



Wikipedia Android app multilingual update

Post-release report

Mikhail Popov

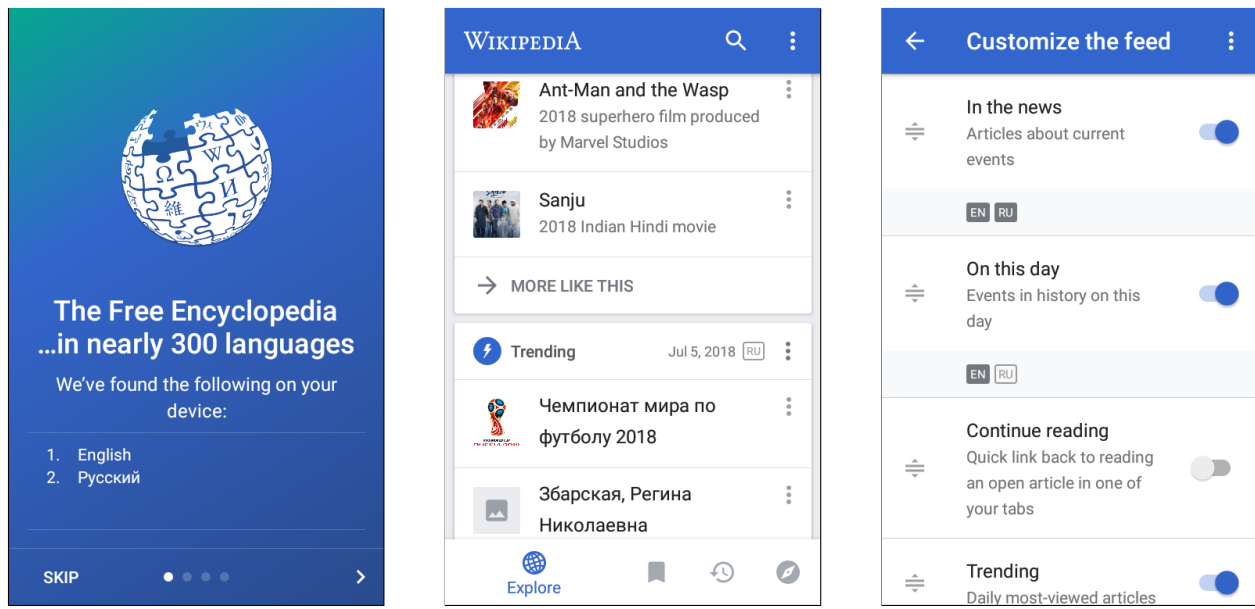
13 July 2018

Summary On 25 June 2018, the Android team released a major update to the Wikipedia app that improved and extended support for multilingual users. Previously, the app's users were limited to a single language and were required to switch between languages if they wished to see content in other languages they could read in. The update enables users to set multiple languages, to get parts of the Explore Feed in those languages, and switch between languages when searching for articles.

The multilingual update appears to be a useful one for many users, and majority of users who open the language settings end up having multiple languages set. So, presumably, they appreciate not having to switch between languages anymore. While we did not see much of a difference in sessions (average session length and average number of sessions per day), we did see an increased number of articles read per session among users who read in more than one language. Furthermore, with nearly 20% of multilingual users switching between languages when searching, we suspect having this ability makes searching more convenient.

Unfortunately, we need to be better at onboarding users to Feed customization because nobody is customizing Feed cards on a language-by-language basis. We suspect this has a negative effect on the usefulness of the Feed to users who read in 3+ languages, as those users were less likely to interact with language-specific cards but were more likely to interact with language-agnostic and English-only cards.

Open report: R code and RMarkdown source document are [publicly available](#) under a [Creative Commons Attribution-ShareAlike 4.0 International License](#).



(A) System and keyboard languages are detected and suggested during onboarding for new users.

(B) Feed can now include cards (such as most read articles) in multiple languages.

(C) Feed can be customized so certain cards show up in some languages but not others.

FIGURE 1: Screenshots of the Wikipedia Android app with the new, improved support for multilingual users.

Introduction

Following the research and design work of Rita Ho, the Android team released a major update to the Wikipedia app (available on [Google Play Store](#)) on 25 June 2018, which improved and extended support for multilingual users (see [T160567](#) for more information). The updated features include:

- Users can not set multiple languages (previously they could only set one and would need to switch)
 - Users are prompted during onboarding (Fig. 1A)
 - Users can change their language settings from search and settings screens
- The Explore Feed can contain content in multiple languages (Fig. 1B)
- Users can customize which Feed cards are enabled on a per-language basis (Fig. 1C)
- Users can switch between languages when searching for articles

The update also included analytics for the new and improved features (discussion and details documented in [T190931](#)). During onboarding (and later in settings) users can opt out of sharing usage data, so the analysis in this report relies on anonymous data from those users who chose to share. Furthermore, certain schemas/funnels use different sampling rates – some at 100%, some at 10%, and some at 1% – and in cases where the sampling rate is lower than 100% we refer to those users as a “random sample” or having been “randomly selected”.

