Interaction Timeline User Interviews

Claudia Lo, Design Researcher, Anti-Harassment Tools Team. 24 Jan 2019

Previously, the Anti-Harassment Tools team developed the [Interaction Timeline tool](#), with the intention of helping administrators understand the sequence of events between two conflicting editors. These semi-structured user interviews were conducted in order to see what potential points of friction existed in the tool, and to better understand how administrators are currently using it.

The Interaction Timeline was initially created to help with cases of conflict, with edit wars and wikihounding as example use cases. It takes, as input, two editor usernames and a wikiproject, and optionally a start and end date. It outputs all edits chronologically, made on the specified wikiproject within the given date range, made by either editor on pages that they have both touched. For example, if User:Apples edits the pages "Pie", "Cake" and "Pancake", and User:Bananas edits the pages "Noodles", "Cake" and "Pancake", the tool would show their edits for the pages "Cake" and "Pancake". These edits are displayed in two columns side-by-side, with small red text noting the time interval between them, as can be seen in Figure 1. Clicking on them will load the diff in question.

For this study, we used a semi-structured interview consisting of a few questions to establish the participant's familiarity with the tool and with conflict resolution, and two short tasks. The first task was to ask the participant to find an example of inter-editor conflict, and to figure out what happened using the tools of their choice. Then, they were asked to do the same, but with the Interaction Timeline. Afterwards, participants were asked to compare the experience of using their preferred tools with Interaction Timeline, and give feedback on what they liked about the tool, and what they thought could be improved.
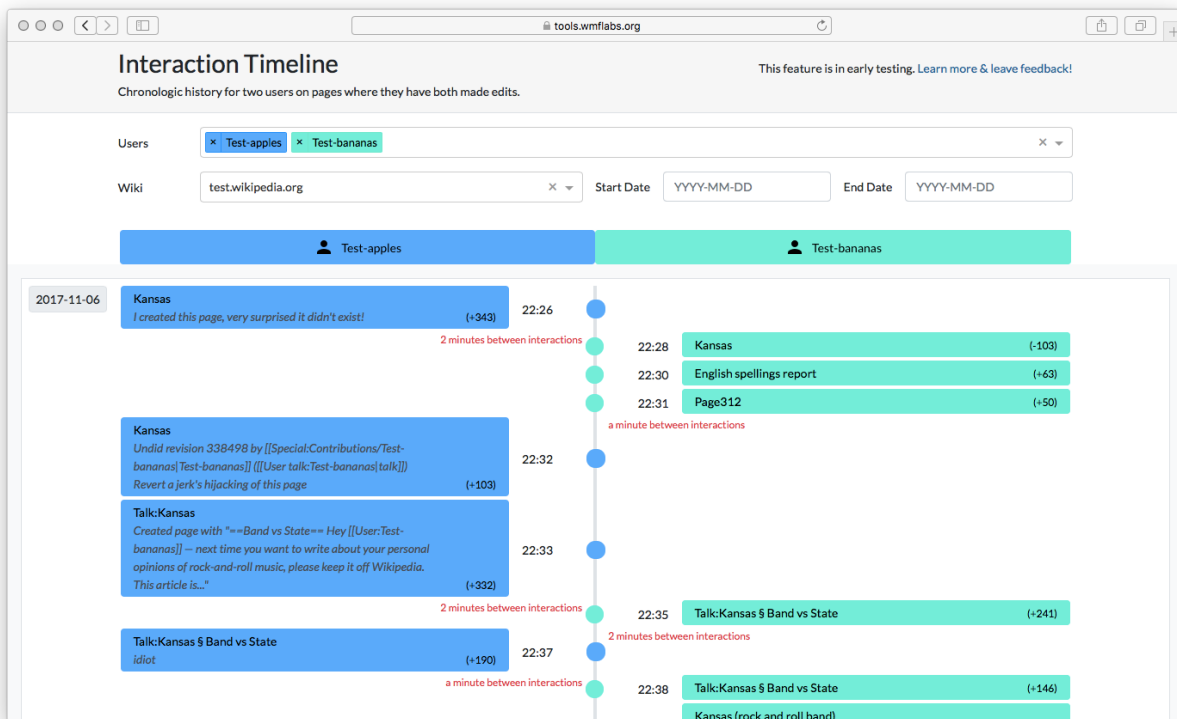
Fig. 1. A screenshot of the Interaction Timeline tool. Taken from the [Interaction Timeline page](#) on meta-wiki.

