://doi.org/10.1177/0002764212469365.

Johnson, I. L., Lin, Y., Li, T. J. J., Hall, A., Halfaker, A., Schöning, J. & Hecht, B. (2016) Not at home on the range: Peer production and the urban/rural divide. In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, May 7-12, 2016, San Jose, California*. New York, ACM. pp. 13-25.

Matthews, J. (2020) Patterns and anti-patterns, principles and pitfalls: Accountability and transparency in AI. *AI Magazine*. 41(1), 82–89. https://doi.org/10.1609/aimag.v41i1.5204.

Steiner, T. (2014) Bots vs. wikipedians, anons vs. logged-ins (redux) a global study of edit activity on wikipedia and wikidata. In: *Proceedings of the International Symposium on Open Collaboration, August 27-29, 2014, Berlin, Germany*. New York, ACM. pp. 1-7. https://doi.org/10.1145/2641580.2641613.

Zheng, L., Albano, C. M., Vora, N. M., Mai, F. & Nickerson, J. V. (2019) The roles bots play in Wikipedia. *Proceedings of the ACM on Human-Computer Interaction.* 3(CSCW). pp. 1-20. https://doi.org/10.1145/3359317.