The data analysis in this report focuses on user experience impacted by the improved support. The metrics in this report are largely broken down by language ability, forming three groups of users: monolingual, bilingual, and multilingual – those who have set one, two, and three or more languages, respectively.

Results

Language settings

Among a random sample of 1.4K users who **engaged with the language settings** (either changed or left them as-is) in Table 1, 27.8% are monolingual, 54.1% are bilingual, and 18.1% are multilingual (3 or more languages). **However**, of 29K randomly-selected users **who have used the new version of the app**, 91.6% are monolingual, 6.14% are bilingual, and 2.26% are multilingual (3 or more languages).

TABLE 1: Breakdown of users by number of languages they have in their settings.

	Number of languages				
	1	2	3	4	5+
Approximate % of users	27.8%	54.1%	12.6%	3.42%	2.11%

During onboarding, when the app detects languages from the user’s system and keyboard settings and suggests them to the user (see Fig. 1A), of 356 sampled users, 8.43% removed 1 or more languages, 73% added 1 or more languages, and 6.74% changed their primary language.

In Tables 2 and 3, we show the top languages and language combinations used by a random sample of users who have used the app at least once in the 2 weeks after the update. Table 2 lists the popularity of each language among users who read in 1–5+ languages. For example, we can see that Hindi is not used by itself and is only used in conjunction with 2 other languages, and Hindi is used by 2.1% of users who read in 3 languages. On the other hand, Swedish is only used by itself and is used by 1% of monolingual users.

Table 3 lists most popular language combinations. For example, German & Russian (at 12%) was the second most popular combination among users who read in 2 languages with neither being English. However, once we look at combinations which included English and at least 2 other languages, the combination of English, German, and Russian (not shown) is used by 2.7% of multilingual users (who read in English in addition to 2 or more other languages).

TABLE 2: Prevalence of languages among a random sample of users. Within each number of languages, the top 10 languages were selected and then those languages were combined into an overall list.

	Overall	Within users who read in this number of languages				
		1	2	3	4	5+
Arabic	0.9%	–	2.4%	2.5%	–	–
Chinese (Simplified)	1.1%	–	–	8.5%	11.9%	3.7%
Chinese (Traditional)	1.5%	–	1.9%	8.5%	11.7%	3.7%
Dutch	1.8%	1.9%	1.7%	–	–	2.6%
English	42.8%	43.9%	46.0%	30.1%	24.0%	13.8%
French	5.6%	5.6%	4.4%	5.6%	7.8%	9.7%
German	18.7%	20.7%	10.9%	8.4%	7.3%	6.3%
Hindi	0.3%	–	–	2.1%	–	–
Italian	5.2%	5.5%	3.7%	3.2%	4.9%	4.1%
Japanese	2.6%	2.6%	2.5%	–	4.2%	3.4%
Portuguese	1.4%	1.4%	–	–	1.8%	4.9%
Russian	4.1%	4.1%	3.9%	4.5%	2.2%	–
Spanish	5.2%	5.3%	4.3%	5.0%	6.0%	7.8%
Swedish	1.1%	1.0%	–	–	–	–

Reading behavior

Currently, when a user stops using the app for more than 30 minutes (with basically any interaction with the app counting as usage), then we consider the current session as over, and send the summary of that session once the user interacts with the app again. Looking at distributions of average session length per user by language ability in Fig. 2, they look largely similar, with 5 minutes being the most common average session length. However, distribution among bilingual users is bimodal with the additional peak around 2 minutes.

The average number of sessions per day in Fig. 3 also shows a similarity between the three language abilities, with the only differences being the top 1% of each group. The top 1% of monolingual users average at 14 or more sessions per day while bilingual and multilingual (3+ languages) users average around 9 or more sessions per day.

This is interesting in the context of Fig. 4 which shows median pages viewed per session. So while bilingual and multilingual users tended to have shorter sessions and fewer sessions per day, they read more articles during those sessions. For example, 59% more multilingual users than monolingual users read 5+ pages per session, and 26% fewer multilingual users than monolingual users do not read any articles (on average).

TABLE 3: Top 3 (unordered) language combinations. Among those who read 2+ languages, we categorized users into 4 groups based on whether they read in 2 or 3+ languages and whether the combination includes English.