Participants were directly solicited via email, based on their involvement with on-wiki conflict resolution. The interviews were conducted over Google Hangouts, with participants sharing their screen. Both video and audio were recorded. Despite best efforts, this study had a very low response rate, and returned two participants.

Neither participant used the Interaction Timeline tool for its intended purpose. Participant 1 commonly used it for Sockpuppet Investigations (SPI), comparing the intersecting activity of an account with its suspected sockpuppet. Participant 2 often oversaw disputes between editors, but these disputes were rooted in content disagreements, rather than the conduct disputes for which the Interactive Timeline tool was developed. Consequently, Participant 2 was not familiar with this tool.

Common Themes

Both participants were generally positive about the tool, but noted a few key areas of improvement. One noted strength was the Interaction Timeline's presentation of information and its ability to compare activity on projects other than English Wikipedia. However, there were a few common stumbling blocks, notably some minor interface confusion, and difficulties in discovering it.

Both participants liked the two-column layout, finding it a logical and clear way to compare edits. Both of them intuitively figure out that clicking on each entry would open its relevant diff. Participant 2, despite their self-professed lack of experience in this work, immediately caught onto the red edit interval text, and later singled it out as a useful feature. Although neither participant specified a date range, this did not seem to negate its usefulness.

The name-entry field proved to be a slight point of friction. Participant 1 initially did not notice that it was limited to two names, and was puzzled as to why they could not paste in a third username for half a minute. Participant 2 attempted to comma-separate usernames entered manually, which meant that entering the second username became tricky as the system attempted to autocomplete an entry beginning with a comma. My speculation is that, since the text box for usernames appears to have room for multiple usernames, users do not notice the text saying that it accepts two names, or do not realize that the tool accepts exactly two names.

The other point of friction was in finding the tool. This was less of an issue for Participant 1, as they had created a custom page in their sandbox linking to tools they used, including the Interaction Timeline. However, Participant 2 assumed that this tool would be under the Gadgets page; upon not finding it there, they became stuck, and was unsure where to go next or even what to search. Compounding this issue is the fact that the main Interaction Timeline project page is hosted on meta-wiki, not English Wikipedia, where these participants would otherwise be.

One improvement both participants suggested was highlighting suspicious edit patterns. For example, highlighting edits of identical size in either direction could help identify suspicious patterns between two editors. In an edit war, multiple successive identically-sized edits in

different directions, done rapidly between two editors, is a sign of a string of reverts: User:A might make a +30 byte change, which User:B reverts, causing a -30 byte change, which User:A reinstates as a +30 change, and so on. For sockpuppetry, the suspicious behaviour would be two accounts, supposedly unrelated, making edits of the same size and direction, might indicate sockpuppet behaviour.

Takeaways

Both participants seemed to like using the Interaction Timeline tool and found its data presentation an improvement to other tools, or to simply scanning through edit summaries. However, given the low response rates, and the fact that of the two respondents, neither used it for its intended purpose, it seems that its impact has been rather low. Even though the tool was liked, it does not seem to provide enough of a benefit for Wikipedia's admins to adjust their workflows to incorporate it.

However, I believe that the API developed alongside this tool has potential, especially in a future reporting system to minimize the work a reporter has to do. This could open up future research into what kinds of information are best automatically captured by a reporting system in the course of making a report.

To reiterate, while this standalone tool was well received, it does not seem to have enough of an impact on its own to warrant significant time or resources spent in further development. Meanwhile, the API on which it runs could have useful future applications for Community Health work, and may be worth investigating in that aspect.