Languages	Proportion of users within group
2 languages, not including English	
Chinese (Simplified), Chinese (Traditional)	13.3%
German, Russian	11.9%
Arabic, French	8.4%
3 or more languages, not including English	
Chinese (Simplified), Chinese (Traditional), German	19.6%
Chinese (Simplified), Chinese (Traditional), Japanese	7.8%
Chinese (Simplified), Chinese (Traditional), Spanish	5.9%
2 languages, including English	
English, German	20.3%
English, Spanish	8.3%
English, French	8.1%
3 or more languages, including English	
Chinese (Simplified), Chinese (Traditional), English	12.6%
Chinese (Simplified), Chinese (Traditional), English, Japanese	3.5%
English, French, German	2.8%

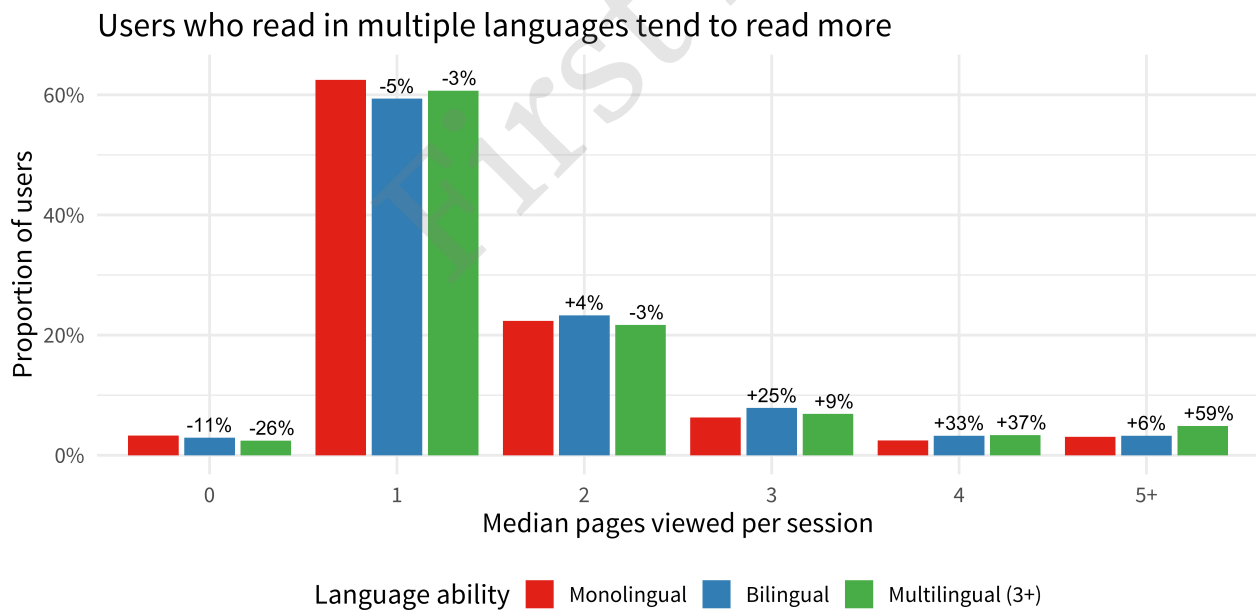


FIGURE 2: For each number of pages, the +/- percentages are relative to proportion of monolingual users who, on average, read that many pages per session.

In Figs. 5 and 6, we explore daily usage of the app by looking at number of days people used the app and the length of the longest streak of days they used the app consecutively. For the most part, monolingual users used the app more frequently (3 or more days) and were more likely to use app 2 or more days in a row than bilingual and multilingual users. An exception to this is that there were 12–14% more bilingual and multilingual users than monolingual users who used the app for 2 days, and there were 4% more bilingual users than monolingual users who used the app consecutively for 7 or more days in a row. Otherwise, users who read in more than one language tended to use the app sporadically with at least one day of no usage between days of usage.

Explore Feed

The Explore Feed (pictured in Fig. 1B) is what the user sees whenever they launch the app. It contains a variety of content such as historical events on this day, article of the day, and picture of the day. This content is presented in cards, the order of which can be customized in the feed customization screen (pictured in Fig. 1C). Some cards are available in multiple languages (if the user has multiple languages in their settings and that content is available in those languages), and those cards can be shown/hidden on a language-by-language basis. For example, a user who reads in Hindi and Marathi might wish to see the news from just Marathi Wikipedia and trending (most read) articles from just Hindi Wikipedia.

Customization

Based on the random sample of 23.155K users from the feed customization analytics, there were 0 users who customized their feed such that certain cards would be shown in some languages but not others. Users hide cards, but not for specific languages (where possible). Either users are not as aware of this feature as we would have hoped (more likely), or they really do want to see the cards in all the languages they have set (less likely). During onboarding – including an additional notice in their feed – we let users know that they can customize what is shown and how types of content are prioritized, but not that they can hide/show on a per-language basis.

Engagement

Note: while writing this report, we noticed several issues with feed analytics instrumentation – most notably with logging impressions – which will require further investigation and repairs. The statistics in Tables 4 and 5 are, therefore, estimates based on a much smaller subset of data than what is actually available. Since multilingual (3+) users are already not an especially large population of users, there were even fewer users who made it into the subset and so the clickthrough rates in Table 5 are 0% for almost all of the cards within this group.

In Table 4, we can see that bilingual users are more likely to engage with the Feed than monolingual users. Interestingly, multilingual (3+ languages) users are less likely to engage with cards which are

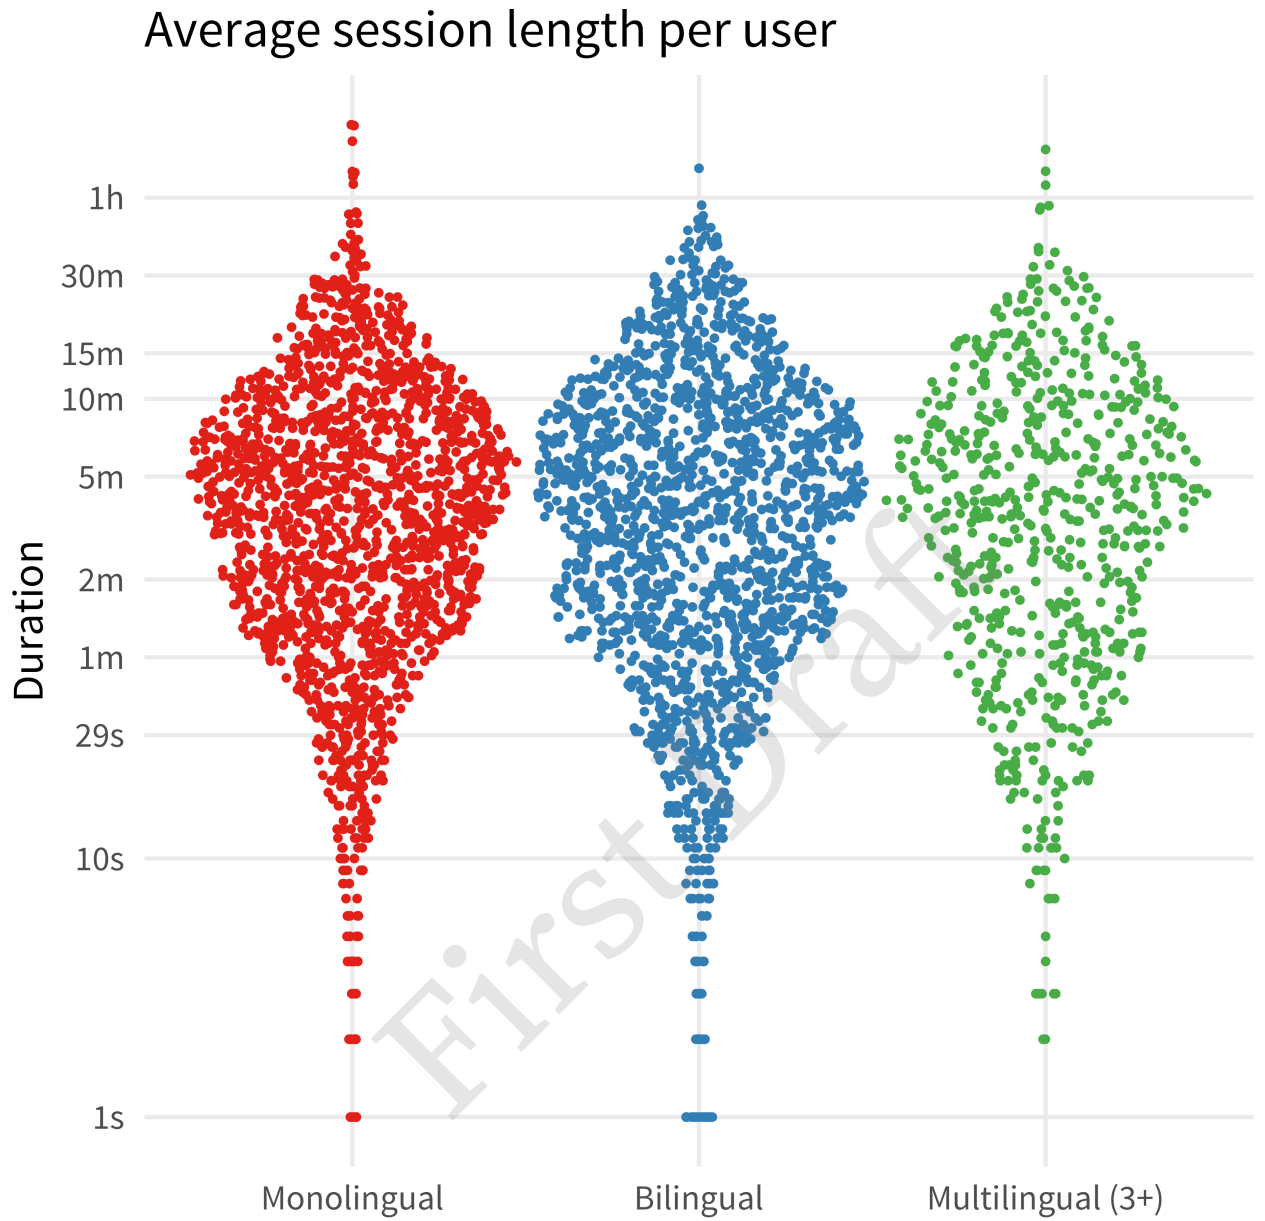
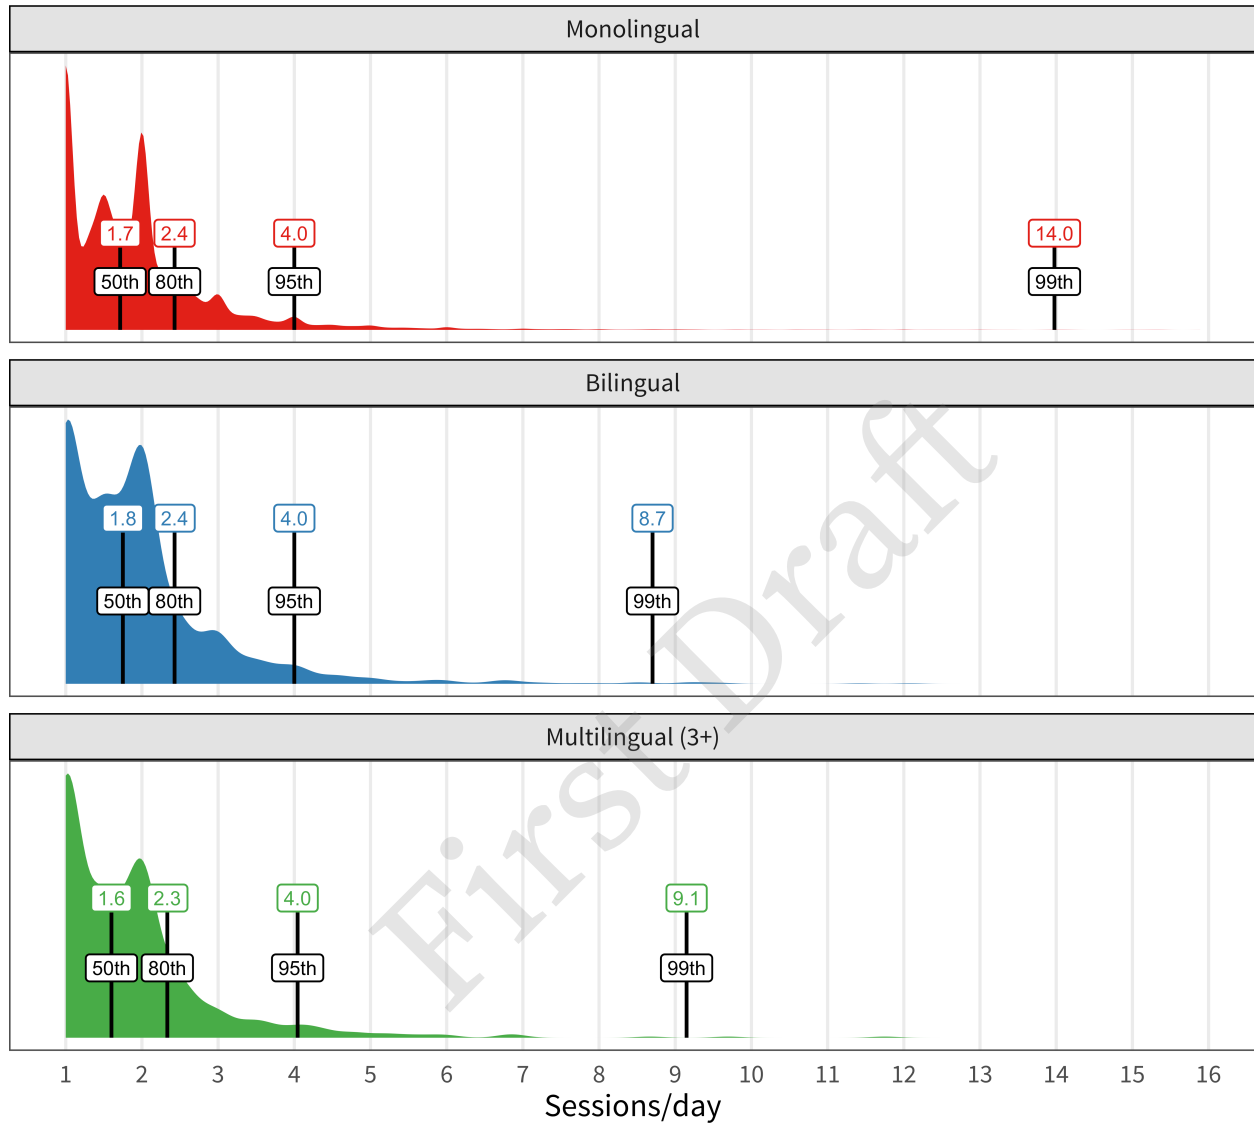


FIGURE 3: Distribution of average session length per user by language ability. For clarity, the monolingual group has been sub-sampled to the size of the bilingual group due to having significantly more users.

Average sessions per user per day



Language ability ■ Monolingual ■ Bilingual ■ Multilingual (3+)

FIGURE 4: Distribution of average number of sessions per day by language ability, with percentiles.

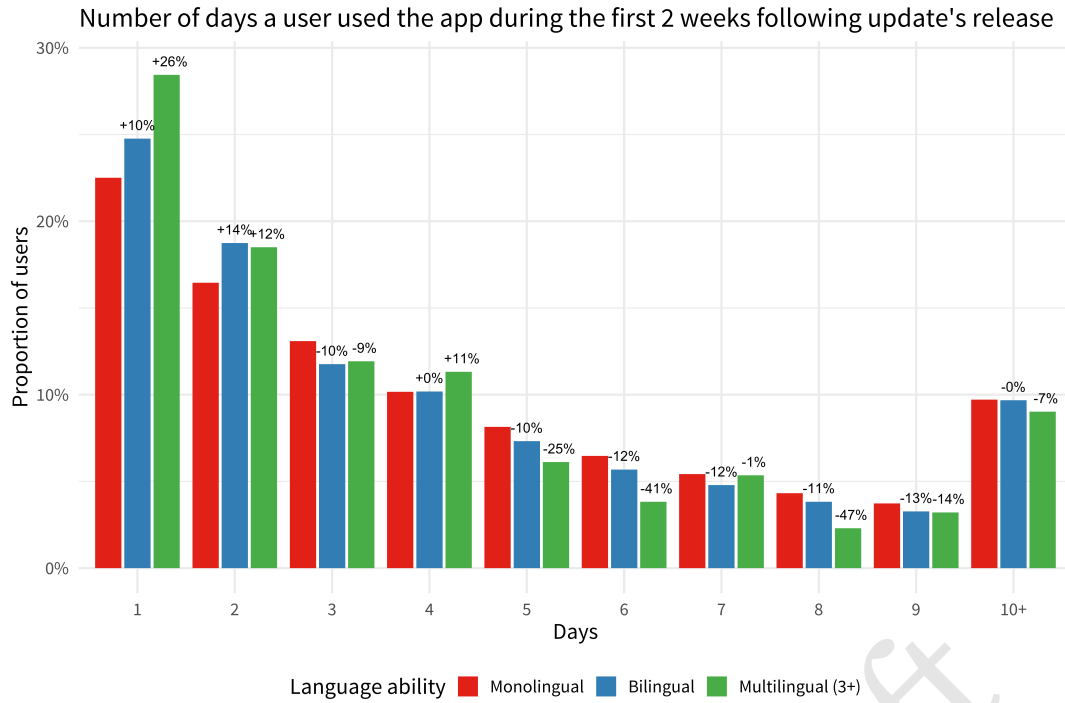


FIGURE 5: For each number of days, the +/- percentages are relative to proportion of monolingual users who, on average, used the app that many days.

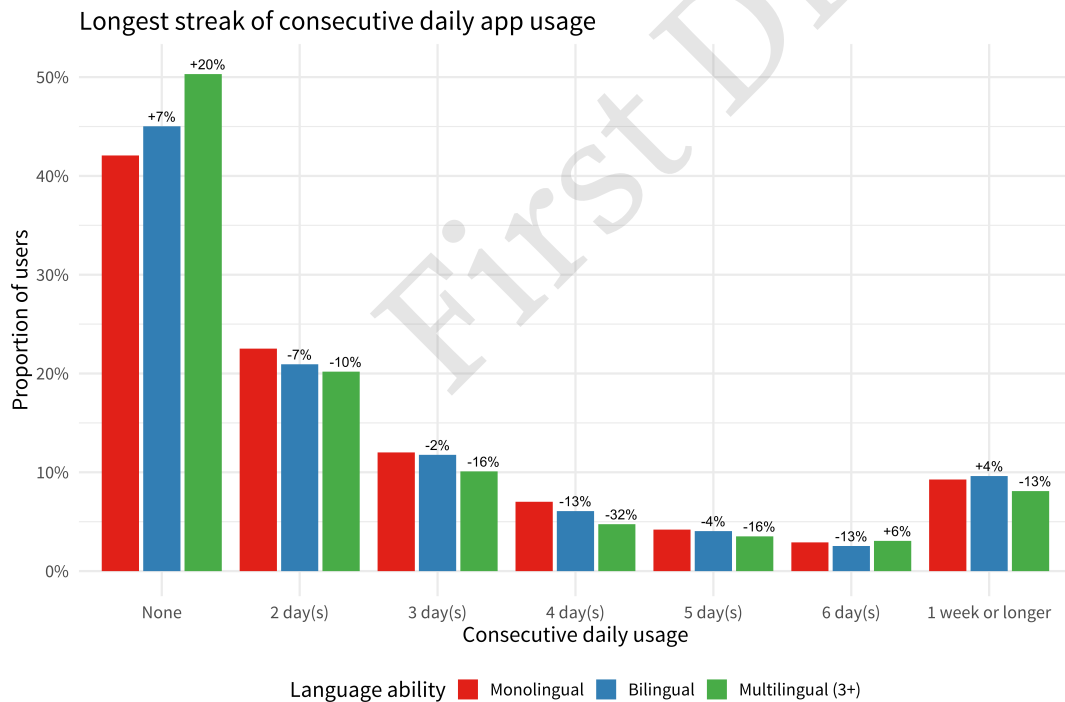


FIGURE 6: For each longest streak of consecutive daily usage, the +/- percentages are relative to proportion of monolingual users. For example, there were 4% more bilingual users than monolingual users whose longest streak of consecutive daily usage was 1 week or longer.

available in different languages but more likely to engage with cards which are language-agnostic or available in English only. Intuitively this makes sense in the context of the customization section above, where we showed that users do not customize the cards on a language-by-language basis. This means that users with 3 or more languages would have 3 (or more) versions of the same card, so they probably hide the card entirely or move it down, giving higher priority to other (language-agnostic and English-only) cards. It would be fascinating to explore this phenomenon in greater depth, but that endeavor is outside the scope of this report.

TABLE 4: Engagement with Feed cards by language ability and whether the cards are language-customizable (e.g. 'In the news') or not – including cards which are language-agnostic such as onboarding cards and cards which are only available in English, such as 'Because you read'.

Language ability	Multilingual-enabled cards CTR	Other (i.e. English-only) cards CTR
Monolingual	3.11%	13%
Bilingual	3.73%	15.1%
Multilingual (3+)	1.83%	18.1%

TABLE 5: Explore Feed engagement (clickthrough rates) by card type, multilingual support (whether the card is available in English only or in other languages), and user's language settings (1, 2, and 3+ languages). **Note:** there was a significant problem with logging events related to news cards, so that type was omitted.

Feed cards by multilingual support	Engagement among users who are		
	Monolingual	Bilingual	Multilingual (3+)
None (English only)			
Because You Read List	4.59%	5.26%	0%
Featured Image	7.82%	9.52%	0%
Main Page	6.17%	10.7%	16.7%
Onboarding - Customize Feed	41.7%	49.1%	22.4%
Onboarding - Reading List Sync	22.4%	15.5%	23.4%
Random	0%	0%	0%
Where available			
Featured Article	2%	3.7%	0%
Most Read List	9.66%	19.1%	12.1%
N/A (depends on primary)			
Continue Reading	3.98%	4.76%	4%

Table 5 breaks down each card type's clickthrough rates by language ability among a random sample of users. There are several notable statistics in this table:

- Nobody uses the randomizer.
- Engagement with featured article card is low across language abilities, although bilingual users show higher engagement since featured articles are shown in multiple languages so

there is a greater chance of seeing a highlighted article that they might be interested in reading.

- Unlike almost all other cards, bilingual users are less likely to click on the card which introduces them to [the reading list syncing feature](#) (read more about this feature on [the Wikimedia Blog](#)).

Searching behavior

Among users who read in more than one language, we estimate 18.6% switch languages when searching. We recently updated the search analytics to provide us with a better estimate of how many users switch languages when they have the option to, but this more reliable data was not available at the time of writing this report. For this reason we have chosen to not include search engagement metrics such as the clickthrough rate in this report.

Conclusion & Discussion

The multilingual update appears to be a useful one for many users, and majority of users who open the language settings end up having multiple languages set so presumably they appreciate not having to switch between languages anymore. While we did not see much of a difference in sessions (average session length and average number of sessions per day), we did see an increased number of articles read per session among users who read in more than one language.

Furthermore, with nearly 20% of multilingual users switching between languages when searching, we suspect having this ability makes searching more convenient. Unfortunately, we need to be better at onboarding users to Feed customization because nobody is customizing Feed cards on a language-by-language basis. We suspect this has a negative effect on the usefulness of the Feed to users who read in 3+ languages, as those users were less likely to interact with language-specific cards but were more likely to interact with language-agnostic and English-only cards.

The results of analyzing daily usage of the app suggest further investigation. With Figures 5 and 6 showing that users who read in one language tend to use the app more frequently than their multilingual counterparts, it seems like an interesting research project to find out why. After all, this seems counter-intuitive because one might expect that having more content available in multiple languages the user reads in would inspire more frequent usage of the app.

Acknowledgements

This report was created using the statistical software and programming language **R** (2018). Data collection relied on the excellent work of the [Mobile Apps](#) – especially the instrumentation work of Sharvani Haran – and data pipelines developed by [Analytics Engineering](#) teams.

References

- Allaire, J., Xie, Y., McPherson, J., Luraschi, J., Ushey, K., Atkins, A., ... Chang, W. (2018). Rmarkdown: Dynamic documents for r. Retrieved from <https://CRAN.R-project.org/package=rmarkdown>
- Bache, S. M. (2015). Import: An import mechanism for r. Retrieved from <https://CRAN.R-project.org/package=import>
- Bache, S. M., & Wickham, H. (2014). Magrittr: A forward-pipe operator for r. Retrieved from <https://CRAN.R-project.org/package=magrittr>
- Henry, L., & Wickham, H. (2018). Purrr: Functional programming tools. Retrieved from <https://CRAN.R-project.org/package=purrr>
- Hester, J. (2017). Glue: Interpreted string literals. Retrieved from <https://CRAN.R-project.org/package=glue>
- Ooms, J. (2014). The jsonlite package: A practical and consistent mapping between json data and r objects. arXiv:1403.2805 [stat.CO]. Retrieved from <https://arxiv.org/abs/1403.2805>
- Pedersen, T. L. (2018). Ggforce: Accelerating 'ggplot2'. Retrieved from <https://CRAN.R-project.org/package=ggforce>
- Popov, M., & Keyes, O. (2018). Wmf: R code for wikimedia foundation internal usage. Retrieved from <https://phabricator.wikimedia.org/diffusion/1821/>
- R Core Team. (2018). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>
- R Special Interest Group on Databases (R-SIG-DB), Wickham, H., & Müller, K. (2018). DBI: R database interface. Retrieved from <https://CRAN.R-project.org/package=DBI>
- Wickham, H. (2016). Ggplot2: Elegant graphics for data analysis. Springer-Verlag New York. Retrieved from <http://ggplot2.org>
- Wickham, H. (2018). Forcats: Tools for working with categorical variables (factors). Retrieved from <https://CRAN.R-project.org/package=forcats>
- Wickham, H., François, R., Henry, L., & Müller, K. (2018). Dplyr: A grammar of data manipulation. Retrieved from <https://CRAN.R-project.org/package=dplyr>
- Wickham, H., & Henry, L. (2018). Tidy: Easily tidy data with 'spread()' and 'gather()' functions. Retrieved from <https://CRAN.R-project.org/package=tidy>
- Xie, Y. (2014). Knitr: A comprehensive tool for reproducible research in R. In V. Stodden, F. Leisch, & R. D. Peng (Eds.), Implementing reproducible computational research. Chapman; Hall/CRC. Retrieved from <http://www.crcpress.com/product/isbn/9781466561595>
- Xie, Y. (2015). Dynamic documents with R and knitr (2nd ed.). Boca Raton, Florida: Chapman; Hall/CRC. Retrieved from <https://yihui.name/knitr/>
- Xie, Y. (2018). Knitr: A general-purpose package for dynamic report generation in r. Retrieved from <https://yihui.name/knitr/>
- Zhu, H. (2018). KableExtra: Construct complex table with 'kable' and pipe syntax. Retrieved from <https://CRAN.R-project.org/package=kableExtra>
- Zhu, H., Tsai, T., & Traviison, T. (2018). Memor: A 'rmarkdown' template that can be highly customized. Retrieved from <https://CRAN.R-project.org/package=memor>